2018 IEEE 59th Annual Symposium on Foundations of Computer Science

Near-Optimal Communication Lower Bounds for Approximate Nash Equilibria (Extended Abstract)

Mika Göös
Harvard University
Cambridge, MA, USA
mika@seas.harvard.edu

Aviad Rubinstein
Harvard University
Cambridge, MA, USA
aviad@cs.stanford.edu

***Abstract*—We prove an $N^{2-o(1)}$ lower bound on the randomized communication complexity of finding an ϵ-approximate Nash equilibrium (for constant $\epsilon > 0$) in a two-player $N \times N$ game.**

***Keywords*-Communication Complexity; Game Theory; Nash equilibrium.**

I. Introduction

How many bits of communication are needed to find an ϵ-approximate Nash equilibrium (for a small constant $\epsilon > 0$) in a two-player $N \times N$ game? More precisely:

- **Alice** holds her payoff matrix $A \in [0, 1]^{N\times N}$ of some constant precision.
- **Bob** holds his payoff matrix $B \in [0, 1]^{N\times N}$ of some constant precision.
- **Output** an *ϵ-approximate Nash equilibrium*: a mixed strategy $\mathcal{A}$ for Alice and a mixed strategy $\mathcal{B}$ for Bob such that neither player can unilaterally change their strategy and increase their expected payoff by more than ϵ. (See Section I-C for a formal definition.)

It is well known that such approximate equilibria have a concise $O(\log^2 N)$-bit description: one may assume wlog that $\mathcal{A}$ and $\mathcal{B}$ are supported on at most $O(\log N)$ actions [1]. There is a trivial upper bound of $O(N^2)$ by communicating an entire payoff matrix. Previous work [2] showed that finding an ϵ-approximate Nash equilibrium requires N^δ bits of communication for a small constant $\delta > 0$. In this work, we improve this to a near-optimal $N^{2-o(1)}$ lower bound. Our main theorem is slightly more general, as it also applies for games of asymmetric dimensions.

Theorem 1. *There exists an $\epsilon > 0$ such that for any constants $a, b > 0$ the randomized communication complexity of finding an ϵ-approximate Nash equilibrium in an $N^a \times N^b$ game is $N^{a+b-o(1)}$.*

It is interesting to note that there is an $\tilde{O}(N)$ protocol for computing an ϵ-correlated equilibrium of an $N \times N$ game [3]. Hence our result is the first that separates approximate Nash and correlated equilibrium.

Our result also implies the first near-quadratic lower bound for finding an approximate Nash equilibrium in the weaker query complexity model, where the algorithm has black-box oracle access to the payoff matrices (previous work established such lower bounds only against deterministic algorithms [4]). In this query complexity model, there is an $\tilde{O}(N)$-queries algorithm for computing an ϵ-coarse correlated equilibrium of an $N \times N$ game [5]. Hence our result is the first that separates approximate Nash and coarse correlated equilibrium in the query complexity model. See Table I for a summary of known bounds[1].

A. Background

Nash equilibrium is the central solution concept in game theory. It is named after John Nash who, more than 60 years ago, proved that every game has an equilibrium [6]. Once players are at an equilibrium, they do not have an incentive to deviate. However, Nash's theorem does not explain how the players arrive at an equilibrium in the first place.

Over the last several decades, many *dynamics*, or procedures by which players in a repeated game update their respective strategies to adapt to other players' strategies, have been proposed since Nash's result (e.g., [7], [8], [9], [10], [11]). But despite significant effort, we do not know any plausible dynamics that converge even to an approximate Nash equilibrium. It is thus natural to conjecture that there are no such dynamics. However, one has to be careful about defining "plausible" dynamics. The first example of dynamics we consider *im*plausible, are "players agree a priori on a Nash equilibrium". The *uncoupled dynamics* model proposed by Hart and Mas-Collel [10] rules out such trivialities by requiring that a player's strategy depends only on her own utility function and the history of other players' actions. Another example of implausible dynamics that converge to a Nash equilibrium are exhaustive search dynamics that enumerate over the entire search space. (Exhaustive search dynamics can converge to an approximate Nash equilibrium in finite time by enumerating over an ϵ-net of the search space.) We thus consider a second natural desideratum, which is that dynamics should converge

[1]We thank Yakov Babichenko for his help in understanding these connections and other insightful communication.

