Informational Bounds on Equilibria (a Survey)

YAKOV BABICHENKO
Technion

Query complexity and communication complexity of equilibria have been actively studied in the past decade. Recent progress in these fields of informational complexity has led to a quite good understanding of equilibria. This survey summarizes the established results for the three most common solution concepts: Nash equilibria, correlated equilibria, and coarse correlated equilibria. The survey provides a high-level idea of the techniques that are utilized to deduce recently developed lower bounds on Nash equilibria.

Categories and Subject Descriptors: Theory of computation [**Theory and algorithms for application domains**]: Algorithmic game theory and mechanism design—*Exact and approximate computation of equilibria, Convergence and learning in games*

General Terms: Complexity of Equilibria

Additional Key Words and Phrases: Communication complexity, Query complexity, Nash equilibrium, Correlated equilibrium, Coarse correlated equilibrium, Potential games

1. INTRODUCTION

Query complexity and communication complexity are two well-established and well-studied complexity models used to study how much information from the input are needed to find an output of a problem. Query complexity studies a setting where an algorithm initially has no information about the input but can (adaptively) ask queries about the input to find an output. Communication complexity studies a setting where the input is distributed between Alice and Bob (or between multiple parties) and Alice and Bob are allowed to communicate in order to find an output. In a game-theoretic context, let S be some solution concept of a normal-form game; E.g., S could be the set of all Nash equilibria, the set of all approximate Nash equilibria, the set of correlated equilibria etc. Query and communication complexity of S are defined as follows.

- **Query complexity.** Initially the algorithm is ignorant of the payoffs in the game. The *query complexity of a solution concept S* captures the number of payoff queries that are needed for an algorithm to find a solution in S.
- **Communication complexity.** Initially every player knows only his own utility function. The *communication complexity of a solution concept S* captures the number of bits that players need to communicate to each other to find a solution in S.

We start by discussing the particular relevance of these two informational models in the game-theoretic context.

Author's address: yakovbab@technion.ac.il

1.1 Informational Complexity and the Speed of Learning

An underlying assumption of any equilibrium notion is that players predict correctly the behavior of their opponents (or the correlated distribution over the action profiles). One justification for this problematic assumption, which appears in the seminal work of [Nash 1951], is that in some scenarios players may learn the behavior of their opponents in cases where the game is played repeatedly. This idea led to an extensive study of learning dynamics and their convergence to equilibria; see, e.g., [Young 2004; Hart and Mas-Colell 2013; Kalai and Lehrer 1993]. One natural, and general, class of adaptive dynamics is that of *uncoupled dynamics* [Hart and Mas-Colell 2003; 2006], where it is assumed that players do not know the utilities of their opponents but observe their past behavior. The possible and the impossible with respect to the existence of uncoupled dynamics that lead to equilibria is quite well understood. Regret minimizing dynamics are known to converge to correlated equilibria [Hart and Mas-Colell 2000; Hart 2005; Blum and Monsour 2007]. Several uncoupled dynamics that converge to approximate Nash equilibria [Foster and Young 2006; Hart and Mas-Colell 2006; Germano and Lugosi 2007; Young 2009; Babichenko 2012] have been introduced. The known dynamics that converge to Nash equilibria and to correlated equilibria have very different characteristics. While regret-minimizing dynamics converge to correlated equilibria persistently and fast (persistency is captured by the fact that the regret of a player forms a supermartingale), the dynamics that converge to Nash equilibria are based on an exhaustive search principle and the convergence is slow.

Informational complexity models and, in particular, communication complexity models constitute a formal framework for studying the rate of convergence of various dynamics. Such models do not specify exact dynamic under consideration but, instead provide bounds on the rate of convergence for all dynamics in a given class. As has been pointed out by [Conitzer and Sandholm 2004], for every solution concept (in particular equilibrium solutions), the communication complexity of a solution is identical (up to a logarithmic factor) to the rate of convergence of any uncoupled dynamics to this solution. Therefore, studying the communication complexity of an equilibrium notion is essentially equivalent to studying the rate of convergence of uncoupled dynamics to this equilibrium notion.

1.2 Informational Complexity versus Computational Complexity

The broad agenda of equilibrium complexity research is to examine whether a given equilibrium notion is an appropriate solution by understanding the hardness of computing the equilibrium. The hardness of computing equilibria in games with many players requires clarification about the input representation. Consider, for instance, n-player binary-action games. A utility of a single player is a mapping $u_i : \{0,1\}^n \to \mathbb{R}$. Hence, a direct representation of a game requires $c \cdot n \cdot 2^n$ bits, when c is the number of bits needed to represent a single payoff in the game.[1] Throughout this survey we assume that c is polylogarithmic in the size of the game. To study the hardness of computing equilibria in games with many players, we shall first

[1]The description of a game with a constant number of players might also be problematic if we consider games with a huge number of actions such as dueling games [Immorlica et al. 2011] and Blotto games [Hart 2008; Behnezhad et al. 2019].

address the issue that a general instance of the problem has an exponential input size.

One approach, which has been adopted in the computational complexity literature, focuses on succinctly representable classes of games such as graphical games, anonymous games, congestion games, and more. The informational complexity literature takes a different approach: it imposes informational restrictions. Query complexity studies a scenario where an algorithm (or a player in the game) is willing to compute an equilibrium in the game, but it (or he) does not know the payoffs. The algorithm (or the player) can deduce the payoff of a specific player (his own payoff or an opponent's payoff) in a specific action profile in constant time. Communication complexity studies a scenario where each player knows only his own payoff.

There are two conceptual differences between informational and computational complexity models. First, the domain of the problem differs. Computational models apply to succinctly representable classes of games. By contrast, informational complexity models typically apply to all games without restrictions, or to non-succinctly representable classes. Moreover, since we do not impose any computational restrictions in the informational complexity models, if the game is succinctly representable we typically can communicate the entire game in a reasonable number of bits and then compute an equilibrium.[2] The second difference is in the type of results. Typical hardness results in the computational model are either PPAD-hardness, PLS-hardness, NP-hardness, or that of another complexity class. Namely, hardness results in the computational model rely on an exponential hypothesis of the corresponding class. By contrast, hardness results in informational models typically do not rely on any hypothesis and provide precise bounds on the amount of information needed for achieving an equilibrium. This is a good place to mention that even though from a conceptual perspective informational complexity models differ from computational complexity models, from a technical perspective the techniques developed to solve informational complexity problems are commonly useful for tackling computational complexity problems and vice versa. Moreover, the existing literature indicates a very tight connection between the computational environment and the informational environment in a variety of settings (e.g., combinatorial allocation, submodular maximization, and many others). Whenever a problem in the computational environment admits an efficient algorithm, its analog in the informational environment typically can be solved efficiently (insofar as the amount of information is concerned). Similarly for negative results. Whenever a problem in the computational environment is shown to be hard, its analog in the informational environment typically requires a large amount of information.

The survey is organized as follows. Section 2 provides the definitions of the central notions in this survey. We summarize the known complexity results in Sections 3, 4, 5, and 6. Section 3 focuses on the most fundamental setting where normal form games are considered and the goal is to classify whether the complexity is polynomial or polylogarithmic in the size of the game. Section 4 discusses specific classes of games. Section 5 provides tighter bounds beyond the polynomial versus

[2] The query model is more delicate because we require payoff queries rather than direct queries of bits from the input. See Section 4.

logarithmic scale. Section 6 discusses the dependence of the results on the approximation value. Section 7 examines which of the solution concepts are total and succinct (i.e., which solution concepts guarantee the existence of a solution that can be represented succinctly for all games). In Sections 8 and 9 we provide high-level ideas of the techniques that have been utilized to prove informational hardness of equilibria. Section 8 starts with very simple and naive techniques and discusses the limitations of these techniques. Section 9 provides the key ingredients of the (much) more involved recent techniques in this field. Finally, Section 10 concludes with some open problems.

