Electronic Colloquium on Computational Complexity, Report No. 71 (2017)

Bounds for the Communication Complexity of Two-Player Approximate Correlated Equilibria

Young Kun Ko *   Ariel Schvartzman †

April 13, 2017

Abstract

In the recent paper of [BR16], the authors show that, for any constant $10^{-15} > \varepsilon > 0$ the communication complexity of ε-approximate Nash equilibria in 2-player $n \times n$ games is $n^{\Omega(\varepsilon)}$, resolving the long open problem of whether or not there exists a polylogarithmic communication protocol. In this paper we address an open question they pose regarding the communication complexity of 2-player ε-approximate correlated equilibria.

For our upper bounds, we provide a communication protocol that outputs a ε-approximate correlated equilibrium after exchanging $\tilde{O}(n\varepsilon^{-2})$ bits, saving over the naive protocol which requires $O(n^2)$-bits. This is in sharp contrast to Nash equilibria where for sufficiently small constant ε, no $o(n^2)$-communication protocol is known. In the m-player, n-action setting, our protocol can be extended to a $\tilde{O}(nm)$-bit protocol.

For our lower bounds, we exhibit a simple two player game that has a logarithmic information lower bound: for any constant $\varepsilon < \frac{1}{8}$ the two players need to communicate $\Omega(\log n)$-bits of information to compute any ε-correlated equilibrium in the game. For the m-players, 2-action setting we show a lower bound of $\Omega(m)$ bits, which matches the upper bound we provide up to polylogarithmic terms and shows that the dependence on the number of players is unavoidable.

*Department of Computer Science, Princeton University, ykk@cs.princeton.edu.
†Department of Computer Science, Princeton University, acohenca@cs.princeton.edu.

ISSN 1433-8092

1 Introduction

While Nash equilibria are arguably the most studied notion of equilibrium in strategic games, recent results regarding their communication and computational complexity have undermined their prevalence as a predictable solution concept when agents are computationally bounded. In particular, these results show that two players cannot converge to any approximate Nash equilibria in the limited communication setting where each player only knows its utility function. While there have been multiple attempts to produce procedures that converge to Nash equilibria of general games [HM06, GL07, FY03], it has been shown that at least $\exp(m)$ bits of communication are required to compute Nash equilibria of m-player, constant action games [HM07]. For the case of 2-player games, it has been recently shown that even computing approximate Nash equilibria requires poly(n) bits of communication [BR16].

In addition even in the setting where all the payoffs matrices are known Nash equilibria seem to be unnatural due to their computational hardness. Computing any exact Nash equilibrium is known to be PPAD-complete, making it unlikely to have any polynomial time algorithm [CDT09, DGP09]. Furthermore, it has been shown that under the Exponential Time Hypothesis (ETH) for the class PPAD, ε-approximate Nash equilibria cannot be computed in time faster than quasi-polynomial in the number of strategies per player [Rub16]. This almost exactly matches the algorithm of [LMM03]. The picture becomes more bleak when we consider m-player games. In this case, the problem of even approximating Nash equilibria becomes PPAD-complete [Rub14]. These results suggest that approximate Nash equilibria may not be efficiently computable.

Correlated equilibria arise as an alternative equilibrium concept. This notion, introduced in the seminal work of [Aum74], allows agents to cooperate in order to reach stability. Informally, a strategy profile is a correlated equilibrium when a referee or trusted party can draw strategy samples according to it and recommend them to the players in such a way that they have no incentive to consistently deviate, assuming everyone else plays according to their recommendation. Computationally, correlated equilibria are in sharp contrast to Nash equilibria: there exists an ellipsoid-based algorithm to compute **exact** correlated equilibrium **in polynomial time** even for multi-player games [PR08], for a large (but not universal) class of games including graphical games, anonymous games, congestion games and scheduling games. Unfortunately this result is still unsettling: one can imagine many settings where a referee may not have access to all utility functions or where players may not want to share such information with a referee.

This is indeed comparable to many interesting communication or distributed computation problems where if one party knows all parts of the input, it is easy to compute the output (e.g. disjointness [BGPW13], equality, gap Hamming distance [CR11]). In particular, **the hardness comes only from the distributional nature of the input not the computational aspect,** unlike Nash equilibria.

With this in consideration it becomes more natural to ask whether there is a communication protocol for computing correlated equilibria better than the naive one where each player sends their payoff matrix to a referee who computes the answer. In the case of exact correlated equilibria for 2-player games a simple reduction from the distributed version of linear programming shows that sending the full payoff matrix is indeed optimal [CS89]. For the exact or approximate m-player, constant action case there are simple procedures that converge quickly and use at most polynomial communication in the natural parameters of the input [HMC00, CL03, CBL06].

In this paper we address the question posed by [BR16] of settling the communication complexity of approximate correlated equilibrium. We make progress in both providing non-trivial protocols and deducing non-trivial lower bounds. The arguments used on the lower bound proofs rely on tools from information theory, which lower bounds communication complexity.

1.1 Our results

Upper Bounds. Our upper bounds are similar in spirit to that of [HM10, GR16]. The protocol we provide is based on a non-adaptive no-regret algorithm by [HMC00]. Unfortunately, this protocol converges to a different notion of approximate correlated equilibrium and **assumes that the number of actions per player is constant**. We overcome both of these barriers with simple arguments using dimension reduction [JL84], avoiding a huge loss in the analysis caused by the dimension of the regret matrix. Our result works for general games and has strong implications for the case of 2-player n-action games and m-player 2-action games. In particular for 2-player n action games, our protocol saves a factor of n over sending the whole payoff matrix.

