The Communication Complexity of Uncoupled Nash Equilibrium Procedures*

Sergiu Hart
Center for the Study of Rationality,
Institute of Mathematics, and
Department of Economics
The Hebrew University of Jerusalem
Jerusalem, ISRAEL
hart@huji.ac.il

Yishay Mansour
School of Computer Science
Tel Aviv University
Tel Aviv, ISRAEL
mansour@tau.ac.il

ABSTRACT

We study the question of how long it takes players to reach a Nash equilibrium in *uncoupled* setups, where each player initially knows only his own payoff function. We derive lower bounds on the *communication complexity* of reaching a Nash equilibrium, i.e., on the number of bits that need to be transmitted, and thus also on the required number of steps. Specifically, we show lower bounds that are exponential in the number of players in each one of the following cases: (1) reaching a pure Nash equilibrium; (2) reaching a pure Nash equilibrium in a Bayesian setting; and (3) reaching a mixed Nash equilibrium. We then show that, in contrast, the communication complexity of reaching a correlated equilibrium is polynomial in the number of players.

General Terms

Theory of Computation, Economics

Categories and Subject Descriptors

F.0 [**Theory – General**]; J.4 [**Economics**]

Keywords

Communication Complexity, Computational Game Theory

*We thank Fabrizio Germano, Adam Kalai, Eyal Kushilevitz, Andreu Mas-Colell, and Noam Nisan for useful discussions, and the Institute of Advanced Studies at the Hebrew University of Jerusalem where some of this work was done. The research was partially supported by grants of the Israel Science Foundation (Hart, Mansour) and by an IBM faculty award (Mansour). The full version of the paper can be found on the authors' websites, `http://www.ma.huji.ac.il/hart` and `http://www.cs.tau.ac.il/~mansour`.

STOC'07, June 11–13, 2007, San Diego, California, USA.

1. INTRODUCTION

Equilibrium is a central concept in interactions between decision-makers. The definition of equilibrium is *static*: it is characterized by the property that the participants ("players") have no incentive to depart from it. No less fundamental, however, are the *dynamic* issues of how such an equilibrium arises (see, e.g., [11, 25]). Since decisions are assumed to be taken independently by the participants, it is only natural to study dynamics in *decentralized* environments, where each decision-maker has only partial information—for instance, he knows only his own preferences and not those of the other players. As a result, no player can find an equilibrium on his own, and the resulting dynamics become complex and need not converge to a rest-point (i.e., an equilibrium).

Significant progress has been made in understanding the dynamic aspects of one equilibrium concept, that of *correlated equilibrium* [1]. A correlated equilibrium obtains when players receive signals before the game is played; these signals, which may be correlated, do not affect the payoffs in the game. Of course, the players may well use these signals when making their strategic choices. To date, there are several efficient algorithms [8, 14, 15, 3, 4, 23, 24, 2, 25, 13] that, in all games, converge fast to (approximate) correlated equilibria.

In contrast, convergence to *Nash equilibrium* is a much more complex and less clear-cut issue.[1] As we have stated above, a natural assumption that most dynamics satisfy is that of *uncoupledness* [16]: each player is assumed to know initially *only his own payoff function*, and not those of the other players.[2] On the one hand, it has been shown that it is impossible for uncoupled dynamics that are deterministic and continuous[3] always to converge to a Nash equilibrium, even when it is unique [16]. On the other hand, there are a number of uncoupled dynamics that converge to Nash equilibria; these dynamics use various techniques such as hypothesis-testing, regret-testing, and other variants of exhaustive or stochastic search [9, 10, 25, 7, 17, 12]. Since

[1]A Nash equilibrium is a fixed point of a nonlinear function, whereas a correlated equilibrium is a solution of finitely many linear inequalities. This may be one reason—though not the only one—that it appears to be more difficult to converge to the former than to the latter.

[2]Similar notions are "distributed" in computer science, "privacy-preserving" in mechanism design, and "decentralized" in economics.

[3]Continuous with respect to both actions and time.

all these dynamics perform some form of search over all action combinations, it follows that the number of steps until a Nash equilibrium is reached is exponential in the number of players (when the number of actions of each player is kept fixed). In the current work we will show that this is a general phenomenon and not a deficiency of the existing literature: *there is an exponential lower bound on the speed of convergence to Nash equilibria.*

To make this precise, define a *Nash equilibrium procedure* as a dynamic process whereby the players reach a Nash equilibrium, whether *pure* or *mixed*.[4] We study the number of steps needed before the procedure terminates at the appropriate equilibrium. Again, we are considering uncoupled procedures: each player's payoff function is private, initially known only to him. We use the theory of *communication complexity* to derive lower bounds on the amount of communication, measured in terms of the number of transmission bits—and thus also the number of steps—that the players have to perform in order to reach a Nash equilibrium. This important connection was first observed in [6], where various lower bounds for *two*-person games are derived (as the number of actions increases). Here we analyze general n-person games.

Our results provide lower bounds that are exponential in the number of players (we keep the number of actions of each player bounded, e.g., two) for the communication complexity in each of the following cases: (1) reaching a pure Nash equilibrium—in general games, and also in the restricted class of games having the "finite improvement property" (Section 3 and Appendix A); (2) reaching a pure Nash equilibrium in a Bayesian setup (Section 4); and (3) reaching a mixed Nash equilibrium (Section 5). In the full version of the paper we exhibit simple procedures that yield upper bounds that are also exponential (for both pure Nash equilibria and mixed Nash equilibria). The proofs omitted from this extended abstract and additional material can also be found there.

These exponential lower bounds may seem unsurprising, given that the size of the input (i.e., the players' private payoff functions) is also exponential. We thus analyze the communication complexity of reaching *correlated equilibria*, and we show that it is, in contrast, only *polynomial* in the number of players (Section 6). Therefore, the exponential communication complexity of Nash equilibrium procedures is a result of the equilibrium requirement, and *not* of the size of the input.