2. DEFINITIONS

2.1 Solution Concepts

The definitions of all but one solution concept are standard. In the case of ϵ-correlated equilibria in games with many actions the definition should be formulated carefully. We briefly define the standard notions and then discuss in more detail the case of ϵ-correlated equilibria.

In an n-player m-action game the set of players is $[n]$. The set of actions of player i is $A_i = [m]$. The set of mixed actions of player i is $X_i = \Delta([m])$. The set of pure action profiles is $A = [m]^n$. The set of mixed action profiles is $X = X_1 \times ... \times X_n$. The set of correlated distributions is $C = \Delta(A)$. A utility of player i is given by $u_i : A \to [0,1]$. This utility function can be extended to mixed action profiles and denoted by $u_i(x) = \mathbb{E}_{a_i \sim x_i} u_i(a_1, ..., a_n)$. The utility can also be extended to correlated distributions and denoted by $u_i(c) = \mathbb{E}_{a \sim c} u_i(a)$. Given $a \in A$ ($x \in X$) we denote by a_{-i} ($x_{-i} \in X_{-i}$) the profile of i's opponents. Given $x_{-i} \in X_{-i}$ we denote the best-reply value of player i against x_{-i} by $br(x_{-i}) = \max_{a_i \in A_i} u_i(a_i, x_{-i}) \in [0,1]$.

Definition 2.1. An action profile $a \in A$ is a pure Nash equilibrium if for every player $i \in [n]$ it holds that $u_i(a) \geq u_i(a'_i, a_{-i})$ for every $a'_i \in A_i$.

Definition 2.2. A mixed action profile $x \in X$ is a Nash equilibrium if for every player $i \in [n]$ it holds that $u_i(x) \geq u_i(a'_i, x_{-i})$ for every $a'_i \in A_i$. Equivalently, $x \in X$ is a Nash equilibrium if for every player $i \in [n]$ it holds that $u_i(x) = br(x_{-i})$.

The first definition in 2.2 can be extended to the following approximate notion.

Definition 2.3. A mixed action profile $x \in X$ is an ϵ-Nash equilibrium if for every player $i \in [n]$ it holds that $u_i(x) \geq u_i(a'_i, x_{-i}) - \epsilon$ for every $a'_i \in A_i$.

The second definition in 2.2 can be extended to the following approximate notion.

Definition 2.4. A mixed action profile $x \in X$ is an ϵ-well-supported Nash equilibrium if for every player $i \in [n]$ it holds that $u_i(a_i, x_{-i}) \geq br(x_{-i}) - \epsilon$ for every action a_i that is played with positive probability (i.e., $x_i(a_i) > 0$).

To define approximate correlated equilibrium we use the notion of a *swapping policy*, which is a function $f : [m] \to [m]$.

Definition 2.5. A correlated distribution $c \in C = \Delta(A)$ is a correlated equilibrium if for every swapping policy f of player i it holds that

$$u_i(c) \geq \mathbb{E}_{(a_i, a_{-i}) \sim c} u_i(f(a_i), a_{-i}).$$

This definition of correlated equilibrium is somewhat overcomplicated, because we could restrict attention to the class of m^2 swapping policies $(f_{j,j})_{j,k\in[m]}$ of the form $f_{j,k}(b) = b$ if $b \neq j$ and $f_{j,k}(j) = k$ (i.e., we could swap only action j for action k when j is recommended). It can be easily shown that for an exact correlated equilibrium the requirement $u_i(c) \geq \mathbb{E}_{(a_i,a_{-i})\sim c} u_i(f(a_i), a_{-i})$ for every swapping policy f is equivalent to the requirement $u_i(c) \geq \mathbb{E}_{(a_i,a_{-i})\sim c} u_i(f(a_i), a_{-i})$ for every swapping policy $f \in \{f_{j,k}\}_{j,k\in[m]}$.

Definition 2.5 can be extended to the following approximate notion.

Definition 2.6. A correlated distribution $c \in C = \Delta(A)$ is an ϵ-correlated equilibrium if for every swapping policy f of player i it holds that

$$u_i(c) \geq \mathbb{E}_{(a_i,a_{-i})\sim c} u_i(f(a_i), a_{-i}) - \epsilon.$$

We emphasize that for a constant value of ϵ and for a large number of actions, the definition that requires that $u_i(c) \geq \mathbb{E}_{(a_i,a_{-i})\sim c} u_i(f(a_i), a_{-i}) - \epsilon$ only for simple swapping policies $f = f_{j,k}$ is meaningless. For instance, the uniform distribution over all action profiles is a $\frac{1}{m}$-correlated equilibrium for all games. Indeed, with probability $1 - \frac{1}{m}$ action j will be recommended and then the utility $u_i(a_i, a_{-i})$ is identical to the utility $u_i(f(a_i), a_{-i})$.

Finally, a coarse correlated equilibrium and its approximate notion can be defined by restricting the family of constant swapping functions to constant swaps $\{f_k\}$ such that $f_k(j) \equiv k$. Formally,

Definition 2.7. A correlated distribution $c \in C = \Delta(A)$ is a coarse correlated equilibrium if for every player i and action $a'_i \in A_i$ it holds that

$$u_i(c) \geq \mathbb{E}_{(a_i,a_{-i})\sim c} u_i(a'_i, a_{-i}).$$

Definition 2.8. A correlated distribution $c \in C = \Delta(A)$ is an ϵ-coarse correlated equilibrium if for every player i and action $a'_i \in A_i$ it holds that

$$u_i(c) \geq \mathbb{E}_{(a_i,a_{-i})\sim c} u_i(a'_i, a_{-i}) - \epsilon.$$

The solution concepts and their approximate notions satisfy the inclusions

$$\text{PNE} \subset \text{NE} \subset \text{CE} \subset \text{CCE} \quad \epsilon\text{-PNE} \subset \epsilon\text{-NE} \subset \epsilon\text{-CE} \subset \epsilon\text{-CCE}.$$

For the case of Nash equilibria we also have that NE $\subset$ ϵ-WSNE $\subset$ ϵ-NE.

2.2 Informational Complexity Models

Query complexity problems are defined are follows. A *query protocol* maps every history of past queries and answers to either an additional query or to an output. In our case queries are pairs (a, i) and the answer is $u_i(a)$.

For deterministic query complexity, this mapping is deterministic, and the *cost* of a protocol is defined to be the maximal number of queries (across all inputs, i.e., games) until an output is produced. A protocol is *correct* if it outputs a correct answer (i.e., an equilibrium) for all games. The query complexity of an equilibrium is defined to be the minimal cost across all correct protocols.

For randomized query complexity the mapping from histories to queries might be random, and the *cost* of a protocol is defined to be the maximal expected number of queries (maximum across all inputs and the expectation is taken with respect to the randomization of the protocol) until an output is produced. A protocol is *correct* if it outputs a correct answer with probability 2/3 (again with respect to the randomization of the protocol) for all games. The query complexity of an equilibrium is defined to be the minimal cost across all correct protocols.

In communication complexity the definitions are similar. A communication protocol maps every history of communication to either a 0/1 message from Alice to Bob, a 0/1 message from Bob to Alice, or a termination. In the case of termination Alice and Bob need to produce an outcome.

For deterministic communication complexity, this mapping is deterministic, and the *cost* of a protocol is defined to be the maximal number of messages (across all inputs, i.e., games) until termination. A protocol is *correct* if the output of Alice coincides with the output of Bob and is correct (i.e., is an equilibrium) for all games. The communication complexity of an equilibrium is defined to be the minimal cost across all correct protocols.