Theorem 1. *There exists a communication protocol Π such that for any m-player n-action game G_n^m, the players compute a ε-CE after exchanging at most $\tilde{O}(nm\varepsilon^{-2})$ bits.*

Lower Bounds. Our lower bound is similar at a first glance to that of [BR16], but our techniques differ significantly due to the nature of the solution concepts studied. As it is pointed out in [BR16] the hardness of proving lower bounds for equilibria lies in being able to hide the solutions (which, by [Nas51, Aum74], must exist). But unlike computing Nash equilibria, which is equivalent to computing fixed points of continuous functions, such equivalence for hard problems under oracle model becomes vague for correlated equilibrium, in part due to its more general nature. This obstacle becomes clearer when we consider the communication complexity of computing correlated equilibrium.

Even in this setting we exhibit a hard game in which $\Omega(\log n)$ bits of communication must be exchanged for two players to agree on an approximate equilibrium. In the m-player 2-action setting, we prove a linear lower bound in the number of players. Note that this proves near-optimality of Theorem 1 and shows that the dependence of the number of players is unavoidable.

Theorem 2. *There exists a 2-player n-action game $G \in \mathbb{G}_n^2$ such that $\Omega(\log n)$ bits of communication are required for the players to agree on a ε-CE with high probability for some small $\varepsilon = \Omega(1)$.*

Theorem 3. *There exists a m-player 2-action game $G \in \mathbb{G}_2^m$ such that $\Omega(m)$ bits of communication are required for the players to agree on a ε-CE with high probability for $\varepsilon < 1/3$.*

1.2 Related Work

The communication complexity of predictable solution concepts has gained a lot of attention and by now most problems pertaining exact and approximate Nash equilibria are well understood. It is known that the communication complexity of computing pure Nash equilibria in 2-player n-action games is poly(n) [CS04]. For m-player binary action games the complexity escalates to $\exp(m)$, even if we relax the solution concepts from pure to exact Nash or correlated equilibria [HM10]. These results were extended to the case of approximate Nash equilibria. In particular, [BR16] showed that the randomized communication complexity of computing ε-Nash equilibria in 2-player n-action games and m-player binary action games is $\Omega(n^\varepsilon)$ and $2^{\Omega(\varepsilon m)}$ for some constant $\varepsilon > 0$.

Some results are known for the communication complexity of computing correlated equilibria for the family of m-player binary action games with bounded, integer payoffs. There is a protocol for the family that computes correlated equilibria after exchanging polynomially many bits in terms of n and the magnitude of the payoffs [HM10]. The former is based on the polynomial time algorithm for computing correlated equilibria by [PR08] and the later is based on a no-regret learning algorithm by [CBL06]. Our results improve upon these for the approximate case by shaving a factor of n and removing the dependence and constraints on the payoff function. It is worth noting that in the same paper they exhibit a family of multiplayer games that do not need to communicate at all to find exact correlated equilibria.

Query Complexity. Another lens through which to consider the cost of computing equilibria is that of query complexity. In this model, a single agent has black box access to the payoff function and can query it on either pure strategies or mixed strategies. A long line of work [FGGS13, HN13, Bab14, Rub16] has recently established that the query complexity of computing approximate 2-player n-action Nash equilibria and approximate m-player 2-action Nash equilibria is $\mathrm{poly}(n)$ and $\exp(m)$, respectively, even for randomized algorithms. For approximate m-player correlated equilibria, there is an exponential gap between the best randomized algorithms and the deterministic lower bounds [HN13, Bab14].

In the case of correlated equilibria (and coarse correlated equilibria), [GR16] show that for m-player binary action games and for any $\varepsilon < 1/2$ the query complexity is $\Theta(\log m)$. They provide an algorithm based on multiplicative weights that uses $\tilde{O}(nm\varepsilon^{-2})$ queries to compute ε-coarse correlated equilibria in m-player n-action games.

Recently and independently of this work [AG17] have shown a lower bound of $\Omega(n)$ for the randomized communication complexity of approximate correlated equilibria in the domain where $\varepsilon < \frac{1}{\mathrm{poly}(n)}$. Their techniques and ideas are similar to ours, except their reduction is directly from the disjointness problem whereas our analysis is based on information-theoretic arguments concerning the games we propose.

1.3 Future Directions

Though it closes the gap for m-player constant action correlated equilibria, our result leaves open the gap for the communication complexity of approximate 2-player n-action correlated equilibria. We are confident that our approach can provide non-trivial communication and query complexity lower bounds of $\Omega(\varepsilon^{-1}\log)$ and $\Omega(n\varepsilon^{-1})$. We conjecture that the right bounds are $\mathrm{poly}(\log n)$ and $n\mathrm{poly}(\log n)$.We share some future directions that might help in settling this question.

- It is known that approximate correlated equilibria with small supports exist [BBP13] with size $O(\log^2 n)$. There is a small gap with the best lower bounds ($O(\log n)$). If the upper bounds were tight, we suspect our techniques could raise lower bounds.

- There exist algorithms to compute approximate correlated and coarse correlated equilibria [BBP13, HN13]. However they either rely on computing exact correlated equilibria, which is prohibitively expensive in the communication setting, or require polynomially many rounds, which already brings the communication cost above our conjectured answer. Progress in algorithms that are distributed in nature and exploit the structure of the solutions could improve on the cost of the protocol we propose.

- Not much is known about the query complexity of 2-player n-action approximate correlated equilibria. The folklore lower bound of $\Omega(n)$ from games with dominant strategies is significantly far from the trivial upper bound of $O(n^2)$. Recent connections between lower bounds in query complexity and lower bounds in communication complexity [GPW15, Göö15] suggest that strong query complexity lower bounds could provide better communication complexity lower bounds.

2 Preliminaries

2.1 Game Theoretic definitions

We consider m-player n-action games where each player has a strategy set A_i and a payoff function $u_i : A \to [0,1]$, where $A = \prod_i A_i$. Let $A^{-i} = \prod_{j \neq i} A_j$. In 2-player games we will refer to the first player as Alice and the second player as Bob.