2575-8454/18/$31.00 ©2018 IEEE
DOI 10.1109/FOCS.2018.00045

(much) faster than exhaustive search. Note that the two restrictions (uncoupled-ness and fast convergence) are still very minimal—it is still easy to come up with dynamics that satisfy both and yet do not plausibly expect to predict players' behavior. But, since we are after an impossibility result, it is fair to say that if we can rule out any dynamics that satisfy these two restrictions and converge to a Nash equilibrium, we have strong evidence against any plausible dynamics.

A beautiful observation by Conitzer and Sandholm [12] and Hart and Mansour [13] is that the *communication complexity* of computing an (approximate) Nash equilibrium, in the natural setting where each player knows her own utility function, precisely captures (up to a logarithmic factor) the number of rounds for an arbitrary uncoupled dynamics to converge to an (approximate) Nash equilibrium. Thus the question of ruling out plausible dynamics is reduced to the question of proving lower bounds on communication complexity. There are also other good reasons to study the communication complexity of approximate Nash equilibria; see e.g. [14].

B. Related work

The problem of computing (approximate) Nash equilibrium has been studied extensively, mostly in three models: communication complexity, query complexity, and computational complexity.

Communication complexity.: The study of the communication complexity of Nash equilibria was initiated by Conitzer and Sandholm [12] who proved a quadratic lower on the communication complexity of deciding whether a game has a pure equilibrium, even for zero-one payoff (note that for pure equilibrium this also rules out any more efficient approximation). Hart and Mansour [13] proved exponential lower bounds for pure and exact mixed Nash equilibrium in n-player games. Roughgarden and Weinstein [15] proved communication complexity lower bounds on the related problem of finding an approximate Brouwer fixed point (on a grid). In [2], in addition to the lower bound for two-player game which we improved, there is also an exponential lower bound for n-player games. The same paper also posed the open problem of settling the communication complexity of approximate correlated equilibrium in two-player games; partial progress has been made by [16], [17], but to date the problem of determining the communication complexity of ϵ-approximate correlated equilibrium remains open (even at the granularity of $\mathrm{poly}(N)$ vs $\mathrm{polylog}(N)$). Babichenko et al. [18] prove communication complexity lower bounds for finding a pure equilibrium in the restricted class of potential games. On the algorithmic side, Czumaj et al. [19] gave a polylogarithmic protocol for computing a 0.382-approximate Nash equilibrium in two-player games, improving upon [20].

Query complexity.: In the query complexity model, the algorithm has black-box oracle access to the payoff matrix of each player. Notice that this model is strictly weaker than the communication complexity model (hence our communication lower bound applies to this model as well). For the *deterministic* query complexity of ϵ-approximate Nash equilibrium in two-player games, Fearnley et al. [21] proved a linear lower bound, which was subsequently improved to (tight) quadratic by Fearnley and Savani [4]. For randomized algorithms, the only previous lower bound was $\Omega(1/\epsilon^2)$, also by [4]; notice that this is only interesting for $\epsilon = o(1)$. As mentioned earlier, Goldberg and Roth [5] can find an approximate coarse correlated equilibrium with $\tilde{O}(N)$ queries (two-player $N \times N$ game) or $\tilde{O}(n)$ queries (n-player game with two actions per player). In the latter regime of n-player and two action per player, a long sequence of works [22], [23], [24], [25] eventually established an $2^{\Omega(n)}$ lower bound on the query complexity of ϵ-approximate Nash equilibrium. (This last result is implied by and was the starting point for the exponential query complexity lower bound in [2]). Finally, some of the aforementioned results are inspired by a query complexity lower bound for approximate fixed point due to Hirsch et al. [26] and its adaptation to ℓ_2-norm in [25]; this construction will also be the starting point of our reduction.

Computational complexity.: For computational complexity, the problem of finding an exact Nash equilibrium in a two-player game is PPAD-complete [27], [28]. Following a sequence of improvements [29], [30], [31], [32], [33], we know that a 0.3393-approximate Nash equilibrium can be computed in polynomial time. But there exists some constant $\epsilon > 0$, such that assuming the "Exponential Time Hypothesis (ETH) for PPAD", computing an ϵ-approximate Nash equilibrium requires $N^{\log^{1-o(1)} N}$ time [25], which is essentially tight by [1].