In summary, this paper may be viewed as providing further evidence of the intrinsic difficulty of reaching Nash equilibria, in contrast to correlated equilibria.[5]

2. PRELIMINARIES

2.1 Game-Theoretic Setting

The basic setting is as follows. There are $n \geq 2$ players, $i = 1, 2, \ldots, n$. Each player i has a finite set of actions A_i with $|A_i| \geq 2$, and the joint action space is $A = \prod_{i=1}^n A_i$. Let Δ_i denote the set of probability distributions over A_i and put $\Delta = \prod_{i=1}^n \Delta_i$. Most of the games we introduce will be *binary-action games*, where the action space of each player i is $A_i = \{0, 1\}$, and so $A = \{0, 1\}^n$; in this case a mixed action of player i is given by $0 \leq p_i \leq 1$, interpreted as the probability that $a_i = 1$.

Each player i has a payoff (or utility) function u_i which maps A to the real numbers, i.e., $u_i : A \to \mathbb{R}$. We extend u_i to Δ in a multilinear way, by defining $u_i(p_1, \ldots, p_n) = \mathbf{E}[u_i(a_1, \ldots, a_n)]$ for each $(p_1, \ldots, p_n) \in \Delta$, where the expectation $\mathbf{E}$ is taken with respect to the product distribution $p_1 \times \cdots \times p_n$ on A. We denote this game by $G = (n, \{A_i\}_i, \{u_i\}_i)$.

For a joint action $a = (a_1, \ldots, a_i, \ldots, a_n) \in A$, let a^{-i} be the joint action of all players *except* player i, i.e., $a^{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$. For each player i, the (pure) *best-reply correspondence* maps a joint action a^{-i} of the other players to the set $\mathrm{BR}(a^{-i}; u_i) = \arg\max_{a_i \in A_i} u_i(a_i, a^{-i})$. A joint action $a \in A$ is a *pure Nash equilibrium* if for every player i we have $u_i(a) \geq u_i(b_i, a^{-i})$ for any $b_i \in A_i$; or equivalently, $a_i \in \mathrm{BR}(a^{-i}; u_i)$ for all i. Similarly, a combination of mixed actions $p \in \Delta$ is a *mixed Nash equilibrium* if $u_i(p) \geq u_i(q_i, p^{-i})$ for every player i and any $q_i \in \Delta_i$.

Finally, we define the concepts of "improvement step" and "improvement path." Given a joint action $a \in A$, an *improvement step* of player i is an action $b_i \in A_i$ such that $u_i(b_i, a^{-i}) > u_i(a)$; we refer to i as the *improving player.* An *improvement path* is a sequence of improvement steps (where the improvement steps can be performed by different players). A game G has the *finite improvement property* if all the improvement paths are finite;[6] such a game always possesses a pure Nash equilibrium.

2.2 Communication Complexity Background

In the "classical" setting in communication complexity there are two agents,[7] one holding an input $x \in \{0, 1\}^K$ and the other holding an input $y \in \{0, 1\}^K$, where K is a finite set. Their task is to compute a joint function of their inputs $f(x, y) \in \{0, 1\}$. The agents send messages to one another, and we assume that at the end of the communication they each have the value of $f(x, y)$. The *communication complexity* of a deterministic communication protocol Π for computing $f(x, y)$ is the number of bits sent during the computation of $f(x, y)$ by Π; denote this number of bits by $\mathrm{CC}(\Pi, f, x, y)$. The communication complexity $\mathrm{CC}(\Pi, f)$ of a protocol Π for computing a function f is defined as the worst case over all possible inputs $(x, y) \in \{0, 1\}^K \times \{0, 1\}^K$, i.e., $\mathrm{CC}(\Pi, f) = \max_{x,y \in \{0,1\}^K} \mathrm{CC}(\Pi, f, x, y)$. Finally, the *communication complexity* $\mathrm{CC}(f)$ *of computing a function* f is the minimum over all protocols Π for computing f, i.e., $\mathrm{CC}(f) = \min_\Pi \mathrm{CC}(\Pi, f)$.

A well-studied function in communication complexity is the *disjointness* function. Let $\mathcal{S}$ be a finite set; the $\mathcal{S}$- *disjointness function* $\mathrm{DISJ}_\mathcal{S}$ is defined on the subsets of $\mathcal{S}$ (i.e., on $\{0, 1\}^\mathcal{S} \times \{0, 1\}^\mathcal{S}$) by $\mathrm{DISJ}_\mathcal{S}(S_1, S_2) = 1$ if the two inputs $S_1, S_2 \subset \mathcal{S}$ are disjoint sets, i.e., $S_1 \cap S_2 = \emptyset$, and $\mathrm{DISJ}_\mathcal{S}(S_1, S_2) = 0$ otherwise. There is a large literature on the communication complexity of the disjointness function (see [20]). We state here one result that will be used to derive bounds in our setting (see [20], Chapter 1.3).

[4]We emphasize that we make no assumptions about the players' incentives. We obtain *lower bounds*, which give the *minimum* it takes to reach an equilibrium—*no matter what the incentives are* (see [6]).

[5]See [17], Section 5(g), particularly the last sentence there.

[6]These are the "generalized ordinal potential games" [21].

[7]We call them "agents" to avoid confusion with the players of the game.

THEOREM 1. *The communication complexity of the $\mathcal{S}$-disjointness function is $|\mathcal{S}|$ bits, i.e.,* $\mathrm{CC}(\mathrm{DISJ}_{\mathcal{S}}) = |\mathcal{S}|$.

2.3 Nash Equilibrium Procedures

A *Nash equilibrium procedure* is a dynamic process by which the players reach a Nash equilibrium of the game, whether pure or mixed (both cases will be considered below). Fix the number of players n and the action spaces A_i; a game G is thus identified with its payoff functions $(u_1, \ldots, u_n)$. Let $\mathcal{G}$ be a family of games to which the procedure should apply. The basic assumption is that of *uncoupledness*: each player knows only his own payoff function u_i [16, 17].

Formally, the n players who participate in a Nash equilibrium procedure have the following information and capabilities. The "input" of the procedure is a game $G = (u_1, \ldots, u_n)$ in the family $\mathcal{G}$. Initially, each player i has access only to his own "private" payoff function[8] u_i. In each round $t = 1, 2, \ldots$, every player i performs an action[9,10] $a_{i,t} \in A_i$. At the end of round t all the players observe each other's actions; i.e., they all observe the joint action $(a_{1,t}, \ldots, a_{n,t}) \in A$.