For randomized communication complexity, we assume that Alice and Bob have access to an infinite string of random public coin flips. The mapping from histories to messages might be random (i.e., based on the realizations of the random public coin flips), and the *cost* of a protocol is defined to be the maximal expected number of queries (maximum across all inputs and the expectation is taken with respect to the random public coin flips) until termination. A protocol is *correct* if it is correct with probability 2/3 (again with respect to the randomization of the protocol) for all games. The communication complexity of an equilibrium is defined to be the minimal cost across all correct protocols.

For the case of more than two players, the definitions are similar, where in each step of the protocol player i sends a public message to all the other players.

We refer the reader to the books of [Kushilevitz and Nisan 1997; Arora and Barak 2009] on the topic of communication complexity.

Note that every query protocol (deterministic or random) with cost h induces a communication protocol with cost ch, where c is the number of bits needed to represent a single payoff in the game. Instead of querying the payoff of player i, player i sends c public messages to the other players and then all players know this payoff, as in the query model. Hence the communication complexity of any problem is at most the query complexity of that problem (multiplied by a negligible factor of the representation size of a single payoff).

It is interesting to note that all negative communication complexity results presented in this survey hold not only for the n-party communication problem but even in the simpler problem where Alice holds as a private input the utilities of players $i = 1, ..., n/2$ and Bob holds the utilities of players $i = n/2 + 1, ..., n$. The positive results, on the other hand, apply to the more restrictive model of n-party communication.

3. POLYNOMIAL VERSUS LOGARITHMIC COMPLEXITY

Note that in each informational model, that of query complexity or communication complexity a trivial upper bound on the complexity of a problem is the input size: players can just query or communicate the entire game.[3] Probably the most fundamental question that can be asked is whether the complexity is polynomial or polylogarithmic in the input size. This survey focuses mainly on this question. A discussion of more precise bounds (i.e., the power of the polynomial or the logarithm) is relegated to Section 5. After almost two decades of study, we currently have a very good understanding of the informational complexity of the most common equilibrium notions. The results are summarized in Tables I and II. These tables include an additional column that specifies whether the solution concept is *succinctly total*. Simply speaking, a solution concept is succinctly total if every game admits a succinctly representable solution, where succinctness is defined to be polylogarithmic in the size of the game. Note that succinctness of the output is a very natural desideratum in this setting, as we simply want the outcome of the algorithm (or the communication) to be of reasonable size. For further discussion on succinct totality see Section 7.

The notion of ϵ-equilibrium in Tables I and II refers to an additive and constant value of approximation. Other values of approximation are discussed in Section 6.

	Suc Tot	Det QC	Rand QC	Det CC	Rand CC
Pure Nash	No	poly	poly	poly	poly
Nash	No	poly	poly	poly	poly
ϵ-Nash	Yes	poly	poly	poly	poly
Correlated	Yes	poly	poly	log	log
ϵ-correlated	Yes	poly	log	log	log

Table I. Complexity of equilibria in n-player binary-action games. Polynomial complexity means that a $2^{\Omega(n)}$ lower bound is known. Logarithmic complexity means that a $\mathsf{poly}(n)$ upper bound is known.

	Suc Tot	Det QC	Rand QC	Det CC	Rand CC
Pure Nash	No	poly	poly	poly	poly
Nash	No	poly	poly	poly	poly
ϵ-Nash	Yes	poly	poly	poly	poly
Correlated	No	poly	poly	poly	poly
ϵ-correlated	Yes	poly	poly	?	?
Coarse correlated	No	poly	poly	poly	poly
ϵ-coarse correlated	Yes	poly	poly	?	log

Table II. Complexity of equilibria in two-player m-action games. Polynomial complexity means that a $\mathsf{poly}(m)$ lower bound is known. Logarithmic complexity means that a $\mathsf{polylog}(m)$ upper bound is known.

We discuss the results of Tables I and II for each of the solution concepts in turn.

Nash equilibria. The results indicate that all variants of Nash equilibrium problems in all informational models require polynomial information in the size of the game. The deduction of these results has a long history, which is summarized below. The study of informational complexity of equilibria was initiated by [Conitzer

[3]We impose no computational restrictions on the players, and so the computation of an equilibrium once the entire game is known comes with no cost.

and Sandholm 2004] who showed an $\Omega(m^2)$ communication complexity bound on computing a pure Nash equilibrium in two-player m-action games. [Conitzer and Sandholm 2004] also presented results about more basic solution concepts such as sequential elimination of dominated strategies. This hardness result is obviously valid also for the query complexity model. [Hart and Mansour 2010] proved the hardness of a pure Nash equilibrium in n-player games, as well as the communicational hardness of an exact Nash equilibrium in n-player games. [Babichenko 2016] proved the first result for a total variant of a Nash equilibrium problem. Specifically, [Babichenko 2016] showed that an ϵ-well-supported Nash equilibrium requires an $\exp(n)$ number of queries. Subsequently, [Chen et al. 2017] extended the result to ϵ-Nash equilibrium for slightly worse bound of $\exp(\frac{n}{\log n})$. Finally, [Rubinstein 2016] improved the bound to $\exp(n)$ for an ϵ-Nash equilibrium. The communication complexity lower bounds on an approximate Nash equilibrium (in the two-player and the n-player settings) was proved in [Babichenko and Rubinstein 2017].

Correlated equilibria. The informational complexity of correlated equilibrium notions is more intricate. The story starts with regret-minimizing algorithms [Littlestone and Warmuth 1994; Cesa-Bianchi and Lugosi 2006] that guarantee convergence of the regret to 0 at a rate that is polynomial (quadratic) in the approximation and logarithmic in the number of actions. [Hart and Mas-Colell 2000; 2001] showed the connection of these regret minimizing algorithms to the learning of correlated equilibrium. The regret-minimizing algorithms are translated to an algorithm whose empirical distribution of play forms an ϵ-correlated equilibrium in time that is logarithmic in the number of players (n), polynomial in the approximation value (ϵ), and (only) polynomial in the number of actions (m). Indeed, in the definition of an approximate correlated equilibrium (Definition 2.6), we require that the regret of the m^m swapping functions be low. This requires $\log(m^m) = m \log m$ iterations of regret-minimizing algorithms. See the discussion after Definition 2.6 on why it is necessary to consider all the m^m swapping functions to achieve an appropriate notion of approximate equilibrium.

In each step of the algorithm, the updating of the regrets requires only nm payoff queries. In the communication setting the updating of the regrets requires no communication. This yields a $\mathsf{poly}(n, m)$ query algorithm and a corresponding $\mathsf{poly}(n, m)$ communication protocol. Note that in n-player games with a constant number of actions, this implies an algorithm that is logarithmic in the input (i.e., $\mathsf{poly}(n)$) for a correlated equilibrium.[4] By contrast, for two-player games with many actions, this only implies an algorithm that is polynomial in the input for a correlated equilibrium.

It is important to notice that regret-minimizing algorithms are randomized and that they converge to an approximate correlated equilibrium. [Hart and Nisan 2018] showed that in the query complexity model, the properties of randomization and of producing an approximate correlated equilibrium are both essential for solving the problem in $\mathsf{poly}(n)$ queries. Namely, the deterministic query complexity of an

[4] Moreover, in an alternative standard query model where an answer to a query is the payoff profile rather than the payoff of a specific player, only $O(\log(n))$ payoff-profile queries are sufficient; see [Goldberg and Roth 2016]. However, note that each single query reveals $\Theta(n)$ information about the game.

ϵ-correlated equilibrium is exponential in n and the randomized query complexity of an exact correlated equilibrium is exponential in[5] n.