C. Definition of ϵ-Nash

A two-player game is defined by two utility functions (or payoff matrices) $U^A, U^B \colon S_A \times S_B \to [0,1]$. A mixed strategy for Alice (resp. Bob) is a distribution $\mathcal{A}$ ($\mathcal{B}$) over the set of actions S_A (S_B). We say that $(\mathcal{A}, \mathcal{B})$ is an ϵ-*approximate Nash equilibrium* (ϵ-ANE) if every alternative Alice-strategy $\mathcal{A}'$ performs at most ϵ better than $\mathcal{A}$ against Bob's mixed strategy $\mathcal{B}$, and the same holds with roles reversed. Formally, the condition for Alice is

$$\mathop{\mathbb{E}}_{\substack{a\sim\mathcal{A}\\ b\sim\mathcal{B}}}\left[U^A(a,b)\right] \;\geq\; \max_{\mathcal{A}' \text{ over } S_A} \; \mathop{\mathbb{E}}_{\substack{a'\sim\mathcal{A}'\\ b\sim\mathcal{B}}}\left[U^A(a',b)\right] - \epsilon.$$

II. Technical Overview

Our proof follows the high-level approach of [2]; see also the lecture notes [34] for exposition. The approach of [2] consists of four steps.

1) Query complexity lower bound for the well-known PPAD-complete EoL problem.

Type of equilibrium	Query Complexity		Communication Complexity	
ϵ-Nash equilibirum	$\Omega(N^{2-o(1)})$	[This paper]	$\Omega(N^{2-o(1)})$	[This paper]
ϵ-correlated equilibrium	$\Omega(N)$	[Folklore]	$O(N \log N)$	[3]
ϵ-coarse correlated equilibrium	$\tilde{\Theta}(N)$	[5]	$\mathrm{polylog}(N)$	[Folklore]

Table I: Query and communication complexities of approximate equilibria. The oracle query model is more restrictive than the communication model. Nash equilibrium is more restrictive than correlated equilibrium, which is yet more restrictive than coarse correlated equilibrium. Hence the complexity of a problem increases as we move up and left in the table. For all problems there are trivial bounds of $\Omega(\log N)$ and $O(N^2)$.

2) Lifting of the above into a communication lower bound for a two-party version of EoL.
3) Reduction from EoL to ϵ-BFP, a problem of finding an (approximate) Brouwer fixed point.
4) Constructing a hard two-player game that combines problems from both Step 2 and Step 3.

In this paper we improve the result from [2] by optimizing Steps 1, 2, and 4. We outline these improvements below; see the full version [35] for proofs and details.

A. Steps 1–2: Lower bound for End-of-Line

The goal of Steps 1–2 is to obtain a randomized communication lower bound for the End-of-Line (or EoL for short) problem: Given an implicitly described graph on $[N]$ where a special vertex $1 \in [N]$ is the start vertex of a path, find an end of a path or a non-special start of a path. The following definition is a "template" in that it does not yet specify the protocols Π_v.

> **EoL template**
> - **Input:** Alice and Bob receive inputs α and β that implicitly describe *successor* and *predecessor* functions $S, P\colon [N] \to [N]$. Namely, for each $v \in [N]$ there is a "low-cost" protocol $\Pi_v(\alpha, \beta)$ to compute the pair $(S(v), P(v))$.
> - **Output:** Define a digraph $G = ([N], E)$ where $(v, u) \in E$ iff $S(v) = u$ and $P(u) = v$. The goal is to output a vertex $v \in [N]$ such that either
>   - $v = 1$ and v is a non-source or a sink in G; or
>   - $v \neq 1$ and v is a source or a sink in G.

The prior work [2] proved an $\tilde{\Omega}(N^{1/2})$ lower bound for a version of EoL where the Π_v had communication cost $c := \Theta(\log N)$. The cost parameter c is, surprisingly, very important: later reductions (in Step 4) will incur a blow-up in input size—and hence a quantitative reduction in the eventual lower bound—that is *exponential* in c. (Namely, when constructing payoff matrices in Step 4, the data defining a strategy for Alice will include a c-bit transcript of some Π_v.)