In a *mixed Nash equilibrium procedure* Π for $\mathcal{G}$, the "output" of each player i is a distribution $p_i \in \Delta_i$, such that $(p_1, \ldots, p_n) \in \Delta$ is a mixed Nash equilibrium of the game $G = (u_1, \ldots, u_n)$ that was given as input.[11] In a *pure Nash equilibrium procedure* Π for $\mathcal{G}$, the "output" of player i is either (1) a pure action $a_i \in A_i$, or (2) a declaration of "no pure Nash equilibrium." In case (1), the joint output $(a_1, \ldots, a_n) \in A$ is a pure Nash equilibrium of G, whereas in case (2) G has no pure Nash equilibrium. Let PNEP and MNEP denote the collection of pure and mixed Nash equilibrium procedures, respectively.

The communication complexity $\mathrm{CC}(\Pi, G)$ of a Nash equilibrium procedure Π applied to a game G is the number of bits communicated until Π terminates when the input is G. Given a family of games $\mathcal{G}$, the communication complexity of a Nash equilibrium procedure Π for the family $\mathcal{G}$ is the worst-case communication complexity of Π over all games $G \in \mathcal{G}$, i.e., $\mathrm{CC}(\Pi, \mathcal{G}) = \max_{G \in \mathcal{G}} \mathrm{CC}(\Pi, G)$. Finally, $\mathrm{CC}(\text{PURE}, \mathcal{G})$, the *communication complexity of pure Nash equilibrium procedures* for a family of games $\mathcal{G}$, is the minimal communication complexity of any pure Nash equilibrium procedure Π for the family of games $\mathcal{G}$, i.e., $\mathrm{CC}(\text{PURE}, \mathcal{G}) = \min_{\Pi \in \text{PNEP}} \mathrm{CC}(\Pi, \mathcal{G})$; similarly, $\mathrm{CC}(\text{MIXED}, \mathcal{G}) = \min_{\Pi \in \text{MNEP}} \mathrm{CC}(\Pi, \mathcal{G})$ is the *communication complexity of mixed Nash equilibrium procedures* for $\mathcal{G}$.

One may measure the communication complexity of Nash procedures also in terms of the number of rounds; this may be more natural from the game-theoretic viewpoint. Formally, the *time communication complexity* $\mathrm{tCC}(\Pi, G)$ of a Nash equilibrium procedure Π applied to a game G is the number of time periods until Π terminates. The two communication complexity measures, CC and tCC, are closely related: in each time period the players transmit at least 1 bit and at most $\sum_i \log |A_i| = \log |A|$ bits.[12]

PROPOSITION 2. *The (bit) communication complexity* CC *and the time communication complexity* tCC *satisfy:*

$$\frac{1}{\log |A|}\,\mathrm{CC} \leq \mathrm{tCC} \leq \mathrm{CC}. \quad (1)$$

(A similar connection for two-player games was observed in [6].)

We are interested in the asymptotic behavior of the communication complexity of Nash equilibrium procedures as the number of players n increases, while the size of the action sets is fixed. Let Γ_s^n be the family of all n-person games where each player has at most s actions, i.e., $|A_i| \leq s$ for all i. We want to estimate the communication complexity of Nash equilibrium procedures on the class Γ_s^n as n increases and s is fixed. Our results will deal with the class Γ_2^n of binary-action games (except for Theorem 4, where we need 4 actions). Since the communication complexity is defined as the worst case over all games, any lower bound for Γ_2^n is clearly also a lower bound for Γ_s^n for every $s \geq 2$. In the complete version of the paper we discuss the extension of our results to $s \geq 2$ actions (we get better lower bounds that depend on s).

3. PURE EQUILIBRIA

In this section we derive exponential lower bounds on the communication complexity of *pure* Nash equilibrium procedures. Our result is

THEOREM 3. *Any pure Nash equilibrium procedure has communication complexity $\Omega(2^n)$, i.e., for every $s \geq 2$,*

$$\mathrm{CC}(\text{PURE}, \Gamma_s^n) \geq \mathrm{CC}(\text{PURE}, \Gamma_2^n) = \Omega(2^n).$$

Proposition 2 implies that the *time communication complexity* of pure Nash equilibrium procedures is $\mathrm{tCC}(\text{PURE}, \Gamma_2^n) = \Omega(2^n/n) = \Omega(2^{n - \log n})$.

At this point one may conjecture that restricting the class of games to those that have pure Nash equilibria may decrease the communication complexity. However, this is not so. Even if one considers only the restricted class $\mathcal{FIP}_s^n$ of n-person s-action games that have the "finite improvement property" (see Section 2.1) and thus always possess pure Nash equilibria, the lower bound remains exponential. Specifically, for games with $s \geq 4$ actions, we have

THEOREM 4. *Any pure Nash equilibrium procedure on the class $\mathcal{FIP}_s^n$ of s-action games with the finite improvement property has communication complexity $\Omega(2^{n/2})$, i.e., for every $s \geq 4$,*

$$\mathrm{CC}(\text{PURE}, \mathcal{FIP}_s^n) \geq \mathrm{CC}(\text{PURE}, \mathcal{FIP}_4^n) = \Omega(2^{n/2}).$$

Theorem 3 will be proved in Section 3.2 below using a simple reduction from the disjointness problem (recall Theorem 1), whereas Theorem 4 will require a much more complex construction, which we present in the Appendix A (the full proof appears in the complete version of the paper).

---

[8]The number of players n, the action spaces A_i, and the set of games $\mathcal{G}$ are fixed and commonly known.

[9]It is natural to consider dynamics in the framework of repeated games, and so the communication proceeds through actions. However, one may well use any set B_i instead of A_i for the period-by-period communication. For binary-action games, $a_{i,t} \in A_i$ just means that the communication of each player in each period is 1 bit.

[10]The procedure is thus *deterministic*; in the complete version of the paper we also discuss stochastic procedures, and show that similar results hold there.

[11]Finite games always possess mixed Nash equilibria.

[12]Throughout this paper log is always $\log_2$.