In this paper we will be interested in studying approximate correlated equilibria (CE), approximate coarse correlated equilibria (CCE) and a different relaxation of exact correlated equilibria due to [HMC00], which we will refer to as approximate Hart-Mas-Colell Correlated Equilibria (HMCE).

A common interpretation of ε-CCE is that a referee draws a strategy profile $a \in A$ according to the correlated distribution x and recommends action a_i to player i. The distribution is a ε-CCE if no single pure deviation from the recommended action yields a benefit greater than ε for any player. There is a similar interpretation for the case of ε-CE: a distribution is a ε-CE if any deviation from the recommended action does not yield a benefit greater than ε for any player. A ε-HMCE only requires that no player benefits more than ε by changing a single recommendation by any other action. We now formally define them in terms of regret, in accordance to [BBP13] (for more standard, equivalent definitions, see e.g. [HN13]).

Definition 1. *Let $R^i_j(a) = u_i(j, a_{-i}) - u_i(a)$ be the regret of player i for playing switching rule j at strategy profile a. A distribution $x \in \Delta(A)$ is an ε-*coarse correlated equilibrium *if $\mathbb{E}_{a \sim x}[R^i_j(a)] \leq \varepsilon$ for all players i and actions $j \in A_i$.*

Definition 2. *Let $R^i_f(a) = u_i(f(a_i), a_{-i}) - u_i(a)$ be the regret of player i for playing switching rule f at strategy profile a. A distribution $x \in \Delta(A)$ is an ε-*correlated equilibrium *if $\mathbb{E}_{a \sim x}[R^i_f(a)] \leq \varepsilon$ for all players i and switching rules $f : A_i \to A_i$.*

Definition 3. *A distribution $x \in \Delta(A)$ is an ε-Hart-Mas-Colell correlated equilibrium if for every player i, every action $j \in A_i$ and every action $k \in A_i$, $\mathbb{E}_{a_{-i} \sim x_{-i}}[R^i_k(a) | a_i = j] \leq \varepsilon$.*

The second and third definitions are relaxations of the definition of exact ($\varepsilon = 0$) correlated equilibria. However, as noted in [BBP13], approximate HMCE are uninteresting to study from a communication perspective. For any game there exists a 0-communication that produces $\frac{1}{k}$-HMCE: independently of the payoff functions the players can agree on a set of k strategies in $\Delta(A)$ and directly output a uniform distribution over them, where $\frac{1}{\varepsilon} \leq k \leq n$. It is not hard to see that this is indeed a $\frac{1}{k}$-HMCE. The advantage of working with this definition is that there exists a non-adaptive no-regret learning algorithms to compute such ε-equilibria for m-player games with a constant number of actions in a number of rounds polynomial in $1/\varepsilon$ [HMC00]. We adapt the algorithm into a communication protocol and revisit their analysis with the consideration that the number of actions per player is part of the input.

2.2 Communication Complexity definitions

In the classical communication problems there are m parties each of which are given inputs $x_i \in \{0,1\}^n$ $x = (x_1, x_2, ..., x_m)$ and who are interested in computing a joint function of their inputs, $f(x)$, where $x = (x_1, x_2, ..., x_m)$. The (randomized) communication complexity of a protocol Π for computing the function $f(x)$ is the (expected) number of bits the two parties need to exchange to compute $f(x)$ by following Π (with high probability). This quantity will be referred to as $\mathrm{CC}(\Pi, f, x)$. The communication complexity of protocol Π for computing f is the worst-case communication complexity for any pair of inputs, i.e. $\mathrm{CC}(\Pi, f) = \max_x \mathrm{CC}(\Pi, f, x)$. The communication complexity of a function f is the minimum communication complexity over all protocols that compute f, $\mathrm{CC}(f) = \min_\Pi \mathrm{CC}(\Pi, f)$.

We will be interested in computing ε-CE, ε-CCE, ε-HMCE of general games $G = (A, u)$ belonging to the family of m-player n-action games $\mathbb{G}_n^m$ with bounded payoff functions. We assume each player only has access to their payoff function u_i. We consider protocols where for every round t, each player communicates as many bits as it wants. We say that Π is a protocol for computing ε-CE of the game G if there exists a number of rounds T after which one of the players outputs a distributions $x \in \Delta(A)$ that forms a ε-correlated equilibria with high probability. We let $\mathrm{CC}(\varepsilon - \mathrm{CE}, \mathbb{G}_n^m) = \min_\Pi \mathrm{CC}(\Pi, \varepsilon - \mathrm{CE}, \mathbb{G}_n^m) = \min_\Pi \max_{G \in \mathbb{G}_n^m} \mathrm{CC}(\Pi, \varepsilon - \mathrm{CE}, G)$. We can analogously define the communication complexity of computing ε-CCE and ε-CE.

2.3 Information Theoretic definitions

Our communication lower bounds are actually based on information theoretic results, so here we provide the tools that will be used in Section 3. Throughout the paper log is the logarithm in base 2 and ln is the natural logarithm. For further references, we refer the reader to [CT12].

Definition 4 (Entropy)**.** *The* entropy *of a random variable A, denoted by $H(A)$ is defined as*

$$\sum_{a \in \mathsf{Supp}(A)} \Pr[A = a] \log \frac{1}{\Pr[A = a]}.$$

Intuitively this quantifies how much uncertainty we have about variable A. This can be extended to define the relation between various variables. For instance suppose we have possibly correlated random variable A and B. Then we can define *conditional entropy* of A given B as $H(A|B) := H(AB) - H(B)$. Note that if $A = B$, the conditional entropy is 0. We formalize this dependency as mutual information.

Definition 5 (Mutual Information)**.** *The* mutual information *between two random variable A and B, denoted by $I(A;B)$ is defined as*

$$I(A;B) := H(A) - H(A|B) = H(B) - H(B|A).$$

The conditional mutual information *between A and B given C, denoted by $I(A;B|C)$, is defined as*

$$I(A;B|C) := H(A|C) - H(A|BC) = H(B|C) - H(B|AC).$$

This quantity measures how much information the random variable B reveals about A and vice-versa (even conditioned on the value of C). Mutual Information is fundamental to the following metric between probability distributions.