In this work, we obtain an optimized lower bound:

Theorem 2. *There is a family of* EoL *instances, whose randomized communication complexity is $\tilde{\Omega}(N)$, and where the Π_v have constant cost $c = O(1)$.*

Note: Since we consider $c = O(1)$, a c-bit transcript of an Π_v cannot even name arbitrary $\log(N)$-bit vertices in $[N]$. Thus we need to clarify what it means for Π_v to "compute" $(S(v), P(v))$. The formal requirement is that the pair $(S(v), P(v))$ is some prescribed function of both v and the c-bit transcript $\Pi_v(\alpha, \beta)$. Concretely, we will fix some *bounded-degree host graph* $H = ([N], E)$ independent of (α, β), and define graphs G as subgraphs of H. In other words, for any vertex v there is a fixed short list of potential predecessor and successor vertices. Now we can, for example, let Π_v announce $S(v)$ as "the i-th out-neighbor of v in H", which takes only $O(1)$ bits to represent.

As in [2], our lower bound is obtained by first proving an analogous result for *query complexity*, and then applying a *lifting theorem* that escalates the query hardness into communication hardness. A key difference is that instead of a generic lifting theorem [36], [37], as used by [2], we employ a less generic, but quantitatively better one [38], [39].

Step 1: Query lower bound.: The query complexity analogue of EoL is defined as follows.

> **Q-EoL$_H$ for host digraph $H = ([N], E)$**
> - **Input:** An input $x \in \{0,1\}^E$ describes a (spanning) subgraph $G = G_x$ of H consisting of the edges e such that $x_e = 1$.
> - **Output:** Find a vertex $v \in [N]$ such that either
>   - $v = 1$ and in-deg$(v) \neq 0$ or out-deg$(v) \neq 1$ in G; or
>   - $v \neq 1$ and in-deg$(v) \neq 1$ or out-deg$(v) \neq 1$ in G.

We exhibit a bounded-degree host graph H such that any randomized decision tree needs to make $\tilde{\Omega}(N)$ queries to the input x in order to solve Q-EoL$_H$. Moreover, the lower bound is proved using *critical block sensitivity* (cbs), a measure introduced by Huynh and Nordström [38] that lower bounds randomized query complexity (among other things); see full version [35] for definitions.

Lemma 3. *There is a bounded-degree host graph* $H = ([N], E)$ *such that* $\mathrm{cbs}(\text{Q-EoL}_H) \geq \tilde{\Omega}(N)$.

It is not hard to prove an $\Omega(N)$ bound for a *complete* host graph (equipped with successor/predecessor pointers), nor an $\Omega(N^{1/2})$ bound for a bounded-degree host graph (by reducing degrees in the complete graph via binary trees). But to achieve both an $\tilde{\Omega}(N)$ bound and constant degree requires a careful choice of a host graph that has good enough routing properties. Our construction uses *butterfly* graphs.

Prior to this work, a near-linear randomized query lower bound was known for a bounded-degree *Tseitin* problem [39], a canonical PPA-complete search problem. Since PPAD $\subseteq$ PPA, our new lower bound is qualitatively stronger (also, the proof is more involved).

Step 2: Communication lower bound.: Let $R \subseteq \{0,1\}^N \times \mathcal{O}$ be a query search problem (e.g., $R =$ Q-EoL$_H$), that is, on input $x \in \{0,1\}^N$ the goal is to output some $o \in \mathcal{O}$ such that $(x, o) \in R$. Any such R can be converted into a communication problem via *gadget composition*. Namely, fix some two-party function $g \colon \Sigma \times \Sigma \to \{0,1\}$, called a *gadget*. The composed search problem $R \circ g$ is defined as follows: Alice holds $\alpha \in \Sigma^N$, Bob holds $\beta \in \Sigma^N$, and their goal is to find an $o \in \mathcal{O}$ such that $(x, o) \in R$ where

$$x \;:=\; g^N(\alpha, \beta) \;=\; (g(\alpha_1, \beta_1), \ldots, g(\alpha_N, \beta_N)).$$

It is generally conjectured that the randomized communication complexity of $R \circ g$ is characterized by the randomized query complexity of R, provided the gadget g is chosen carefully. This was proved in [37], but only for a *non-constant-size* gadget where Alice's input is $\Theta(\log N)$ bits. This is prohibitively large for us, since we seek protocols Π_v of *constant* communication cost. We use instead a more restricted lifting theorem due to [39] (building on [38]) that works for a constant-size gadget, but can only lift critical block sensitivity bounds.