3.1 Reductions

We now show how to reduce the disjointness problem to the problem of finding pure Nash equilibria. Divide the player set $\{1, \ldots, n\}$ into two sets T_1 and T_2 of size $n/2$ each (assume for simplicity that n is even), say $T_1 = \{1, \ldots, n/2\}$ and $T_2 = \{n/2 + 1, \ldots, n\}$. It will be convenient to rename the players such that the players in T_ℓ are (ℓ, i) for $i \in \{1, \ldots, n/2\}$ and $\ell \in \{1, 2\}$. For any two sets $S_1, S_2 \subset \mathcal{S}$—an input of the $\mathcal{S}$-disjointness problem—the *reduction* will define a game $G = (n, \{A_i\}_i, \{u_i\}_i)$, such that two properties are satisfied:

- *Reducibility*: $S_1 \cap S_2 \neq \emptyset$ if and only if G has a pure Nash equilibrium.

- *Constructibility*: The payoff function of each player (ℓ, i) in T_ℓ is constructible from S_ℓ (i.e., for every $a \in A$ the number $u_{\ell,i}(a)$ is computable, by a finite algorithm, from a, S_ℓ, and i).

The reducibility property enables us to relate the outcome of a pure Nash equilibrium procedure on G with the outcome of the $\mathcal{S}$-disjointness function on S_1 and S_2. Namely, if the players reach a pure Nash equilibrium in G then the sets S_1 and S_2 are not disjoint, and if they do not reach a pure Nash equilibrium then the sets are disjoint. The constructibility property ensures that given a pure Nash equilibrium procedure Π^{NE} we are able to generate a protocol Π_D for the disjointness problem, with the same communication complexity. More specifically, given Π_{NE} we create a protocol Π_D by having agent $\ell \in \{1, 2\}$ simulate all the players in T_ℓ (he can do so by the constructibility property). We summarize this in the following claim, which is based on Theorem 1.

Claim 5. *Assume that there exists a reduction from the $\mathcal{S}$-disjointness problem to n-person pure Nash equilibrium procedures that satisfies the reducibility and constructibility properties. Then any pure Nash equilibrium procedure has communication complexity of at least $|\mathcal{S}|$ bits.*

3.2 Matching Pennies Reduction

We now provide a simple reduction, which we call the *matching pennies reduction*, and establish Theorem 3.

Take $\mathcal{S} = \{0, 1\}^n$; for each $S_1, S_2 \subset \mathcal{S}$ the reduction will generate a binary-action game G in Γ_2^n as follows. The action spaces are $A_i = \{0, 1\}$ for all i, and a joint action is thus $a \in A = \{0, 1\}^n$. The payoff $u_{\ell,i}(a)$ of each player (ℓ, i) in T_ℓ will be high (specifically, 2) if the joint action a lies in the set S_ℓ, and low (specifically, 0) if it does not. In the latter case, two distinguished players in T_ℓ, say $(\ell, 1)$ and $(\ell, 2)$, will in addition play a matching pennies game between themselves.

Formally, for $\ell = 1, 2$, the payoff function $u_{\ell,i}$ of a player (ℓ, i) in T_ℓ is defined as follows. For $i \geq 3$, put[13]

$$u_{\ell,i}(a) = \begin{cases} 2, & \text{if } a \in S_\ell, \\ 0, & \text{if } a \notin S_\ell; \end{cases}$$

as for players $(\ell, 1)$ and $(\ell, 2)$ in T_ℓ, their payoff functions are

$$u_{\ell,1}(a) = \begin{cases} 2, & \text{if } a \in S_\ell, \\ 1, & \text{if } a \notin S_\ell \text{ and } a_{\ell,1} = a_{\ell,2}, \\ 0, & \text{if } a \notin S_\ell \text{ and } a_{\ell,1} \neq a_{\ell,2}; \end{cases}$$

$$u_{\ell,2}(a) = \begin{cases} 2, & \text{if } a \in S_\ell, \\ 0, & \text{if } a \notin S_\ell \text{ and } a_{\ell,1} = a_{\ell,2}, \\ 1, & \text{if } a \notin S_\ell \text{ and } a_{\ell,1} \neq a_{\ell,2}. \end{cases}$$

Claim 6. *For $n \geq 4$, the reducibility and constructibility properties hold for the matching pennies reduction.*

Proof. The payoff functions of the players in T_ℓ depend on S_ℓ only, and so the constructibility property holds. For the reducibility property, note that a is a pure Nash equilibrium if and only if $a \in S_1 \cap S_2$ (indeed, if $a \in S_1 \cap S_2$, then every player gets the maximal payoff of 2; otherwise, $a \notin S_\ell$ for some ℓ, and then either $(\ell, 1)$ or $(\ell, 2)$ benefits by deviating). ∎

We can now prove Theorem 3.

Proof of Theorem 3: Follows from Claims 5 and 6 (recall that $\mathcal{S} = \{0, 1\}^n$). ∎

4. PURE EQUILIBRIA IN A BAYESIAN SETTING

We now consider a Bayesian setting where the game (i.e., the payoff functions) is chosen according to a *probability distribution* that is known to all players. While the communication complexity of pure Nash equilibrium procedures has been shown to be exponential in the worst case, it is conceivable that the *expected* communication complexity will be smaller (where the expectation is taken over the randomized selection of the payoff functions). However, that turns out not to be the case. We will exhibit a simple distribution for which the expected communication complexity of pure Nash equilibrium procedures is exponential. Our result is the following.

Theorem 7. *There exists a probability distribution over games such that any pure Nash equilibrium procedure has expected communication complexity $\Omega(2^n)$; i.e., there exists a probability distribution $\mathbf{P}$ over the family of binary-action games Γ_2^n such that*

$$\mathbf{E}[\mathrm{CC}(\textsc{pure}, G)] = \Omega(2^n),$$

where the expectation $\mathbf{E}$ is taken over games $G \in \Gamma_2^n$ chosen according to the probability distribution $\mathbf{P}$.

(Note that Theorem 3 is implied by Theorem 7.) Unlike the results in the previous section, here we will not apply a reduction, but rather provide a direct proof, using techniques from "distributional communication complexity" (see [20], Chapter 3.4).