Definition 6 (Kullback-Leibler Divergence)**.** *Given two probability distributions μ_1 and μ_2 on the same sample space Ω such that $(\forall \omega \in \Omega)(\mu_2(\omega) = 0 \Rightarrow \mu_1(\omega) = 0)$, the* Kullback-Leibler Divergence *(KL-Divergence for short) between μ_1 and μ_2 is defined as (also known as relative entropy)*

$$D(\mu_1 || \mu_2) = \sum_{\omega \in \Omega} \mu_1(\omega) \log \frac{\mu_1(\omega)}{\mu_2(\omega)}.$$

We now provide useful properties that will be relevant to our proofs. For KL-divergence we use the following properties.

Fact 1. *For random variables A, B and C we have*

$$I(A; B|C) = \mathbb{E}_{B,C}\left[D(A|B, C||A|C)\right].$$

where $A|B, C$ is the distribution of random variable A conditioned on $B = b, C = c$ and similarly for $A|C$.

Fact 2 (Chain Rule for KL-Divergence)**.** *Consider two distributions $P(x_1, \ldots, x_n)$ and $Q(x_1, \ldots, x_n)$. Then*

$$D(P||Q) \geq \sum_{i=1}^{n} D(P(x_i)||Q(x_i)).$$

Definition 7 (Information Complexity)**.** *The* Information Cost *of a 2-party protocol* Π *that computes f is defined as*

$$IC(\Pi) = I(\Pi; A|B) + I(\Pi; B|A),$$

where A is the input to the first party and B is the to the second party. The information cost of f is simply the minimum information cost over all protocols that compute f.

It is easy to check that $\mathrm{CC}(f) \geq \mathrm{IC}(f)$, since 1-bit can carry at most 1-bit of information.

3 Lower Bounds

3.1 Warm up: Best CE requires $\Omega(\log n)$ bits of communication

Let BCE (”best” CE) be the problem of computing the largest-welfare correlated equilibrium. In this section we exhibit a simple bimatrix game that requires at least $\Omega(\log n)$ bits of (randomized) communication for players to agree on the BCE. Our proof relies on an information-theoretic argument, which always lower bounds communication. We show this weak lower bound in order to motivate and ease the main result of this section, the proof of Theorem 2.

Theorem 4. *There exists a game $G \in \mathbb{G}_n^2$ such that $CC((BCE, G)) = \Omega(\log n)$.*

Proof. Consider the following permutation game G_P: Alice and Bob get random $n \times n$ permutation matrices A, B each with the promise that there is exactly one entry such that $A_{ij} = B_{ij} = 1$. There is a unique largest-welfare correlated equilibrium: that in which Alice plays i, Bob plays j and they obtain a combined welfare of 2. We show that $\Omega(\log n)$ bits of communication are required through an information-theoretic argument. We bound the information cost from Bob's perspective, $I(\Pi; A|B)$. This suffices due to the symmetry of the construction. The information cost of game will be twice that from Bob.

$$I(\Pi; A|B) = E_{\Pi,B}\left[D(A|\Pi, B||A|B)\right]. \tag{1}$$

Let F be the distribution of the index i such that $A_{i,j} = B_{i,j} = 1$. Then we can rewrite (1) as

$$\begin{aligned} E_{\Pi,B}\left[D(A|\Pi, B||A|B)\right] &= E_{\Pi,B}\left[D(F|\Pi, B||F|B)\right] + E_{\Pi,F,B}\left[D(X|\Pi, F, B||X|F, B)\right], \\ &\geq E_{\Pi,B}\left[D(F|\Pi, B||F|B)\right] \end{aligned}$$

where the inequality holds due to non-negativity of KL divergence. We know $(F|\Pi, Y)$ must be a distribution that concentrates all of it's mass on i and $F|Y$ is a uniform distribution is simply the uniform distribution over n elements. Therefore we get that $D(F|\Pi, B||F|B) = \log n$. Similarly, we have $I(\Pi; B|A) \geq \log n$. Thus $\mathrm{CC}((\mathrm{BCE}, G)) \geq \mathrm{IC}((\mathrm{BCE}, G)) = \Omega(\log n)$. $\square$

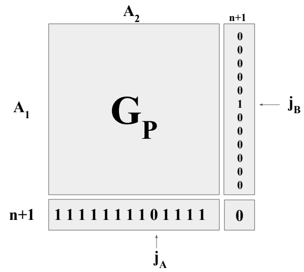

Figure 1: An example of what the modified payoff matrix looks like for Alice.

Even though this bound is possibly not tight it will carry the main intuition of Section 3.2. It might be tempting to think that the same game should work for the general ε-CE case. Unfortunately there is a simple 0-communication protocol: both players can immediately output the uniform distribution over the strategies. No player can gain more than $\frac{1}{n}$ by consistently deviating. Since we are only concerned with the domain in which $n > \frac{1}{\varepsilon}$ this protocol correctly outputs a $\varepsilon-$CE.

3.2 Unrestricted lower bound for G_n^2

It turns out that a simple modification of the game suffices to carry the result from BCE to ε-CE. We add a small game on the side which dissuades from large support, 0-communication strategies. In particular, we show that for $\varepsilon < \frac{1}{8}$ any ε-CE must allocate most of its mass on the original game. This allows us to give a simple argument similar to that in the previous section.

Construction Consider the game G_P from the proof of Theorem 4 with a slight adjustment: give each player an additional action $n+1$. Choose $j_a, j_b \in [n]$ independently at random. We refer to G_P as the main part of the game and the remainder as the auxiliary part of the game.