Lemma 4 ([39]). *There is a fixed gadget* $g \colon \Sigma \times \Sigma \to \{0,1\}$ *such that for any* $R \subseteq \{0,1\}^N \times \mathcal{O}$ *the randomized communication complexity of* $R \circ g$ *is at least* $\Omega(\mathrm{cbs}(R))$.

Theorem 2 now follows by combining Lemma 3 and Lemma 4. We need only verify that the composed problem Q-EoL $\circ\, g$ fits our EoL template. For $v \in [N]$ consider the protocol Π_v that computes as follows on input $(\alpha, \beta) \in \Sigma^E \times \Sigma^E$:

1) Alice sends all symbols $\alpha_e \in \Sigma$ for e incident to v.
2) Bob privately computes all values $x_e = g(\alpha_e, \beta_e)$ for e incident to v.
3) Bob announces $S(v)$ as the first out-neighbor of v in the subgraph determined by x if such an out-neighbor exists; otherwise Bob announces $S(v) := v$. Similarly for $P(v)$.

This protocol has indeed cost $c = O(1)$ because H is of bounded degree and $|\Sigma|$ is constant.

B. Step 3: Reduction to ϵ-BFP

By Brouwer's fixed point theorem, any continuous function $f \colon [0,1]^m \to [0,1]^m$ has a fixed point, that is, x^* such that $f(x^*) = x^*$. The BFP query problem is to find such a fixed point, given oracle access to f. We will consider the easier ϵ-BFP problem, where we merely have to find an x such $f(x)$ is ϵ-close to x.

A theorem of [25] reduces Q-EoL to ϵ-BFP with $m = O(\log(N))$. For our purposes, there are two downsides to using this theorem. First, it is a reduction between query complexity problems, which seems to undermine the lifting to communication we obtained in Step 2. (This obstacle was already encountered in [15] and resolved in [2].)

The second issue with [25]'s reduction is that it blows up the search space. We can discretize $[0,1]$ to obtain a finite search space. But even if the discretization used one bit per coordinate (and in fact we need a large constant number of bits), the dimension m is still larger than $\log_2 N$ by yet another constant factor due to the seemingly-unavoidable use of error correcting codes. All in all we have a polynomial blow-up in the size of the search space, and while that was a non-issue for [25], [2], it is crucial for our fine-grained result.

Our approach for both obstacles is to postpone dealing with them to Step 4. But for all the magic to happen in Step 4, we need to properly set up some infrastructure before we conclude Step 3. Concretely, without changing the construction of f from [25], we observe that it can be computed in a way that is "local" in two different ways (we henceforth say that f is *doubly-local*). Below is an informal description of what this means; see full version [35] for details.

- First, every point $x \in [0,1]^m$ corresponds to a vertex v from the host graph of the Q-EoL problem[2]. We observe that in order to compute $f(x)$, one only needs local access to the neighborhood of v of the Q-EoL (actual, not host) graph. A similar sense of locality was used in [2].
- Second, if we only want to compute the i-th coordinate of $f(x)$, we do not even need to know the entire vector x. Rather, it suffices to know x_i, the values

[2]In fact, each x corresponds to zero, one, or two vertices from the host graph, where the two vertices are either identical or neighbors. For simplicity, in this informal discussion we refer to "the corresponding vertex".

of x on a random subset of the coordinates, and the local information of the Q-EoL graph described in the previous bullet (including v). This is somewhat reminiscent of the local decoding used in [25] (but our locality is much simpler and does not require any PCP machinery).

Theorem 5 (Q-EoL to ϵ-BFP, informal [25])**.** *There is a reduction from* Q-EoL *over N vertices to* ϵ-BFP *on a function $f\colon [0,1]^m \to [0,1]^m$, where f is "doubly-local".*

C. Step 4: Reduction to ϵ-Nash

The existence of a Nash equilibrium is typically proved using Brouwer's fixed point theorem. McLennan and Tourky [40] proved the other direction, namely that the existence of a Nash equilibrium in a special *imitation game* implies an existence of a fixed point. Viewed as a reduction from Brouwer fixed point to Nash equilibrium, it turns out to be (roughly) approximation-preserving, and thus extremely useful in recent advances on hardness of approximation of Nash equilibrium in query complexity [41], [24], [25], computational complexity [42], [25], and communication complexity [15], [2].