Some further background from communication complexity is needed at this point. A *combinatorial rectangle* is $\mathcal{X} = X_1 \times \cdots \times X_n$, where each X_i is a subset of inputs of player i. Every sequence of messages in a communication protocol can be described by a combinatorial rectangle, namely, all inputs generating that sequence of messages. Given a function f of n inputs $x_1, \ldots, x_n$, a combinatorial rectangle $\mathcal{X}$ is called *monochromatic* if $f(x)$ has the same value for all $x = (x_1, \ldots, x_n) \in \mathcal{X}$. A *minimal covering* of

[13] Alternatively: put $u_{\ell,i}(a) = 0$ for all $a \in A$ and all $i \geq 3$.

a function f using combinatorial rectangles is the minimum number of monochromatic combinatorial rectangles needed to represent f (i.e., the minimal number of monochromatic rectangles whose union covers the space of all possible inputs). Clearly, the logarithm of this number is a lower bound on the communication complexity of f (since, roughly speaking, every bit of communication can only split combinatorial rectangles into two; for more details see [20], Chapter 1).

In our setting, the combinatorial rectangles are $\mathcal{U} = U_1 \times \cdots \times U_n$, where each U_i is a set of payoff functions of player i. A monochromatic combinatorial rectangle is labeled by either (1) a pure joint action $a \in A$ (when a is a Nash equilibrium for every game $(u_1, \ldots, u_n) \in \mathcal{U}$), or (2) "no pure Nash equilibrium" (when no game $(u_1, \ldots, u_n) \in \mathcal{U}$ has a pure Nash equilibrium).

Informally, the lower bound on the expected communication complexity of pure Nash equilibrium procedures will be a consequence of the fact that it will be "hard" for the players to agree that there is no pure Nash equilibrium. We will construct a probability distribution over payoff functions such that, first, the probability that there is no pure Nash equilibrium is bounded away from 0 as the number of players n increases. And second, we will show that any combinatorial rectangle that is labeled "no pure Nash equilibrium" has a low probability. This will yield a lower bound on the number of monochromatic combinatorial rectangles, and thus on the communication complexity.

Formally, our probability distribution $\mathbf{P}$ is defined on the family Γ_2^n of binary-action games (i.e., $A_i = \{0,1\}$ for all i). The payoff function u_i of player i is selected randomly as follows. For every $a^{-i} \in \{0,1\}^{n-1}$, with probability $1/2$ put $u_i(0, a^{-i}) = 0$ and $u_i(1, a^{-i}) = 1$, and with probability $1/2$ put $u_i(0, a^{-i}) = 1$ and $u_i(1, a^{-i}) = 0$; these choices are made independently over all a^{-i} and over all i. Note that for every $a \in \{0,1\}^n$ each player i has a unique best reply, and $\mathbf{P}[u_i : a_i \in \mathrm{BR}(a^{-i}; u_i)] = \mathbf{P}[u_i : a_i \notin \mathrm{BR}(a^{-i}; u_i)] = 1/2$.

We start by showing that the probability that there are no pure Nash equilibria is bounded away from 0.

Lemma 8. *There exists a constant $\alpha > 0$ such that the probability that there are no pure Nash equilibria is at least α for all $n \geq 2$.*

Next we show that every combinatorial rectangle labeled "no pure Nash equilibrium" has low probability.

Lemma 9. *Let $\mathcal{U} = U_1 \times \cdots \times U_n$ be a combinatorial rectangle labeled "no pure Nash equilibrium." Then*

$$\mathbf{P}[(u_1, \ldots, u_n) \in \mathcal{U}] \leq 2^{-2^{n-1}}.$$

Combining the two lemmata allows us to prove Theorem 7.

Proof of Theorem 7: By Lemma 8, the total probability of the event that there is no pure Nash equilibrium is bounded from below by $\alpha > 0$. By Lemma 9, each combinatorial rectangle labeled "no pure Nash equilibrium" has probability at most $2^{-2^{n-1}}$. Therefore R, the number of such rectangles, satisfies $R \geq \alpha 2^{2^{n-1}}$; this gives a lower bound on the expected communication complexity of $\log R = \Omega(2^n)$ (see [20], Chapter 2.1, for details). ∎

5. MIXED EQUILIBRIA

Before we introduce our result for mixed Nash equilibrium procedures, a certain preliminary discussion is in order. In the case of mixed Nash equilibria the values of the payoff functions play a crucial role. Consider the following variant of the matching pennies game

$1, 0$	$0, 1$
$0, 1$	$M, 0$

where M is a positive integer. There is a unique Nash equilibrium: $(1/2, 1/2)$ for the row player and $(M/(M+1), 1/(M+1))$ for the column player. Since the parameter M appears only in the payoff function of the row player, and in equilibrium the column player needs to know the precise value of M, it follows that $\log M$ bits have to be communicated. This is a somewhat unsatisfactory result, since the number of bits needed to encode one of the values of the payoff function of the row player is also $\log M$. However, had it been commonly known, for instance, that the payoff functions under consideration have either 1 or M in that entry, then only one bit would have sufficed. We therefore distinguish between two concepts, "magnitude" and "encoding."

Let U_i be a family of payoff functions of player i. The *magnitude* of a rational number ρ is $\mathfrak{mag}(\rho) = \log |M| + \log |K|$, where $\rho = M/K$ is a reduced fraction (i.e., M and K have no common divisor higher than 1), and the magnitude of the family U_i is $\mathfrak{mag}(U_i) = \max_{u_i \in U_i, a \in A} \mathfrak{mag}(u_i(a))$. For each $a \in A$, the *encoding* of the payoff of player i at a is $\mathfrak{enc}(U_i, a) = \log |\{u_i(a) : u_i \in U_i\}|$; i.e., the number of bits required to encode the possible values of $u_i(a)$ as u_i varies over U_i; the encoding of the family U_i is $\mathfrak{enc}(U_i) = \max_{a \in A} \mathfrak{enc}(U_i, a)$. For example, if every payoff function u_i in U_i has two values 1 and M (i.e., $u_i(a) \in \{1, M\}$ for all $u_i \in U_i$ and all $a \in A$), then the encoding of U_i is $\mathfrak{enc}(U_i) = 1$ bit, whereas its magnitude is $\mathfrak{mag}(U_i) = \log M$ bits. Finally, if $\mathcal{U} = U_1 \times \cdots \times U_n$ is a family of games, then $\mathfrak{enc}(\mathcal{U}) = \max_{1 \leq i \leq n} \mathfrak{enc}(U_i)$ and $\mathfrak{mag}(\mathcal{U}) = \max_{1 \leq i \leq n} \mathfrak{mag}(U_i)$.