In contrast to query complexity, the communication complexity of a correlated equilibrium is polynomial in n even for exact correlated equilibrium, and even for deterministic communication protocols. Based on the ellipsoid against hope ideas of [Papadimitriou and Roughgarden 2008], [Jiang and Leyton-Brown 2015] showed a deterministic algorithm that computes an exact correlated equilibrium with $\mathsf{poly}(n)$ communication.

Coarse correlated equilibria. Note that we did not specify any results for coarse correlated equilibria in Table I. The reason for that is that in binary-action games a coarse correlated equilibrium (and its approximate notion) is identical to a correlated equilibrium. Hence, the same results apply. However, for two-player games with many actions, these equilibria notions differ.

Regret-minimizing algorithms converge to an approximate coarse correlated equilibrium in a rate that is logarithmic in the number of actions. Together with the observation that in the communication complexity setting players know their own regret (in particular, without the need for payoff queries), we can deduce that regret-minimizing algorithms require just $\mathsf{polylog}(m)$ rounds to converge to an ϵ-coarse correlated equilibrium. Each round requires us to specify a single played action which again requires $\log(m)$ communication.

The negative results for the query complexity of an approximate coarse correlated equilibrium (Corollary 8.1) and the communication complexity of an exact coarse correlated equilibrium (Corollary 8.2) appear later in this survey.

4. SPECIFIC CLASSES OF GAMES

The hardness of Nash equilibria in general games raises the question of whether it can be computed in classes of games. For this problem to be meaningful in an informational complexity setting, we need to choose a class of games that does not admit a succinct representation. Obviously, any succinctly representable game can be solved in the communication model by a protocol in which every player communicates his succinctly representable utility to the others: it requires low communication. In the query model the situation is slightly more complicated. The standard query model we have considered so far uses payoff queries. If queries are allowed to ask about input, this solves the problem. But what if queries are restricted to asking about payoffs? One can ask how the restriction to payoff queries might affect equilibrium computation in a variety of classes of games, such as congestion games [Fearnley et al. 2015], graphical games, polymatrix games, anonymous games, et cetera. Interestingly, there is a generic answer to all of these questions that has been provided in [Goldberg and Roth 2016]. [Goldberg and Roth 2016] focus on the most general formulation of *succinctly representable games*: a collection of 2^k games (which might have no relation to each other). [Goldberg and Roth 2016] prove the existence of a robust query protocol that computes an ϵ-Nash equilibrium in a $\mathsf{poly}(k)$ number of queries. The algorithm is based on the

[5] In parallel, [Babichenko and Barman 2015] showed the weaker result that the deterministic query complexity of an exact correlated equilibrium is exponential in n.

idea of having a hypothetical game in mind (in their case it is defined to be the median among all consistent games in the class), computing an equilibrium for this game (which incurs no cost in the query model), and checking whether the computed equilibrium is an approximate equilibrium of the actual game. The key point is that whenever the computed equilibrium of the hypothetical game is not an equilibrium, a large (constant) fraction of the consistent games are eliminated.

4.1 Potential Games

There are few of the fundamental classes of games in game theory that have no succinct representation. One important class is that of potential games [Monderer and Shapley 1996], which includes as a special case the class of congestion games.[6] Potential games are characterized by the property of the existence of a single potential function that maps action profiles to real numbers. This potential function captures gains and losses from unilateral deviations by any single player; Namely, the difference between the potential values of two action profiles that is obtained by unilateral deviation is equal to the difference between the utilities of the deviating player in these two action profiles. Pure Nash equilibia in potential games are the local maxima (with respect to unilateral deviations) of the potential function and, in particular, a pure Nash equilibrium is guaranteed to exist. The query complexity of a pure Nash equilibrium in potential games has been studied by [Nisan 2009], who showed a $\tilde{\Omega}(2^n)$ lower bound for the case of n-player binary-action games. Recently, [Babichenko et al. 2019] showed the hardness of a pure Nash equilibrium in the communication model. Concretely, [Babichenko et al. 2019] show a $\mathsf{poly}(m)$ lower bound in two-player m-action games and a $2^{\Omega(n)}$ lower bound for n-player binary-action games. Interestingly, the problem of finding a pure Nash equilibrium in potential games might be viewed as a succinctly total problem; see Section 7.1.

For an ϵ-Nash equilibrium there exists an efficient communication protocol that requires only $\frac{2n}{\epsilon}n(n+\log m)$ communication in n-player m-action games. The same protocol requires $\frac{2}{\epsilon}n^2 m$ queries, which is efficient for games with a constant number of actions.[7] This protocol implements the ϵ-*better-reply dynamic*, where in each step a single player updates his action to an ϵ-better one if such an action exists. This dynamic yields an ϵ-improvement of the potential every time a player updates an action. A single improvement requires $n + \log m$ communication. Communication of n bits is needed to find a player that can make an improvement, and then $\log m$ bits are needed to describe the better action. The key observation is that the potential function of a potential game with payoffs in $[0, 1]$ has a potential function is bounded in $[-n, n]$ and, therefore, the number of ϵ-improvements is bounded by $\frac{2n}{\epsilon}$.

[6]In fact, potential games are equivalent to congestion games where the number of resources might not be polynomial in the number of players and hence their representation size is not polynomial in the number of players.

[7]One cannot expect to have polylogarithmic dependence on m in the query model. The hidden dominant strategy (see Section 8.1) generates a potential game.

5. TIGHTER BOUNDS

We have a better understanding of complexity for certain equilibrium notions in two-player games. Table III summarizes the known lower and upper bounds.

	Det QC	Rand QC	Det CC	Rand CC
Pure Nash	m^2	m^2	m^2	m^2
Nash	m^2	m^2	m^2	m^2
ϵ-Nash	m^2	m^2	m^2	m^2
correlated	m^2	m^2	m^2	m^2
ϵ-correlated	m, m^2	m, m^2	?	$?, m$
coarse correlated	m, m^2	m, m^2	m, m^2	m, m^2
ϵ-coarse correlated	m, m^2	m	?	$\log^2(m)$

Table III. Bounds on the complexity of equilibria in two-player m-action games. A single number indicates matching of lower and upper bounds up to a logarithmic factor. Two numbers indicate differing lower and upper bounds.

The first row in Table III follows from [Conitzer and Sandholm 2004], who showed that $\Omega(m^2)$ communication is needed to determine the existence of a pure Nash equilibrium. This result implies the same bound for the more restricted query model. The second and third rows follow from a recent result by [Göös and Rubinstein 2018], who proved that not only that polynomial communication is needed for an ϵ-Nash equilibrium in two-player games but also that, in fact, communication of $\Omega(m^{2-c})$ is needed for every $c > 0$ is needed; i.e., a communication of almost the entire game. Again, this result carries over to the query model.[8] The fourth row follows from the representation size of a correlated equilibrium. It is shown in [Viossat 2008; Nitzan 2005] that the set of games with a unique correlated equilibrium (which is also a Nash equilibrium) is rich enough, and contains an open ball. Using these games one can deduce games where the representation size of a correlated equilibrium requires $\Omega(m^2)$ bits for representation. The two additional cells that have some bounds (beyond those that are presented in Corollaries 8.1 and 8.2) are the randomized communication complexity of an ϵ-correlated equilibrium and the randomized query complexity of an ϵ-coarse correlated equilibrium. For an ϵ-correlated equilibrium, regret-minimizing algorithms require $O(\log(m^m)) = O(m\log(m))$ steps to converge to an equilibrium (because the number of swapping functions is m^m). Each step requires only $\log m$ communication. For an ϵ-coarse correlated equilibrium, it has been observed by [Goldberg and Roth 2016] that regret-minimizing algorithms require $O(m \log m)$ queries.