For $i \neq j_a$ make $u_A(n+1, i) = 1$ and 0 otherwise. For $j \neq j_b$, make $u_B(j, n+1) = 1$ and 0 otherwise. Make $u_A(j_b, n+1) = u_B(n+1, j_a) = 1$ and $u_A(i, n+1) = 0$ for all $i \in [n+1] \backslash \{j_b\}$, and $u_B(n+1, j) = 0$ for all $j \in [n+1] \backslash \{j_a\}$ (see Figure 2 for an example).

As a simple exercise note that after this amendment the uniform distribution is no longer a ε-correlated equilibrium for $\varepsilon < \frac{1}{2}$. Any player can unilaterally switch to the new strategy and gain $\frac{1}{2} - \frac{1}{n}$ from deviating.

Claim 1. *For any ε-CE, there is at least $\Omega(1)$ fraction of the mass on the main part of the game. In particular, no ε-CE can have $1 - o(1)$ fraction of the mass on the auxiliary part of the game.*

Proof. We consider Alice's perspective. The argument is symmetric for Bob. Let p_0 be the mass this player assigns to $j_b \times (n+1)$, p_1 the mass on $([n] \backslash j_b) \times (n+1)$, and p_2 the mass on $(n+1) \times (n+1)$.

By construction, $u_B(j_b, n+1) = u_B(n+1, n+1) = 0$ so the payoff for Bob, conditioned on this column, is p_1. By deviating to $k \in [n]$ such that $u_B(j_b, k) = 1$, Bob can gain a payoff of p_0. On the other hand, by deviating to $j_a \in [n]$, Bob will gain a payoff of at least p_2, since $u_B(n+1, j_a) = 1$. So any ε-CE needs to satisfy the following set of inequalities.

$$
\begin{aligned}
p_1 &> p_0 - \varepsilon, \\
p_1 &> p_2 - \varepsilon, \\
p_0 &+ p_1 + p_2 = 1.
\end{aligned}
$$

Combining them, we have that

$$2p_1 > p_0 + p_2 - 2\varepsilon = 1 - p_1 - 2\varepsilon. \tag{2}$$

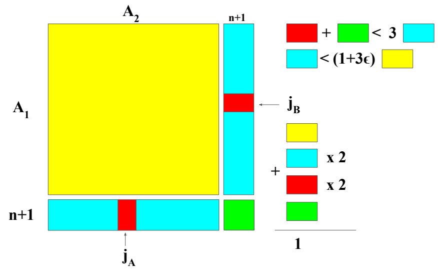

Figure 2: Coloring argument.

Rearranging (2), we get that

$$p_1 > \frac{1-2\varepsilon}{3} > \frac{1}{4}. \tag{3}$$

Then (3) implies that the total mass on $[n]\times(n+1)$ is at most 4 times the mass on $([n]\backslash j_b)\times(n+1)$. Symmetrically, we get that total mass on $(n+1)\times[n]$ is at most 4 times the mass on $(n+1)\times([n]\backslash j_a)$.

Let M be the mass on G_P, M_b as the mass on $([n]\backslash j_b)\times[n]$ and M_i as the mass on $i\times[n]$. Note that for any recommendation $i\in[n]\backslash j_b$ to Alice, the mass on $n+1$ can be at most $1/2+\varepsilon/2$. Otherwise, by deviating to j_b, Alice gains at least ε gain in payoff. In other words, the mass on $i\times(n+1)$ is at most $\frac{1+\varepsilon}{1-\varepsilon}M_i \le (1+3\varepsilon)M_i$. Then the mass on $([n]\backslash j_b)\times(n+1)$ is bounded by

$$(1+3\varepsilon)\sum_{i\in[n]\backslash j_b} M_i = (1+3\varepsilon)M_b < (1+3\varepsilon)M. \tag{4}$$

Combining (3) and (4) we have that mass on $[n+1]\times(n+1)$ and $(n+1)\times[n+1]$ is at most $4(1+3\varepsilon)M$. Moreover, $p_1 > p_2-\varepsilon$. Concretely we have that

$$\begin{aligned}(1+3\varepsilon)M >& p_1,\\ 3p_1 >& p_0+p_2.\end{aligned} \tag{5}$$

The sum of all the masses is 1, which combined with the upper bounds derived and the fact that $\varepsilon < \frac{1}{8}$ yields

$$1 \le 2p_0+2p_1+2p_2+M < (1+8(1+3\varepsilon))\,M < 12M,$$

which is what we wanted to show. ☐

Note that Claim 1 states that for any ε-CE, there is a substantial mass on the main part of the game. Now we are ready to argue that computing ε-CE in the main part of the game requires $\Omega(\log n)$ communication.

Proof of Theorem 2. For each recommendation $i\in[n]$ to Alice on any ε-CE, let $\vec{d^i}$ denote the distribution on the column. Note that the payoff for following the recommendation is $\vec{d^i}\cdot a_i = \vec{d^i}_{\pi_A(i)}$.

Denote $j = \arg\max_{k\in[n+1]} \vec{d^i}_k$, then note that the payoff for deviating in $[n]$ is $\vec{d_i}\cdot a_{\pi_A^{-1}(j)} = \max_{k\in[n+1]}\vec{d^i}_k$ while the payoff for deviating to $n+1$ is $\vec{d_i}\cdot a_{n+1}$. We split into two cases for $\vec{d^i}$: either $\max_{k\in[n+1]}\vec{d^i}_k \ge 1/8$ or $\max_{k\in[n+1]}\vec{d^i}_k < 1/8$.

In the first case we need $\vec{d^i}_{\pi_A(i)} \ge 1/4-\varepsilon > 1/8$ in order for Alice not to deviate. Her prior is the uniform distribution over n strategies since the input distribution is the uniform distribution over permutation matrices. This combined with Fact 1,

$$I(\Pi; Y_i|X) = D(\vec{d^i}||\mathcal{U}_n) = \sum_{k\in[n]} \vec{d^i}_k \log \vec{d^i}_k n \geq \frac{\log n}{8} + H(1/8) = \Omega(\log n),$$

where H refers to the binary entropy function. Thus the information cost is $\Omega(\log n)$. Since information precludes communication, the lower bound for communication follows.