When deriving lower bounds on the communication complexity of mixed Nash equilibrium procedures, one would like the encoding as well as the magnitude to be as low as possible (so that a high complexity will *not* be just a trivial consequence, as in the example above). Specifically, we will construct a large family of games $\mathcal{U}$ that has an encoding of 1 bit and a magnitude of $O(n)$ bits, such that each game in $\mathcal{U}$ will have a different unique Nash equilibrium. This will imply that, in order to reach the correct Nash equilibrium, the number of bits to be transmitted must be at least the logarithm of the size of the family $\mathcal{U}$. Formally, our result is

Theorem 10. *For every $n \geq 2$ there exists a family of binary-action games $\mathcal{U}^n \subset \Gamma_2^n$ whose encoding is 1 bit and whose magnitude is $O(n)$ bits (i.e., $\mathfrak{enc}(\mathcal{U}^n) = 1$ and $\mathfrak{mag}(\mathcal{U}^n) = O(n)$), such that any mixed Nash equilibrium procedure over $\mathcal{U}^n$ has communication complexity $\Omega(2^n)$, i.e.,*

$$\mathrm{CC}(\textsc{mixed}, \mathcal{U}^n) = \Omega(2^n).$$

Our construction is based on a generalization of Jordan's game [19] in which we modify the payoff of one of the players. For $n \geq 2$, the n-person *Jordan game* J_n is a binary-action game with payoff functions $u_i(a) = \mathbf{1}_{\{a_i = a_{i-1}\}}(a)$ for all players $i \neq 2$ and $u_2(a) = \mathbf{1}_{\{a_2 \neq a_1\}}(a)$ for player 2 (we write $\mathbf{1}_X$ for the indicator function of the event X; e.g., $\mathbf{1}_{\{a_1 = a_n\}}(a) = 1$ if $a_1 = a_n$ and $\mathbf{1}_{\{a_1 = a_n\}}(a) = 0$ otherwise;

and we put $a_0 \equiv a_n$). Thus player 2 wants to "mismatch" the action of player 1, whereas every other player $i \neq 2$ wants to "match" the action of the previous player $i-1$.[14]

Let f be a real function from $\{0,1\}^{n-2}$ to the half-open interval $[0,1)$, i.e., $f : \{0,1\}^{n-2} \to [0,1)$; we define the *f-modified Jordan game* $J_n(f)$ by

$$\begin{aligned} u_i(a) &= \mathbf{1}_{\{a_i=a_{i-1}\}}(a), \quad \text{for } i \neq 2\text{; and} \\ u_2^f(a) &= \mathbf{1}_{\{a_2 \neq a_1\}}(a) + \mathbf{1}_{\{a_1=a_2=1\}}(a) \cdot f(a_3, \ldots, a_n) \end{aligned} \quad (2)$$

(only the payoff of player 2 has been modified).

The following lemma shows that a modified Jordan game has a unique Nash equilibrium, and gives an explicit formula for it. For every function f as above, let

$$\mu(f) = \frac{1}{2^{n-2}} \sum_{(a_3,\ldots,a_n)\in\{0,1\}^{n-2}} f(a_3, \ldots, a_n)$$

be the average of the values of f; equivalently, this is the expected value of f when every player i randomizes uniformly, i.e., $p_i = 1/2$ for all i.

LEMMA 11. *The modified Jordan game $J_n(f)$ has a unique Nash equilibrium $(p_1, \ldots, p_n)$, where $p_i = 1/2$ for all players $i \neq 1$, and*[15]

$$p_1 = \frac{1}{2-\mu(f)} \ . \quad (3)$$

To construct our family of games, we vary the function f over a set $\mathcal{F}$ of functions; thus, for each $i \neq 2$, the family $U_i = \{u_i\}$ is a singleton, whereas the family $U_2 = \{u_2^f : f \in \mathcal{F}\}$ consists of all payoff functions u_2^f of player 2 that are obtained for all $f \in \mathcal{F}$. The property of the family $\mathcal{F}$ will be that, for each function $f \in \mathcal{F}$, when we substitute f in (3) we get a different value for p_1. The lower bound on the communication complexity will follow from the fact that for each $f \in \mathcal{F}$ the communication to player 1 must be different. (Indeed, player 1 needs to reach a different value of p_1 for each f, and always starts with the same information.) This will imply that the number of bits that have to be communicated is at least $\log|\mathcal{F}|$. To formalize this, we will call a set of functions $\mathcal{F}$ *separating* if for any two functions $f_1 \neq f_2$ in $\mathcal{F}$ we have $\mu(f_1) \neq \mu(f_2)$. Thus

CLAIM 12. *Let $\mathcal{U}$ be given as above by a separating set of functions $\mathcal{F}$. Then the communication complexity of any mixed Nash equilibrium procedure on $\mathcal{U}$ is at least* $\log|\mathcal{F}|$.

We now construct our family of functions. For every $x = (x_1, \ldots, x_{n-2})$ in $\{0,1\}^{n-2}$, let $[x]_2 = \sum_{i=1}^{n-2} x_i 2^{n-2-i}$ be the integer corresponding to the binary string x. Let $\mathcal{H}$ be the set of Boolean functions $h : \{0,1\}^{n-2} \to \{0,1\}$. For every $h \in \mathcal{H}$, define a function f_h on $\{0,1\}^{n-2}$ by

$$f_h(x) = h(x)\frac{1}{\mathfrak{prime}([x]_2)}$$

for each $x \in \{0,1\}^{n-2}$, where $\mathfrak{prime}(k)$ is the k-th prime, starting for convenience with $\mathfrak{prime}(0) = 2$ (thus $\mathfrak{prime}(1) = 3$, $\mathfrak{prime}(2) = 5$, and so on; note that indeed $f_h(x) \in [0,1)$). Let $\mathcal{F}_\mathcal{H} = \{f_h : h \in \mathcal{H}\}$. The following lemma shows that $\mathcal{F}_\mathcal{H}$ is a separating family.