For n-player games no bounds beyond the logarithmic versus polynomial are known, except for the case of a pure Nash equilibrium where it is known that $\Omega(2^n)$ communication is needed; see [Hart and Mansour 2010].

6. THE DEPENDENCE OF THE BOUNDS ON THE APPROXIMATION VALUE

This survey have focused on approximate notions of equilibria with a small and constant value of approximation ϵ. One may ask how the complexity of the problem

[8]Prior to this excellent work, no better bound than $\Omega(n)$ was known even for the simpler query model.

depends on ϵ. There are to way to approach this problem. The first is to start with problems that are tractable (i.e., of logarithmic complexity) for constant ϵ and ask what is the affect or reducing ϵ to $\epsilon = 1/\mathsf{poly}(n)$ for n-player games or to $\epsilon = 1/\mathsf{poly}(m)$ for two-player m-action games. The second approach is to start with problems that are intractable (i.e., of polynomial complexity) and ask how much should we increase ϵ to make the problem tractable.

We focus first on the latter approach. The positive results for n-player binary-action games (see Table II) remain true not only for constant ϵ but also for $\epsilon = 1/\mathsf{poly}(n)$. This follows from the polynomial dependence on the approximation value of the regret-minimizing algorithms and of the ellipsoid against hope algorithm. For two-player m-action games, setting $\epsilon = 1/\mathsf{poly}(m)$ makes all the problems intractable, including even the simplest problem of an ϵ-coarse correlated equilibrium in a randomized communication complexity model. One can obtain this hardness result on ϵ-coarse correlated equilibrium by analyzing the hide-and-seek game with bad actions (see Section 8.2) for a $1/\mathsf{poly}(m)$-coarse correlated equilibrium. See also [Ganor and Karthik 2018].

Regarding the second approach of increasing ϵ and finding a sufficiently large value that makes the problem tractable, the overall picture is that we do not have a good understanding of the values of ϵ where the transition from tractable to intractable occurs. The gap between the positive and the negative results is huge. The current negative results are theoretical in nature. The deduction of these results does not involve trying to optimize the arguments with respect to ϵ and, therefore, results in bounds of order $\epsilon \approx 10^{-6}$ or in many cases even much smaller. Little is known about the positive results that significantly improve upon trivial protocols.[9] The best values for which communicationally efficient algorithms are known are $\epsilon = 0.38$ for the Nash equilibrium and $\epsilon = 0.65$ for the well-supported-Nash equilibrium [Goldberg and Pastink 2014; Czumaj et al. 2019].

7. SUCCINCT TOTALITY OF EQUILIBRIA

The existence of an (approximate) equlibrium in every game is guaranteed for all solution concepts (Nash, correlated, and coarse correlated equilibrium) is guaranteed except for pure Nash equilibrium. However, the existence of a succinct equilibrium is not guaranteed for all the (approximate) solution concepts. In informational complexity problems, it is desirable that the output be of much smaller size than the input (i.e., polylogarithmic). We briefly discuss which of the above solution concepts are guaranteed to have a succinctly representable solution.

Which problems are not total? All the exact solution concepts in two-player m-action games are not total. The description of an exact Nash equilibrium requires us to specify the probability distribution of m actions, which requires a polynomial description in the size of the game. Exact correlated and coarse correlated equilibria requires to specify a distribution over the m^2 profiles. It is easy to construct games (e.g., zero-sum games) where all correlated (coarse correlated) equilibria

[9] One example of a trivial protocol might be: Alice sends to Bob the action profile (a, b) of her best outcome in the game, and then Bob plays with equal probability the action b and the action b^* that is the best reply to a. This protocol leads to a $\frac{1}{2}$-Nash equilibrium.

have support $\Omega(m)$ and hence the description of such an equilibrium is $\mathsf{poly}(m)$.

An exact mixed Nash equilibrium in n-player binary-action games requires us to specify only n real numbers in $[0,1]$, which in principle looks to be succinct. However, [Hart and Mansour 2010] showed that even for games with outcomes in $\{0,1,2\}$, the description of an exact Nash equilibrium might require a doubly-exponential precision of the mixed strategy. Hence this problem is not total.

Which problems are total? Correlated equilibria in n-player binary-action games are given by a linear program with 2^n variables but only $2n$ constraints. Hence the existence of a correlated equilibrium (an extreme point of the feasible set) with support $2n$ is guaranteed.

The existence of an ϵ-Nash equilibrium whose mixed strategies are located on a grid of size $\Omega(\frac{1}{n})$ is guaranteed by a simple rounding of the probabilities to the grid points.[10]

In two-player m-action games, [Lipton et al. 2003] showed that the existence of ϵ-Nash equilibrium where both players randomize uniformly over a multi-set of $O(\log(m))$ actions. Such a strategy requires only $O(\log(m)^2)$ bits of representation. This ϵ-Nash equilibrium is, in particular, an ϵ-correlated and an ϵ-coarse correlated equilibrium whose support is succinct over the action profiles[11] $O(\log^2(m))$.

Non-deterministic communication complexity. Simply speaking, the non-deterministic complexity of a problem counts the output size of the problem and the amount of communication needed to verify the solution. Note that verification of equilibria requires a single bit of communication: a player who knows his own utility function can verify whether he has a better response (or an ϵ-better response) and send this information to the other players. Hence, the non-deterministic communication complexity for all the succinctly total equilibrium notions is low (logarithmic in the input).

7.1 Potential Games

We recall that a problem is succinctly total if it has a succinct solution for every input. Potential games have a restriction on the input: the game needs to have the potential structure. Therefore, the problem of finding a pure Nash equilibrium in potential games is prima facie not total. However, by combining the following two observations:

(a) There is succinct evidence that a game is not a potential game. This evidence comes in the form of a *unilaterally deviating cycle* where the sum of gains and losses from the deviations is not equal to 0; see [Monderer and Shapley 1996]. A unilaterally deviating cycle is a cycle of profiles of length at most $2n$ that is

[10] More involved arguments can show the existence of an ϵ-Nash equilibrium on a grid of size $\Omega(\frac{1}{\log(n)})$; see [Babichenko et al. 2016]. The existence of an ϵ-Nash equilibrium on a grid of size $\Omega(1)$ remains an open question.

[11] It is interesting to notice that standard linear programming arguments fail to prove the existence of an ϵ-correlated equilibrium with support $\mathsf{polylog}(m)$. From a mathematical point of view, this means that we use Brouwer's fixed-point theorem to prove the existence of a sparsely supported correlated equilibrium. I am not aware of a proof that avoids the use of Brouwer's fixed-point theorem.

obtained by unilateral deviations.

(b) There is a communicationally efficient protocol to verify whether a game is a potential game; otherwise, it outputs a succinct evidence of unilaterally deviating cycle; see [Babichenko et al. 2019].

we can define a close variant of the pure Nash equilibrium problem: the algorithm should either output a pure Nash equilibrium or it should output a unilaterally deviating cycle where gains and losses do not sum up to 0. This problem is total, of low non-deterministic complexity, and is communicationally hard; see [Babichenko et al. 2019]. Very recently, [Babichenko and Rubinstein 2019] extended the hardness of finding a Nash equilibrium in potential games from the case of a pure Nash equilibrium to the general case of any (possibly mixed) equilibrium.

8. SIMPLE TECHNIQUES

There are several very simple techniques to "hide" equilibria in games. I find it useful to start the technical discussion with these simple techniques and to understand their limitations (i.e., in which environments they do not work). An advantage of these simple techniques is that in cases where they succeed in providing a lower bound, this lower bound applies to the weakest solution concept of coarse correlated equilibrium.