In the other case note that the current payoff is at most $1/8$, while the payoff for deviating to $n+1$ is at least $1/2$. This is a contradiction since $1/2 - 1/8 > \varepsilon$, so this can not happen on a ε-CE. □

Remark 1. *Indeed, one can come up with an easier construction if Alice has a dominant strategy and Bob's objective is to find Alice's dominant strategy. However, this is one-sided in a sense that Bob is the only player who learns from the transcript, compared to our construction where both Alice and Bob learn from the transcript.*

3.3 Unrestricted lower bound for G_2^m

In this section we exhibit a game $G \in \mathbb{G}_2^m$ whose ε-correlated equilibrium communication complexity is $\Omega(m)$, proving Theorem 3.

Construction Without loss of generality suppose m is a even number. Each player is equipped with two actions: 0 and 1. We will refer to the first $m/2$ as 'state setters' and define their payoffs as follows: let $\vec{R} \in \{0,1\}^{m/2}$ be a string of random boolean variables where each coordinate is set independently at random at with probability $1/2$. Then

$$u_i(a_i, \vec{a}_{-i}) = \begin{cases} 1 & \text{if } a_i = r_i \\ 0 & \text{otherwise} \end{cases}$$

We refer to the last $m/2$ players as 'imitators', and define their payoffs as follows:

$$u_i(a_i, \vec{a}_{-i}) = \begin{cases} 1 & \text{if } \vec{a}_{i-m/2} = a_i \\ 0 & \text{otherwise} \end{cases}$$

In particular, imitators only get a positive payoff if they manage to successfully imitate their corresponding state setters. We now characterize the structure of ε-CE for this game.

Claim 2. *For any ε-CE, state setter i must have $> 1-\varepsilon$ mass on the recommended action r_i.*

Proof. The payoff for playing r_i with probability p is p. Hence putting all the mass in this action gives a payoff of 1. Therefore for any ε-CE, $p > 1-\varepsilon$. □

With a similar proof, we can also show the following:

Claim 3. *For any ε-CE, imitator i must have mass $> 1-\varepsilon$ on its setter's action, $r_{i-m/2}$.*

Proof of Theorem 3. Suppose we merge all imitators to one player, Bob, and all state setters to the other player, Alice. This only makes the communication task easier since for any m-player protocol, these players can just simulate their own $m/2$ players. Then let A denote the input for state setters and B denote the input for imitators, then indeed $I(\Pi; B|A) = 0$. But

$$I(\Pi; A|B) = I(\Pi; A) = I(\Pi; R)$$

since B is a constant and A is completely determined by R. Then

$$I(\Pi; R) = D(R_\pi || R) \geq \sum_{i \in [m/2]} D(R_{i_\pi} || R_i) \geq m \cdot \frac{1 - H(\varepsilon)}{2},$$

where the equality is Fact 1 and the first inequality stems from Fact 2. Taking ε to be a constant $< 1/3$ we get the desired information lower bound, which in turn implies multi-player communication lower bound. □

4 Upper Bounds for General Games

In this section we present a communication protocol for computing ε-CE based on an adaptive algorithm of [HMC00] that converges to ε-HMCE. The overall communication cost of the protocol is $\tilde{O}(mn)$, where the largest factor comes from having to carry out $O(n)$ rounds in order to guarantee convergence to approximate correlated equilibria. Note that the naive upper bound is $O(mn^m)$ since every player can simply write their utility functions on a common blackboard.

4.1 Warm up: A $O(m \log n)$ protocol for $CC(\varepsilon - \mathbf{HMCE}, \mathbb{G}_n^m)$

We present a protocol based on an algorithm of [HMC00] to compute ε-HMCE. It is worth noting that in their paper they present two algorithms: an adaptive one and a non-adaptive one. As it turns out, both require the same number of rounds to converge to ε-HMCE but the adaptive strategy is technically more challenging to show, which is why we opt for the simpler algorithm.

At a high level, on round t each player looks at the history of strategies $h(t) = \{s'_t : s \in A, (t' < t)\}$ played up until then and computes a matrix that measures the average regret of not having played action k whenever action j was played, for all actions $k, j \in A^i$ from round 1 up to round $t - 1$. They then compute the stationary distribution of this matrix (which [HMC00] shows always exists), play an action according to it and announce it to everyone else. After $T = \frac{1}{\varepsilon^2}$ rounds, the first player (or any player) outputs the empirical distribution of actions played z_T, where for a given $s \in A$ the probability it gets played is $z_T(s) = \frac{1}{T}|\{t \leq T : s = s_t\}|$. The beauty of the procedure is that players only need to communicate the action they perform at the current time period, using $O(\log n)$ bits of communication.

Theorem 5. $CC(\varepsilon - HMCE, \mathbb{G}_n^m) = O(m \log n)$

Proof. We reproduce the algorithm and notation in [HMC00].

The matrix A_t simply counts the regret of not having played action k at time t when action j was played. The matrices D_t, R_t average the regret and ignore actions with negative regret up until time t, respectively. It is clear that the communication cost of the protocol will be $O(mT \log n)$ where T is the number of rounds. The proof in [HMC00] shows that $O(\frac{1}{\varepsilon^2})$ rounds suffice to converge to the set of ε-HMCE. The following corollary concludes the proof.