In the basic imitation game, we think of Alice's and Bob's action space as $[0,1]^m$, and define their utility functions as follows. First, Alice chooses $x^{(\mathbf{a})} \in [0,1]^m$ that should imitate the $x^{(\mathbf{b})} \in [0,1]^m$ chosen by Bob:

$$U^A\left(x^{(\mathbf{a})}; x^{(\mathbf{b})}\right) := -\left\|x^{(\mathbf{a})} - x^{(\mathbf{b})}\right\|_2^2.$$

Notice that Alice's expected utility decomposes as

$$\mathop{\mathbb{E}}_{x^{(\mathbf{b})}}\left[U^A\left(x^{(\mathbf{a})}; x^{(\mathbf{b})}\right)\right] = -\left\|x^{(\mathbf{a})} - \mathbb{E}\left[x^{(\mathbf{b})}\right]\right\|_2^2 - \mathrm{Var}\left[x^{(\mathbf{b})}\right],$$

where the second term does not depend on Alice's action at all. This significantly simplifies the analysis because we do not need to think about Bob's mixed strategy: in expectation, Alice just tries to get as close as possible to $\mathbb{E}\left[x^{(\mathbf{b})}\right]$. Similarly, Bob's utility function is defined as:

$$U^B\left(x^{(\mathbf{b})}; x^{(\mathbf{a})}\right) := -\left\|f\left(x^{(\mathbf{a})}\right) - x^{(\mathbf{b})}\right\|_2^2.$$

It is not hard to see that in every Nash equilibrium of the game, $x^{(\mathbf{a})} = x^{(\mathbf{b})} = f\left(x^{(\mathbf{a})}\right)$.

For our reduction, we need to make some modifications to the above imitation game. First, observe that Bob's utility must not encode the entire function f—otherwise Bob could find the fixed point (or Nash equilibrium) with zero communication from Alice! Instead, we ask that Alice's action specifies a vertex $v^{(\mathbf{a})}$, as well as her inputs to the lifting gadgets associated to (edges adjacent to) $v^{(\mathbf{a})}$. If $v^{(\mathbf{a})}$ is indeed the vertex corresponding to $x^{(\mathbf{a})}$, Bob can use his own inputs to the lifting gadgets to locally compute $f\left(x^{(\mathbf{a})}\right)$ (this corresponds to the first type of "local").

The second issue is that for our fine-grained reduction, we cannot afford to let Alice's and Bob's actions specify an entire point $x \in [0,1]^m$. Instead, we force the equilibria of the game to be strictly mixed, where each player chooses a small (pseudo-)random subset of coordinates $[m]$. Then, each player's *mixed strategy* represents $x \in [0,1]^m$, but each action only specifies its restriction to the corresponding subset of coordinates. By the second type of "local", Bob can locally compute the value of $f(x)$ on the intersection of subsets. Inconveniently, the switch to mixed strategies significantly complicates the analysis: we have to make sure that Alice's mixed strategy is consistent with a single $x \in [0,1]^m$, deal with the fact that in any approximate equilibrium she is only approximately randomizing her selection of subset, etc.

Finally, the ideas above can be combined to give an $N^{1-o(1)}$ lower bound on the communication complexity (already much stronger than previous work). The bottleneck to improving further is that while we are able to distribute the vector x across the support of Alice's mixed strategy, we cannot do the same with the corresponding vertex v from the EoL graph. The reason is that given just a single action of Alice (not her mixed strategy), Bob must be able to compute his own utility; for that he needs to locally compute $f(x)$ (on some coordinates); and even with the doubly-local property of f, that still requires knowing the entire v. Finally, even with the most succinct encoding, if Alice's action represents an arbitrary vertex, she needs at least N actions. To improve to the desired $N^{2-o(1)}$ lower bound, we observe that when Bob locally computes his utility he does have another input: his own action. We thus split the encoding of v between Alice's action and Bob's action, enabling us to use an EoL host graph over N^2 vertices. (More generally, for an asymmetric $N^a \times N^b$ game we can split the encoding unevenly.)