8.1 A Hidden Dominant Strategy

In a two-player game (and in fact even in a single-player decision problem), we can choose uniformly at random a single action $a^* \in [m]$ of Player 1, and set his utilities to be $u_1(a_1, a_2) = \mathbf{1}_{a_1=a^*}$. In the query model, Player 1 will need $O(m)$ queries (even in the randomized model) to find this hidden strategy. This observation shows that even for the weakest solution concept, namely, the ϵ-coarse correlated equilibrium, the query complexity is polynomial in the input.

COROLLARY 8.1. *The randomized query complexity of an ϵ-coarse correlated equilibrium in two-player m-action games is $\Omega(m)$.*

Note that in the communication complexity setting this technique fails: Player 1 knows his entire utility function and hence identifies the dominant strategy immediately. The following example applies to the communication complexity setting.

8.2 Hide-and-Seek Game with Bad Actions of the Hider

Consider a hide-and-seek zero-sum game with $u_2(a_1, a_2) = -u_1(a_1, a_2) = \mathbf{1}_{a_1=a_2}$ where Player 1 is the Hider and Player 2 is the Seeker. It is easy to check that any exact coarse-correlated equilibrium has uniform marginals for both players.

We next slightly modify this game by making half of the Hider's actions dominated. We pick a subset $B \subset [m]$ with $|B| = m/2$ uniformly at random, and we set $u_1(a_1, a_2) = -2 \cdot \mathbf{1}_{a_1 \notin B} - \mathbf{1}_{a_1=a_2}$. We do not change the utilities of the Seeker, which remain $u_2(a_1, a_2) = \mathbf{1}_{a_1=a_2}$. Simple arguments of elimination of dominated strategies show that in any coarse correlated equilibrium the marginals for both players are the uniform distribution over B. Note that the Seeker has no information about B from his utility, but a coarse correlated equilibrium identifies B.

Namely, in any communication protocol the Hider has to communicate B to the Seeker, which requires $O(m)$ communication.

Corollary 8.2. *The randomized communication complexity of an exact correlated equilibrium in two-player m-action games is $\Omega(m)$.*

Note that this construction fails to provide bounds for approximate notions of equilibria. The following approximate Nash equilibrium requires no communication: the Hider chooses a location uniformly at random from B and the Seeker chooses a location uniformly at random from $[m]$. By best-replying to the Hider's strategy, the Seeker can increase his expected payoff from $1/m$ to $2/m$. Hence this profile of mixed actions forms a $1/m$-Nash equilibrium.[12]

9. RECENT LOWER BOUND TECHNIQUES FOR NASH EQUILIBRIA

Several recent papers have succeeded in proving lower bounds on Nash equilibria in the hard-to-prove environment where the problem turns out to be total (and of low non-deterministic communication complexity); see [Babichenko 2016; Babichenko and Rubinstein 2017; Göös and Rubinstein 2018; Babichenko et al. 2019; Babichenko and Rubinstein 2019]. Even though the results are in different settings (query or communication complexity, two-player or n-player games, general or potential games), all the proofs share a common structure. This common structure consists of the following six ingredients which are discussed in more details thereafter.

(1) Start with a query-hard end-of-line problem.
(2) In the communication model, "lift" the query-hard problem to a communicationally hard end-of-line problem.
(3) Embed the line in the end-of-line problem as a continuous function.
(4) Embed the function as a continuous-action imitation game.
(5) In the communication model, introduce into the imitation game incentives to report truthfully the private local information about the line.
(6) Discretize of the imitation game.

1. Query-hard end-of-line problem. The starting point of the above listed reductions is some end of (a single) line problem over a low-degree graph. The starting point of the line is known and the task is to find its end. This low-degree graph might be directed or undirected, and the line can be metered or unmetered depending on the application.[13] *Local behavior of a line* in a given vertex is specified by whether the line passes through this vertex and, if so, what are its previous visit and its next visit. It is crucial that the underlying graph will have a low degree (in

[12]The same technique can be applied to prove the communicational hardness of an ϵ-Nash equilibrium for $\epsilon = \mathsf{poly}(1/m)$, but not beyond that.

[13]A line is metered if the vertex through which the line passes indicates on the distance of the line from its origin. For instance, a line over a positively directed two-dimensional grid is metered because the pair of coordinates (x, y) indicate that the line has passed through $x + y$ vertices so far.

fact, in many proofs it has a constant degree) because then the local behavior of the line will have a very succinct (constant) representation.

2. Communicationally hard end-of-line problem. For communication complexity problems, the reduction should start with some communicationally hard problem. The tool that is utilized here is simulation theorems. Simulation theorems are a beautiful tool of recent development; see [Raz and McKenzie 1999; Göös and Pitassi 2014; Göös et al. 2015; 2017], just to mention a few. The idea is to "lift" a query-hard problem to a communicationally hard problem by carefully splitting the information about the problem between Alice and Bob. An outstanding fact about simulation theorems that they can be applied to any problem. The typical gadget that is used to carefully distribute information is the *index gadget*: for each element in the input Alice holds an array of possible inputs and Bob holds an index of the correct element.[14] Note that this distribution of information is very different from the one we have in game-theoretic settings, where each player simply knows his own utility function; therefore additional (and quite substantial) work is needed in order to apply this beautiful tool to game-theoretic problems.

To summarize, the second step in the established proof techniques (based on existing literature on simulation theorems) defines a communicationally hard instance of end-of-line where Alice holds arrays that contain information on possible local behavior of the line in each vertex and Bob holds indices on where the correct local behavior of the line is hidden in Alice's arrays.

3. Embedding a line as a continuous function. In the next step it will become clearer why we insist on "making the end-of-line problem continuous"; for now, we will just describe what it means. We embed our host graph in some well structured graph as the δ-grid of $[0,1]^k$. We embed the line (that currently is over the δ-grid) to a Lipschitz function $f:[0,1]^k \to [0,1]^k$ (or, in some applications to a potential function $f:[0,1]^k \to [0,1]$). The key properties that we want from this construction are *locality* and *reducibility*.

Locality means that the value of the function at a point $x \in [0,1]^k$ can be calculated from the local behavior of the line in the neighborhood of x when, roughly speaking, the local behavior of the line with respect to a continuous point $x \in [0,1]^k$ is defined by the local behavior at the closest grid point.

Reducibility means that all solutions of f should be located close to the end of the line, where solutions are interpreted as fixed points in the case of $f:[0,1]^k \to [0,1]^k$ and are interpreted as local maxima in the case of a potential function $f:[0,1]^k \to [0,1]$.

Such a (highly non-trivial) embedding was introduced in [Hirsch et al. 1989] for the case of $f:[0,1]^k \to [0,1]^k$. Latter modifications of this construction appear in [Rubinstein 2016; Chen and Deng 2008]. In the case of a potential function, variants of such constructions appear in [Hubáček and Yogev 2017; Babichenko and Rubinstein 2019].

[14] In some applications (e.g., [Göös and Rubinstein 2018; Babichenko et al. 2019]) it is crucial that the size of these arrays will be constant. A simulation theorem for an index gadget of constant size is known only for specific problems. Luckily for us, one of these problems is a variant of an end-of-line problem; see [Göös and Rubinstein 2018].

This is a good place to mention the closely related literature on the query and communication complexity of finding a fixed point.[15] In the fixed point problem the input is a continuous function $f : A \to A$, where A is a compact convex set and the output is an ϵ-fixed point of the function f. In the query model [Hirsch et al. 1989] show an exponential, in the dimension, lower bound for this query problem (i.e., an $\exp(n)$ lower bound for the case of $A = [0,1]^n$), even for an ϵ-fixed point with a constant ϵ. [Hirsch et al. 1989] show also a polynomial in a $\frac{1}{\epsilon}$ lower bound for the two-dimensional problem (i.e., a $\mathsf{poly}(\frac{1}{\epsilon})$ lower bound for the case of $A = [0,1]^2$). Recently, communication variants of this problem have been studied. For instance, in the *decomposition problem* Alice holds a function $f : A \to A$, Bob holds a function $g : A \to A$, and their goal is to compute a fixed point of the decomposition $f \circ g$. [Ganor et al. 2019; Roughgarden and Weinstein 2016] show that communication variants of fixed point computation are as hard as the query problem; namely, they show that the lower bounds of [Hirsch et al. 1989] apply to the communication problem as well.