Corollary 1. *Suppose that at each period $t + 1$ every player i chooses strategies according to the stationary distribution of R_t^i. Then the empirical distribution of plays z_t converges at a rate of $O(\frac{1}{\varepsilon^2})$ almost surely as $t \to \infty$ to the set of exact correlated equilibria.*

□

Algorithm 1 Protocol Π to compute ε-HMCE

- At time $t = 1$, each player plays according to some arbitrary initial distribution p_1^i.
- From $t = 2$ to $t = T$:
- Each player i computes the matrices A_t^i, D_t^i, R_t^i where

$$A_t^i(j,k) = \mathbf{1}_{s_t^i = j}[u^i(k, s_t^{-i}) - u^i(j, s_t^{-i})],$$

$$D_t^i(j,k) = \frac{1}{t}\sum_{\tau=1} A_\tau(j,k),$$

$$R_t^i(j,k) = \max\{0, D_t(j,k)\}.$$

- Each player computes the stationary distribution of R_t^i, p_t^i, plays according to it and announces his move to everyone else.
- Player 1 outputs the empirical distribution z_T of all strategies played.

4.2 From ε-HMCE to ε-CE

There are two issues that prevent us from using the same protocol as a black box for the more general case. In the first place, the algorithm converges to ε-HMCE, not ε-CE. It can be shown that **ε-HMCE are, in the worst case, $\sqrt{n}\varepsilon$-CE and this is tight**. This implies that in order to converge to ε-CE we need at most $\frac{n}{\varepsilon^2}$ rounds of the protocol. The other problem is that the algorithm assumes that the number of actions per player is constant and the dependence in the size of the matrix is hidden in the running time of the procedure. A revision of their analysis gives a protocol whose communication cost is $\tilde{O}(n\varepsilon^{-2})$ per player.

The following lemma characterizes convergence to ε-CE in terms of the regret matrix R_t^i.

Lemma 1. *Consider any sequence of plays and let $\varepsilon \geq 0$. If $\limsup_{t\to\infty} \max_j\{\sum_k R_t^i(j,k)\} \leq \varepsilon$ for all players i and all actions $j \in A^i$, the sequence of empirical distributions converges to the set of ε-CE.*

Proof. For each player i and every $j \in A_i$ we have

$$\sum_{k\neq j} D_T^i(j,k) = \frac{1}{T}\sum_{\tau=1}^{T}\sum_{k\neq j} A_\tau(j,k) = \sum_{s\in A: s_i = j} z_T(s)\sum_{k\neq j}\left(u^i(k,s^{-i}) - u^i(j,s^{-i})\right) = \mathbb{E}_{s\sim z_T}(R_f^i(s)) \leq \varepsilon,$$

for any switching rule f where the last equality follows from the fact that a switching rule is just a linear combination of single deviations. Since the sum of the regret of all individual actions is bounded by ε so is any convex combination of them. ☐

Lemma 2. *If the regret matrix for each player i satisfies $\sqrt{\sum R_T^i(j,k)^2} \leq \varepsilon$ then we obtain a $\sqrt{n}\varepsilon$-CE.*

Proof. By 1 it suffices to show that $\max_k\{R_t^i(j,k)\} \leq \sqrt{n}\varepsilon$.

$$\max_j\{\sum_k R_t^i(j,k)\} \leq \sqrt{n}\max_j\sqrt{\sum_k R_T^i(j,k)^2} \leq \sqrt{n}\sqrt{\sum_{j,k} R_T^i(j,k)^2} \leq \sqrt{n}\varepsilon,$$

where the first inequality follows from the relationship between the ℓ_1 and ℓ_2 norms, and the second follows from the relationship between ℓ_∞ and ℓ_1.

□

4.3 A small detour: Blackwell's Approachability Theorem

Before we proceed to the last component of our analysis of the protocol we will need to introduce Blackwell's Approachability Theorem [Bla56]. The standard setup considers an agent i who plays actions $a_i \in A_i$ and gets payoff vectors in $\mathbb{R}^L$, that depends on his action and another action $a^{-i} \in A^{-i}$ chosen by an opponent, possibly adversarially. In other words, agent i's payoff is of the form $v_i : A^i \times A^{-i} \to \mathbb{R}^L$. The game is played for many rounds and the agents goal is to make her average payoff vector, $D_T = \frac{1}{T}\sum_{t=1}^{T} v_i(a_t^i, a_t^{-i})$, approach some given set $C \in \mathbb{R}^L$. We say a convex, closed set C is approachable if there is a procedure that almost surely guarantees that the distance between D_T, C approaches 0 as T goes to ∞, irrespective of the opponents actions. Blackwell's Approachability Theorem states necessary and sufficient conditions under which this can be done.

Theorem 6. *Let $C \in \mathbb{R}^L$ be a closed, convex set with support function w_C. Then C is approachable if and only if for all $\lambda \in \mathbb{R}^L$ there is a mixed strategy $q_\lambda \in \Delta(A^i)$ such that $\lambda \cdot u_i(q_\lambda, a^{-i}) \leq w_C(\lambda)$, for all $a^{-i} \in A^{-i}$.*

Moreover, the following strategy suffices: at time $t+1$ play $q_{\lambda(D_t)}$ if $D_t \notin C$ and arbitrarily otherwise. The rate of convergence is $O(\frac{1}{\sqrt{T}})$.

[HMC00] use the result directly with $L = n^2$ and the individual payoffs v_i mimicking the entries of the regret matrix R_T^i. They show that the stationary distribution of the regret matrices is exactly the vector q_λ that guarantees convergence in Blackwell's Theorem.

In order to understand how the dimension of the vectors plays a role in this we need to take a close look into the proof of the Theorem. Let ρ_t be the distance between the average payoff vector of agent i and the convex set C. The analysis relies on a recursively bounding ρ_{t+1} as a function of ρ_t. In particular, the recursion looks like this:

$$(t+1)^2\rho_{t+1}^2 \leq t^2\rho_t + ||v_i(a^i, a^{-i}) - \pi_C(D_t)||^2,$$

where $\pi_C(D_T)$ is the projection of the average regret vector at time t onto set C. A standard analysis would bound the rightmost term by a constant (arguing that all points belong to some ball of bounded radius) and, with the use of a telescoping argument, show that the distance converges at a rate of $O(\frac{1}{\sqrt{T}})$. However, in our case, if the dimension is a parameter we care about then an appropriate upper bound on the rightmost term would be $O(L)$ (this is because the vectors lie on L-dimensional space and are entry-wise bounded due to the nature of the utility functions). Carrying the analysis as is would then give a regret of $O(\sqrt{\frac{n}{T}})$, which would in turn significantly blow up the cost of our communication protocol. We claim that projecting these vectors into $O(\log(T))$-dimensional spaces using the Johnson-Lindenstrauss Lemma preserves the validity of the proof at a small cost on the convergence rate.