Acknowledgment

Both authors were supported by Rabin Postdoctoral Fellowships.

References

[1] R. Lipton, E. Markakis, and A. Mehta, "Playing large games using simple strategies," in *Proceedings of the 4th Conference on Electronic Commerce (EC)*. ACM, 2003, pp. 36–41.

[2] Y. Babichenko and A. Rubinstein, "Communication complexity of approximate Nash equilibria," in *Proceedings of the 49th Symposium on Theory of Computing (STOC)*. ACM, 2017, pp. 878–889.

[3] A. Blum and Y. Mansour, "From external to internal regret," *Journal of Machine Learning Research*, vol. 8, pp. 1307–1324, 2007. [Online]. Available: http://www.jmlr.org/papers/v8/blum07a.html

[4] J. Fearnley and R. Savani, "Finding approximate Nash equilibria of bimatrix games via payoff queries," *Transactions on Economics and Computation*, vol. 4, no. 4, pp. 25:1–25:19, 2016.

[5] P. Goldberg and A. Roth, "Bounds for the query complexity of approximate equilibria," *Transactions on Economics and Computation*, vol. 4, no. 4, pp. 24:1–24:25, 2016.

[6] J. Nash, "Non-cooperative games," *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, 1951.

[7] G. Brown, "Iterative solutions of games by fictitious play," *Activity Analysis of Production and Allocation*, pp. 374–376, 1951.

[8] J. Robinson, "An iterative method of solving a game," *Annals of Mathematics*, vol. 54, no. 2, pp. 296–301, 1951.

[9] E. Kalai and E. Lehrer, "Rational learning leads to Nash equilibrium," *Econometrica*, vol. 61, no. 5, pp. 1019–1045, 1993.

[10] S. Hart and A. Mas-Colell, "Uncoupled dynamics do not lead to Nash equilibrium," *American Economic Review*, vol. 93, no. 5, pp. 1830–1836, 2003.

[11] D. Foster and P. Young, "Regret testing: learning to play Nash equilibrium without knowing you have an opponent," *Theoretical Economics*, vol. 1, no. 3, pp. 341–367, 2006. [Online]. Available: http://econtheory.org/ojs/index.php/te/article/view/20060341

[12] V. Conitzer and T. Sandholm, "Communication complexity as a lower bound for learning in games," in *Proceedings of the 21st International Conference on Machine Learning (ICML)*. ACM, 2004, p. 24.

[13] S. Hart and Y. Mansour, "How long to equilibrium? The communication complexity of uncoupled equilibrium procedures," *Games and Economic Behavior*, vol. 69, no. 1, pp. 107–126, 2010.

[14] T. Roughgarden, "Barriers to near-optimal equilibria," in *Proceedings of the 55th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2014, pp. 71–80.

[15] T. Roughgarden and O. Weinstein, "On the communication complexity of approximate fixed points," in *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2016, pp. 229–238.

[16] A. Ganor and Karthik C. S., "Communication complexity of correlated equilibrium in two-player games," arXiv, Tech. Rep., 2017.

[17] Y. K. Ko and A. Schvartzman, "Bounds for the communication complexity of two-player approximate correlated equilibria," Electronic Colloquium on Computational Complexity (ECCC), Tech. Rep. TR17-071, 2017. [Online]. Available: https://eccc.weizmann.ac.il/report/2017/071

[18] Y. Babichenko, S. Dobzinski, and N. Nisan, "The communication complexity of local search," *CoRR*, vol. abs/1804.02676, 2018. [Online]. Available: http://arxiv.org/abs/1804.02676

[19] A. Czumaj, A. Deligkas, M. Fasoulakis, J. Fearnley, M. Jurdziński, and R. Savani, "Distributed methods for computing approximate equilibria," in *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, 2016, pp. 15–28.

[20] P. Goldberg and A. Pastink, "On the communication complexity of approximate Nash equilibria," *Games and Economic Behavior*, vol. 85, pp. 19–31, 2014.