4. Embedding the function as a continuous-action imitation game. The main obstacle in proving hardness results on mixed equilibrium notions comes from the fact that it is typically hard to prove the nonexistence of malicious equilibria (with large support). To demonstrate this point we suggest some naive (and not completely specified) approach to solve end-of-line with an equilibrium of some game.

For sake of simplicity, assume that the given line passes through all the vertices of G. We propose the following simple game: Alice and Bob choose a vertex. We can design a game whose incentives reflect that Alice wants to be the successor of Bob, and Bob wants to match Alice. Indeed if we focus on pure Nash equilibria the only stable scenario is the the profile where Alice and Bob choose the end-of-line vertex where both are happy (the end of the line is defined to be the successor of itself). However, once we focus on approximate Nash equilibria the action profile where both players are randomizing uniformly is an approximate equilibrium. Indeed, Alice would gain only $1/|V|$ by deviating to the end-of-line vertex. This equilibrium does not provide any information about the location of the end of the line, an obstacle common to many such naive approaches.

However, one can overcome this obstacle when the problem is continuous. Given a continuous function $f : [0,1]^k \to [0,1]^k$ there is a very simple two-player imitation game with a continuum of actions all of whose Nash equilibria are pure and correspond to fixed points of f: Alice and Bob choose points $x, y \in [0,1]^k$. Alice wants to match $f(y)$ and Bob wants to match x. More specifically, their utilities are given by $u_A(x,y) = -\|x - f(y)\|_2^2$ and $u_B(x,y) = -\|y - x\|_2^2$. Even if Bob is playing a mixed strategy $\beta \in \Delta([0,1]^k)$, Alice has a unique best response, which is $x^* = \mathbb{E}_{y\sim\beta}[f(y)]$. This simply follows from the fact that expectation is the unique minimizer of the square error. Similarly, Bob has a unique best reply. Therefore, any Nash equilibrium in this game is pure. Once we understand that, it is immediate to verify that the pure Nash equilibrium must be a fixed point. This idea of an

[15]Indeed, the close connection between Brouwer's fixed-point Theorem and Nash equilibria is well known in both directions [Nash 1951; Shmaya 2012].

imitation game was introduced in [McLennan and Tourky 2005; Shmaya 2012]. In some applications, more complicated imitation techniques are required (see, e.g., [Babichenko and Rubinstein 2019]).

5. Incentivizing truthful reporting. In the communication model, neither Alice nor Bob knows the line. In other words, they do not know f, and hence Alice's utility cannot be defined simply by $-\|x - f(y)\|_2^2$. However, for any point x Alice knows her local information at point x (see step 2). We provide Alice with a strong incentive to report this information truthfully. The same holds for Bob with respect to his point y. The key point is that if Alice and Bob are choosing points close to each other, they can combine this information to deduce the local information about the line, and by the locality of f they can compute $f(y)$ (in fact, only Alice needs to perform this computation). In other words, Alice's utility is defined with respect to the combined reported information rather than with respect to the actual f.

6. Discretization. Finally, to conclude the reduction it is necessary to the convert the action space $[0,1]^k$ back to discrete. The elegant consequences of the purity of Nash equilibria in step 4 translates to the observation that in every approximate well-supported Nash equilibrium Alice and Bob choose actions in a small cube of the δ-grid, in which the merging of the local information remains possible.

To strengthen the reduction to approximate Nash equilibria rather than to approximate well-supported Nash equilibria, additional techniques are needed. [Chen et al. 2017] have suggested the technique of replicating players, and [Rubinstein 2016] have suggested the technique of error-correcting codes.

10. OPEN PROBLEMS

(1) Arguably the most fundamental questions that remain open are those of logarithmic versus polynomial complexity. As indicated in Table II, most of the problems regarding the communication complexity of approximate correlated (or coarse correlated) equilibria in two-player games remain open. Specifically, for an ϵ-correlated equilibrium in the deterministic and the randomized communication settings, is there a polynomial lower bound? Is there a polylogarithmic protocol? And do they exist for ϵ-coarse correlated equilibrium in the deterministic communication model?

(2) For an ϵ-Nash equilibrium in n-player binary-action games no algorithm that improves upon 2^n (in the query or communication complexity model) is known. The negative results, on the other hand, prove exponential hardness for $(1+\delta)^n$ for a constant but very small δ. It will be interesting to close these gaps (in both the query and communication models). As indicated in Table III, there are also gaps in the polynomial complexity of correlated notions of equilibria in two-player games.

(3) As mentioned in Section 6, the complexity of an ϵ-Nash equilibrium as a function of constant values of ϵ is far from being understood.[16]

(4) As indicated in Table III, there are several instances of correlated equilibria in two-player m-action games where the answer is known to be bounded in between m and m^2, but the correct power of m is not known.

(5) Beside potential games there is another interesting class of not succinctly representable games for which the existence of a pure approximate Nash equilibrium is guaranteed. Due to non-succinctness, informational complexity questions are relevant. [Azrieli and Shmaya 2013] introduced a class of λ-*Lipschitz games* where the influence of player i on the utility of player j is bounded by λ. They show that in every n-player λ-Lipschitz game with $\lambda = \tilde{O}(n^{-\frac{1}{2}})$, existence of a pure approximate Nash equilibrium is guaranteed. The complexity (informational or computational) of finding such an approximate pure Nash equilibrium is not known.

(6) Quantum computation allows us to:

(a) Define new solution concepts as *quantum correlated equilibria* that lie in between Nash equilibria and correlated equilibria; see [Deckelbaum 2014].

(b) Look at more powerful communication environments such as quantum communication complexity; see [Brassard 2001].

Little is known about the complexity of these quantum solution concepts, and many problems remain open in regard to quantum complexity models in game-theoretic settings. The reader is referred to [Rubinstein 2018] for a discussion of these issues.

REFERENCES

Arora, S. and Barak, B. 2009. *Computational complexity: a modern approach.* Cambridge University Press.

Azrieli, Y. and Shmaya, E. 2013. Lipschitz games. *Mathematics of Operations Research 38,* 2, 350–357.

Babichenko, Y. 2012. Completely uncoupled dynamics and nash equilibria. *Games and Economic Behavior 76,* 1, 1–14.

Babichenko, Y. 2016. Query complexity of approximate nash equilibria. *Journal of the ACM (JACM) 63,* 4, 36.

Babichenko, Y. and Barman, S. 2015. Query complexity of correlated equilibrium. *ACM Transactions on Economics and Computation 3,* 4, 22.

Babichenko, Y., Barman, S., and Peretz, R. 2016. Empirical distribution of equilibrium play and its testing application. *Mathematics of Operations Research 42,* 1, 15–29.

Babichenko, Y., Dobzinski, S., and Nisan, N. 2019. The communication complexity of local search.

Babichenko, Y. and Rubinstein, A. 2017. Communication complexity of approximate nash equilibria. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing.* ACM, 878–889.

Babichenko, Y. and Rubinstein, A. 2019. Communication complexity of (mixed) nash equilibrium in potential games. *In submission.*

Behnezhad, S., Blum, A., Derakhshan, M., Hajiaghayi, M., Papadimitriou, C. H., and Seddighin, S. 2019. Optimal strategies of blotto games: Beyond convexity. In *Proceedings of the 20th ACM conference on Economics and Computation.* ACM.