Claim 4. *After T rounds of the protocol Π the average regret matrix for player i, R_T^i, is such that $\sqrt{\sum R_T^i(j,k)^2} \leq 32\frac{\log T}{\sqrt{T}}$.*

Proof. We can apply the Johnson-Lindenstraus Lemma [JL84] (with $\delta = \frac{1}{2}$)[1] to the T n^2-dimensional vectors that arise from carrying out the telescoping sum in the analysis of Blackwell's Theorem.

[1]The JL states that for any set of m points in $\mathbb{R}^n$ there exists a projection $f : \mathbb{R}^n \to \mathbb{R}^d$, where $d > 8\frac{\log m}{\varepsilon^2}$, that preserves pairwise distances between the points up to a multiplicative ε factor.

This reduces their dimension to $32 \log T$ and approximately preserves the ℓ_2 norm. The rest of the analysis of the recursion carries through and we get an additional $\sqrt{\log(T)}$ factor on the convergence rate. $\square$

Theorem 7. *Protocol Π with $T = 160^2 \frac{n}{\varepsilon^2} \log^2 n$ rounds produces a ε-CE.*

Proof. By Claim 4, $160^2 \frac{n}{\varepsilon^2} \log^2 n$ rounds produce regret matrices such that $\sqrt{\sum R_T^i(j,k)^2} \leq \frac{\varepsilon}{\sqrt{n}}$. Combining this with Lemma 2 finishes the proof. $\square$

References

[AG17] K. C. S. A. Ganor. Communication complexity of correlated equilibrium in two-player games. 2017.

[Aum74] Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, March 1974.

[Bab14] Yakov Babichenko. Query complexity of approximate nash equilibria. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 535–544, 2014.

[BBP13] Yakov Babichenko, Siddharth Barman, and Ron Peretz. Small-support approximate correlated equilibria. *CoRR*, abs/1308.6025, 2013.

[BGPW13] Mark Braverman, Ankit Garg, Denis Pankratov, and Omri Weinstein. From information to exact communication. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 151–160, New York, NY, USA, 2013. ACM.

[Bla56] David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific J. Math.*, 6(1):1–8, 1956.

[BR16] Yakov Babichenko and Aviad Rubinstein. Communication complexity of approximate nash equilibria. *CoRR*, abs/1608.06580, 2016.

[CBL06] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.

[CDT09] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3):14:1–14:57, May 2009.

[CL03] Nicolò Cesa-Bianchi and Gábor Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51(3):239–261, 2003.

[CR11] Amit Chakrabarti and Oded Regev. An optimal lower bound on the communication complexity of gap-hamming-distance. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 51–60, New York, NY, USA, 2011. ACM.

[CS89] J. Chu and G. Schnitger. The communication complexity of several problems in matrix computation. In *Proceedings of the First Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '89, pages 22–31, New York, NY, USA, 1989. ACM.

[CS04] Vincent Conitzer and Tuomas Sandholm. Communication complexity as a lower bound for learning in games. In *Proceedings of the Twenty-first International Conference on Machine Learning*, ICML '04, pages 24–, New York, NY, USA, 2004. ACM.

[CT12] Thomas M Cover and Joy A Thomas. *Elements of information theory.* John Wiley & Sons, 2012.

[DGP09] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[FGGS13] John Fearnley, Martin Gairing, Paul Goldberg, and Rahul Savani. Learning equilibria of games via payoff queries. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 397–414, New York, NY, USA, 2013. ACM.

[FY03] Dean P. Foster and H. Peyton Young. Learning, hypothesis testing, and nash equilibrium. *Games and Economic Behavior*, 45(1):73–96, 2003.

[GL07] Fabrizio Germano and Gábor Lugosi. Global nash convergence of foster and young's regret testing. *Games and Economic Behavior*, 60(1):135–154, 2007.

[Göö15] Mika Göös. Lower bounds for clique vs. independent set. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1066–1076. IEEE, 2015.

[GPW15] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1077–1088. IEEE, 2015.

[GR16] Paul W. Goldberg and Aaron Roth. Bounds for the query complexity of approximate equilibria. *ACM Trans. Econ. Comput.*, 4(4):24:1–24:25, August 2016.

[HM06] Sergiu Hart and Andreu Mas-Colell. Stochastic uncoupled dynamics and nash equilibrium. *Games and Economic Behavior*, 57(2):286–303, 2006.

[HM07] Sergiu Hart and Yishay Mansour. The communication complexity of uncoupled nash equilibrium procedures. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 345–353. ACM, 2007.

[HM10] Sergiu Hart and Yishay Mansour. How long to equilibrium? the communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69(1):107–126, 2010.

[HMC00] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[HN13] Sergiu Hart and Noam Nisan. The query complexity of correlated equilibria. *CoRR*, abs/1305.4874, 2013.

[JL84] William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.

[LMM03] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, EC '03, pages 36–41, New York, NY, USA, 2003. ACM.

[Nas51] John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

[PR08] Christos H Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):14, 2008.

[Rub14] Aviad Rubinstein. Inapproximability of nash equilibrium. *CoRR*, abs/1405.3322, 2014.

[Rub16] Aviad Rubinstein. Settling the complexity of computing approximate two-player nash equilibria. *CoRR*, abs/1606.04550, 2016.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il