[21] J. Fearnley, M. Gairing, P. Goldberg, and R. Savani, "Learning equilibria of games via payoff queries," *Journal of Machine Learning Research*, vol. 16, pp. 1305–1344, 2015. [Online]. Available: http://jmlr.org/papers/v16/fearnley15a.html

[22] S. Hart and N. Nisan, "The query complexity of correlated equilibria," *Games and Economic Behavior*, 2016.

[23] Y. Babichenko, "Completely uncoupled dynamics and Nash equilibria," *Games and Economic Behavior*, vol. 76, no. 1, pp. 1–14, 2012.

[24] X. Chen, Y. Cheng, and B. Tang, "Well-supported vs. approximate Nash equilibria: Query complexity of large games," in *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, 2017, pp. 57:1–57:9.

[25] A. Rubinstein, "Settling the complexity of computing approximate two-player Nash equilibria," in *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2016, pp. 258–265.

[26] M. Hirsch, C. Papadimitriou, and S. Vavasis, "Exponential lower bounds for finding Brouwer fix points," *Journal of Complexity*, vol. 5, no. 4, pp. 379–416, 1989.

[27] C. Daskalakis, P. Goldberg, and C. Papadimitriou, "The complexity of computing a Nash equilibrium," *SIAM Journal on Computing*, vol. 39, no. 1, pp. 195–259, 2009.

[28] X. Chen, X. Deng, and S.-H. Teng, "Settling the complexity of computing two-player Nash equilibria," *Journal of the ACM*, vol. 56, no. 3, pp. 14:1–14:57, 2009.

[29] S. Kontogiannis, P. Panagopoulou, and P. Spirakis, "Polynomial algorithms for approximating Nash equilibria of bimatrix games," *Theoretical Computer Science*, vol. 410, no. 17, pp. 1599–1606, 2009.

[30] C. Daskalakis, A. Mehta, and C. Papadimitriou, "A note on approximate Nash equilibria," *Theoretical Computer Science*, vol. 410, no. 17, pp. 1581–1588, 2009.

[31] ——, "Progress in approximate Nash equilibria," in *Proceedings of the 8th Conference on Electronic Commerce (EC)*. ACM, 2007, pp. 355–358.

[32] H. Bosse, J. Byrka, and E. Markakis, "New algorithms for approximate Nash equilibria in bimatrix games," *Theoretical Computer Science*, vol. 411, no. 1, pp. 164–173, 2010.

[33] H. Tsaknakis and P. Spirakis, "An optimization approach for approximate Nash equilibria," *Internet Mathematics*, vol. 5, no. 4, pp. 365–382, 2008.

[34] T. Roughgarden, "Complexity theory, game theory, and economics," arXiv, Tech. Rep., 2018.

[35] M. Göös and A. Rubinstein, "Near-optimal communication lower bounds for approximate nash equilibria," *CoRR*, vol. abs/1805.06387, 2018. [Online]. Available: http://arxiv.org/abs/1805.06387

[36] M. Göös, S. Lovett, R. Meka, T. Watson, and D. Zuckerman, "Rectangles are nonnegative juntas," *SIAM Journal on Computing*, vol. 45, no. 5, pp. 1835–1869, 2016.

[37] M. Göös, T. Pitassi, and T. Watson, "Query-to-communication lifting for BPP," in *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2017, pp. 132–143.

[38] T. Huynh and J. Nordström, "On the virtue of succinct proofs: Amplifying communication complexity hardness to time–space trade-offs in proof complexity," in *Proceedings of the 44th Symposium on Theory of Computing (STOC)*. ACM, 2012, pp. 233–248.

[39] M. Göös and T. Pitassi, "Communication lower bounds via critical block sensitivity," in *Proceedings of the 46th Symposium on Theory of Computing (STOC)*. ACM, 2014, pp. 847–856.

[40] A. McLennan and R. Tourky, "From imitation games to Kakutani," University of Melbourne, Tech. Rep., 2005. [Online]. Available: http://cupid.economics.uq.edu.au/mclennan/Papers/kakutani60.pdf

[41] Y. Babichenko, "Query complexity of approximate Nash equilibria," *Journal of the ACM*, vol. 63, no. 4, pp. 36:1–36:24, 2016.

[42] A. Rubinstein, "Inapproximability of Nash equilibrium," in *Proceedings of the 47th Symposium on Theory of Computing (STOC)*. ACM, 2015, pp. 409–418.