[16]The complexity of an ϵ-Nash equilibrium as a function of the constant values of ϵ is far from being understood also in the computational complexity model.

Blum, A. and Monsour, Y. 2007. Learning, regret minimization, and equilibria.

Brassard, G. 2001. Quantum communication complexity (a survey). *arXiv preprint quant-ph/0101005*.

Cesa-Bianchi, N. and Lugosi, G. 2006. *Prediction, learning, and games.* Cambridge university press.

Chen, X., Cheng, Y., and Tang, B. 2017. Well-supported vs. approximate nash equilibria: Query complexity of large games. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017).* Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Chen, X. and Deng, X. 2008. Matching algorithmic bounds for finding a brouwer fixed point. *Journal of the ACM (JACM) 55,* 3, 13.

Conitzer, V. and Sandholm, T. 2004. Communication complexity as a lower bound for learning in games. In *Proceedings of the twenty-first international conference on Machine learning.* ACM, 24.

Czumaj, A., Deligkas, A., Fasoulakis, M., Fearnley, J., Jurdziński, M., and Savani, R. 2019. Distributed methods for computing approximate equilibria. *Algorithmica 81,* 3, 1205–1231.

Deckelbaum, A. 2014. Can quantum entanglement implement classical correlated equilibria? *Quantum Information & Computation 14,* 5-6, 493–516.

Fearnley, J., Gairing, M., Goldberg, P. W., and Savani, R. 2015. Learning equilibria of games via payoff queries. *The Journal of Machine Learning Research 16,* 1, 1305–1344.

Foster, D. P. and Young, H. P. 2006. Regret testing: learning to play Nash equilibrium without knowing you have an opponent. *Theoretical Economics 1,* 3 (Sept.), 341–367.

Ganor, A. and Karthik, C. 2018. Communication complexity of correlated equilibrium with small support. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018).* Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Ganor, A., Karthik, C., and Domotor, P. 2019. On the communication complexity of fixed point computation. *In submission*.

Germano, F. and Lugosi, G. 2007. Global nash convergence of foster and young's regret testing. *Games and Economic Behavior 60,* 1, 135–154.

Goldberg, P. W. and Pastink, A. 2014. On the communication complexity of approximate nash equilibria. *Games and Economic Behavior 85*, 19–31.

Goldberg, P. W. and Roth, A. 2016. Bounds for the query complexity of approximate equilibria. *ACM Transactions on Economics and Computation (TEAC) 4,* 4, 24.

Göös, M. and Pitassi, T. 2014. Communication lower bounds via critical block sensitivity. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing.* ACM, 847–856.

Göös, M., Pitassi, T., and Watson, T. 2015. Deterministic communication vs. partition number. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on.* IEEE, 1077–1088.

Göös, M., Pitassi, T., and Watson, T. 2017. Query-to-communication lifting for bpp. *arXiv preprint arXiv:1703.07666*.

Göös, M. and Rubinstein, A. 2018. Near-optimal communication lower bounds for approximate nash equilibria. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS).* IEEE, 397–403.

Hart, S. 2005. Adaptive heuristics. *Econometrica 73,* 5, 1401–1430.

Hart, S. 2008. Discrete colonel blotto and general lotto games. *International Journal of Game Theory 36,* 3-4, 441–460.

Hart, S. and Mansour, Y. 2010. How long to equilibrium? the communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior 69,* 1, 107–126.

Hart, S. and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica 68,* 5, 1127–1150.

Hart, S. and Mas-Colell, A. 2001. A general class of adaptive strategies. *Journal of Economic Theory 98,* 1, 26–54.

HART, S. AND MAS-COLELL, A. 2003. Uncoupled dynamics do not lead to nash equilibrium. *American Economic Review 93,* 5, 1830–1836.

HART, S. AND MAS-COLELL, A. 2006. Stochastic uncoupled dynamics and nash equilibrium. *Games and Economic Behavior 57,* 2, 286–303.

HART, S. AND MAS-COLELL, A. 2013. *Simple adaptive strategies: from regret-matching to uncoupled dynamics.* Vol. 4. World Scientific.

HART, S. AND NISAN, N. 2018. The query complexity of correlated equilibria. *Games and Economic Behavior 108*, 401–410.

HIRSCH, M. D., PAPADIMITRIOU, C. H., AND VAVASIS, S. A. 1989. Exponential lower bounds for finding brouwer fix points. *J. Complexity 5,* 4, 379–416.

HUBÁČEK, P. AND YOGEV, E. 2017. Hardness of continuous local search: Query complexity and cryptographic lower bounds. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms.* SIAM, 1352–1371.

IMMORLICA, N., KALAI, A. T., LUCIER, B., MOITRA, A., POSTLEWAITE, A., AND TENNENHOLTZ, M. 2011. Dueling algorithms. In *Proceedings of the forty-third annual ACM symposium on Theory of computing.* ACM, 215–224.

JIANG, A. X. AND LEYTON-BROWN, K. 2015. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior 91*, 347–359.

KALAI, E. AND LEHRER, E. 1993. Rational learning leads to nash equilibrium. *Econometrica: Journal of the Econometric Society*, 1019–1045.

KUSHILEVITZ, E. AND NISAN, N. 1997. Communication complexity. In *Advances in Computers.* Vol. 44. Elsevier, 331–360.

LIPTON, R. J., MARKAKIS, E., AND MEHTA, A. 2003. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic commerce.* ACM, 36–41.

LITTLESTONE, N. AND WARMUTH, M. K. 1994. The weighted majority algorithm. *Information and computation 108,* 2, 212–261.

MCLENNAN, A. AND TOURKY, R. 2005. From imitation games to kakutani. *Manuscript, available at http://www.econ.umn.edu/ mclennan/Papers/papers.html.*

MONDERER, D. AND SHAPLEY, L. S. 1996. Potential games. *Games and economic behavior 14,* 1, 124–143.

NASH, J. 1951. Non-cooperative games. *Annals of mathematics*, 286–295.

NISAN, N. 2009. The computational complexity of pure nash. In *Turing's Invisible Hand blog, https://agtb.wordpress.com/2009/11/19/the-computational-complexity-of-pure-nash.*

NITZAN, N. 2005. Tight correlated equilibrium. Tech. rep.

PAPADIMITRIOU, C. H. AND ROUGHGARDEN, T. 2008. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM) 55,* 3, 14.

RAZ, R. AND MCKENZIE, P. 1999. Separation of the monotone NC hierarchy. *Combinatorica 19,* 3, 403–435.

ROUGHGARDEN, T. AND WEINSTEIN, O. 2016. On the communication complexity of approximate fixed points. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS).* IEEE, 229–238.

RUBINSTEIN, A. 2016. Settling the complexity of computing approximate two-player nash equilibria. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS).* IEEE, 258–265.

RUBINSTEIN, A. 2018. Will quantum communication make us more cooperative? In *Theory Dish, Stanford's CS Theory Research Blog, https://theorydish.blog/2018/06/26/will-quantum-communication-make-us-more-cooperative/.*

SHMAYA, E. 2012. Brouwer implies nash implies brouwer. In *The Leisure of the Theory Class, http://theoryclass.wordpress.com/2012/01/05/brouwer-implies-nash-implies-brouwer/.*

VIOSSAT, Y. 2008. Is having a unique equilibrium robust? *Journal of Mathematical Economics 44,* 11, 1152–1160.

YOUNG, H. P. 2004. *Strategic learning and its limits.* OUP Oxford.

YOUNG, H. P. 2009. Learning by trial and error. *Games and economic behavior 65,* 2, 626–643.