LEMMA 13. *The family $\mathcal{F}_\mathcal{H}$ is separating; i.e., for any two Boolean functions $h_1 \neq h_2$ in $\mathcal{H}$ we have $\mu(f_{h_1}) \neq \mu(f_{h_2})$.*

Next, the magnitude of $\mathcal{F}_\mathcal{H}$ is $O(n)$ bits, since $\mathfrak{prime}(k) = O(k \log k)$ by the Prime Number Theorem and so $\log(\mathfrak{prime}([x]_2)) \leq \log(\mathfrak{prime}(2^{n-2})) = O(n)$; whereas the encoding of $\mathcal{F}_\mathcal{H}$ is just 1 bit, since $f_h(x)$ has only two possible values, $1/\mathfrak{prime}([x]_2)$ and 0. The same therefore holds for the resulting family of games $\mathcal{U} \equiv \mathcal{U}^n := \{J_n(f) : f \in \mathcal{F}_\mathcal{H}\}$ (see (2)). We have thus established

CLAIM 14. *The family $\mathcal{U}$ satisfies $\mathfrak{enc}(\mathcal{U}) = 1$ and $\mathfrak{mag}(\mathcal{U}) = O(n)$.*

We can now complete the Proof of Theorem 10.

Proof of Theorem 10: There are $2^{2^{n-2}}$ Boolean functions h in $\mathcal{H}$, so $|\mathcal{F}_\mathcal{H}| = |\mathcal{H}| = 2^{2^{n-2}}$. Combining this with Claims 12 and 14 and Lemma 13 proves Theorem 10. ∎

6. CORRELATED EQUILIBRIA

In this section we study the communication complexity of reaching a correlated equilibrium, and prove that it is polynomial rather than exponential. This shows that the exponential bounds for Nash equilibrium procedures are not due just to the complexity of the input, i.e., to the payoff functions being of exponential size, but rather to the intrinsic complexity of reaching Nash equilibria.

Based on the polynomial-time algorithm of Papadimitriou [22] for computing correlated equilibria of certain "succinct polynomial games," we derive a correlated equilibrium procedure with polynomial communication complexity, for *all games* with integer payoffs. Specifically, let $\mathcal{U}_\mathfrak{u}^n \subset \Gamma_2^n$ be the family of n-person binary-action games with integer payoffs of magnitude at most $\mathfrak{u}$ bits, i.e., $\max_{1\leq i\leq n} \mathfrak{mag}(u_i) \leq \mathfrak{u}$; our correlated equilibrium procedure will have a commmunication complexity that is polynomial in the number of players n and the magnitude of the payoffs $\mathfrak{u}$ (for simplicity we again consider only binary-action games; otherwise, it would be polynomial in $n, \mathfrak{u}$, and $\max_{1\leq i\leq n}|A_i|$).

We start by recalling the definition of a correlated equilibrium; see Aumann [1]. Given a game $G = (n, \{A_i\}_i, \{u_i\}_i)$, a distribution Q over the space of joint actions $A = \prod_{i=1}^n A_i$ is (the distribution of) a *correlated equilibrium* of G if for each player i and all actions $b_i, b'_i \in A_i$, we have $\mathbf{E}_Q[u_i(b_i, a^{-i})\ \mathbf{1}_{\{a_i=b_i\}}] \geq \mathbf{E}_Q[u_i(b'_i, a^{-i})\,\mathbf{1}_{\{a_i=b_i\}}]$ (where $\mathbf{E}_Q$ denotes expectation with respect to the distribution Q). Equivalently, consider the "extended game" where, before G is played, a joint action $a = (a_1, ..., a_n) \in A$ is randomly chosen according to Q and each player i is given a "recommendation" to play a_i, his coordinate of the chosen a; then Q is a correlated equilibrium of G if and only if the combination of strategies where each player always plays according to his recommendation constitutes a Nash equilibrium of the extended game.

A *correlated equilibrium procedure* Π is defined in the same way as a Nash equilibrium procedure, except that now the output of each player is a distribution Q, such that Q is a correlated equilibrium of the game $G = (u_1, \ldots, u_n)$ that was given as input.[16] Let CEP be the collection of correlated equilibrium procedures. Similarly to CC(MIXED, $\mathcal{G}$)

[14]This game has a unique Nash equilibrium $(1/2, ..., 1/2)$ (this also follows from Lemma 11 below).

[15]Recall that p_i stands for the probability of action 1, i.e., $p_i = \mathbf{P}[a_i = 1]$.

[16]Finite games always possess correlated equilibria.

and CC(PURE, $\mathcal{G}$), we define the *communication complexity of correlated equilibrium procedures* for a family of games $\mathcal{G}$ as $CC(\text{CORRELATED}, \mathcal{G}) = \min_{\Pi \in \text{CEP}} CC(\Pi, \mathcal{G}) = \min_{\Pi \in \text{CEP}} \max_{G \in \mathcal{G}} CC(\Pi, G)$.

We come now to the construction of [22], which consists of running an ellipsoid algorithm in the Hart-Schmeidler setup [18]. In our communication complexity framework, every player can run internally the computations of the algorithm at no cost. However, since the payoff function u_i is known only to player i, only i can compute his own expected payoffs—which he can then broadcast to all players. The communication complexity counts only the number of bits transmitted, and therefore, as we will see, there is no need to restrict ourselves to "succinct games of polynomial type" as in [22].

We define the procedure Π_{CORR} as follows. All players simulate the algorithm of [22]. At each step of the ellipsoid algorithm, an n-tuple of mixed strategies $p = (p_1, ..., p_n) \in \Delta = \prod_{i=1}^{n} \Delta_i$ is generated (the whole vector p is computed internally by—and thus known to—each player). Every player i then computes his expected payoff $u_i(p)$ and broadcasts it. In terms of communication complexity, again, the local computation of p and $u_i(p)$ has no cost; only the transmission of $u_i(p)$ counts.

Papadimitriou [22] proves, first, that a correlated equilibrium is reached in a number of steps that is bounded by a polynomial in n and $\mathfrak{u}$; and second, that the n-tuples of mixed strategies $p \in \Delta$ generated at every step have a magnitude $\mathfrak{mag}(p) = O(n\mathfrak{u})$ bits. Therefore, when the payoffs $u_i(a)$ for all $a \in A$ are integers of at most $\mathfrak{u}$ bits, the expected payoff $u_i(p)$ for $p \in \Delta$ requires at most $O(n\,\mathfrak{mag}(p) + \mathfrak{u} + n) = O(n^2\mathfrak{u})$ bits (since it is a weighted sum of 2^n entries). Altogether, this implies that the total number of bits transmitted in the procedure Π_{CORR} is bounded by a polynomial in n and $\mathfrak{u}$, and we have shown

THEOREM 15. *For every $n \geq 2$ and $\mathfrak{u} \geq 1$, let $\mathcal{U}_{\mathfrak{u}}^n \subset \Gamma_2^n$ be the family of n-person binary-action games with integer payoffs of magnitude at most $\mathfrak{u}$, i.e., $\max_{1 \leq i \leq n} \mathfrak{mag}(u_i) \leq \mathfrak{u}$. Then, there exists a correlated equilibrium procedure Π_{CORR} whose communication complexity over $\mathcal{U}_{\mathfrak{u}}^n$ is polynomial in n and $\mathfrak{u}$, i.e.,*

$$CC(\text{CORRELATED}, \mathcal{U}_{\mathfrak{u}}^n) \leq CC(\Pi_{\text{CORR}}, \mathcal{U}_{\mathfrak{u}}^n) \leq poly(n, \mathfrak{u}).$$

In the full version of the paper we present further results on the communication complexity of reaching correlated equilibria. Specifically, in the classes of games of Sections 4 and 5 where the communication complexity of reaching Nash equilibria was shown to be exponential, that of correlated equilibria turns out to be quite low. We also analyze procedures for reaching correlated *approximate* equilibria.

7. NASH APPROXIMATE EQUILIBRIA

An *approximate* equilibrium requires each player's gain from deviating to be small. Formally, given $\varepsilon > 0$, a *Nash ε-equilibrium* is a combination of mixed actions $p = (p_1, ..., p_n) \in \Pi_{i=1}^{n} \Delta_i = \Delta$ such that $u_i(p) \geq u_i(q_i, p^{-i}) - \varepsilon$ for every player i and any mixed action $q_i \in \Delta_i$ of i.

It is clearly of interest to study the communication complexity of reaching Nash *approximate* equilibria, and determine whether or not it is also exponential in the number of players.

However, the techniques we have developed in this work do not seem to be able to deal with the communication complexity of approximate Nash equilibrium procedures. Indeed, our analysis is based on games that have no pure Nash equilibria (Sections 3 and 4), or whose mixed Nash equilibria require large descriptions (Section 5) —whereas there always exist approximate Nash equilibria, and moreover with succinct representations.

8. REFERENCES

[1] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

[2] A. Blum and Y. Mansour. From external to internal regret. In *COLT*, pages 621–636, 2005.

[3] A. Cahn. General procedures leading to correlated equilibria. *International Journal of Game Theory*, 33:21–40, 2004.

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

[5] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning and Games*. Cambridge University Press, 2006.

[6] V. Conitzer and T. Sandholm. Communication complexity as a lower bound for learning in games. In *ICML*, 2004.

[7] D. Foster and S. M. Kakade. Deterministic calibration and Nash equilibrium. In *COLT*, pages 33–48, 2004.

[8] D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21:40–55, 1997.

[9] D. Foster and H. P. Young. Learning, hypothesis testing, and Nash equilibrium. *Games and Economic Behaviour*, 45(1):73–96, Oct 2003.

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

[11] Drew Fudenberg and David K. Levine. *The theory of learning in games*. MIT press, 1998.

[12] F. Germano and G. Lugosi. Global Nash convergence of Foster and Young's regret testing. *Games and Economic Behavior*, 2006. To appear.

[13] S. Hart. Adaptive heuristics. *Econometrica*, 73:1401–1430, 2005.

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

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

[16] S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.

[17] S. Hart and A. Mas-Colell. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57:286–303, 2006.

[18] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14:18–25, 1989.

[19] J. Jordan. Three problems in learning mixed equilibria. *Games and Economic Behaviour*,

5:368–386, 1993.

[20] E. Kushilevitz and N. Nisan. *Commincation Complexity*. Cambridge University Press, 1997.

[21] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behaviour*, 14:124–143, 1996.

[22] C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In *Proceedings of the thirty-seventh annual ACM Symposium on Theory of Computing*, 2005.

[23] G. Stoltz and G. Lugosi. Internal regret in on-line portfolio selection. *Machine Learning*, 59(1-2):125–159, 2005.

[24] G. Stoltz and G. Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 2006. To appear.

[25] H. P. Young. *Strategic Learning and Its Limits*. Oxford University Press, 2004.

APPENDIX

A. POTENTIAL GAME REDUCTION

In Section 3.2 we provided a reduction—the matching pennies reduction—from procedures for the 2^n-disjointness problem to n-person pure Nash equilibrium procedures. We now construct another reduction, which we call the *potential game reduction*, whose additional property[17] is that whenever the two sets in the disjointness problem intersect, the corresponding game has the finite improvement path property. This reduction will establish Theorem 4.

Before we describe the potential game reduction, it is worthwhile to investigate why an alternative naive reduction fails. Let us start with a few notations, which will be useful later on. The *Hamming distance* $d_H(w,v)$ between two vectors $w,v \in \{0,1\}^k$ equals the number of coordinates in which they differ; for a set $V \subset \{0,1\}^k$, put $d_H(w,V) = \min_{v\in V} d_H(w,v)$.

Recall from Section 3.1 that n is assumed even and the set of n players is partitioned into two sets of $n/2$ players, T_1 and T_2; the players of T_ℓ are denoted (ℓ,i) for $i \in \{1,\ldots,n/2\}$ and $\ell \in \{1,2\}$. Take $\mathcal{S} = \{0,1\}^{n/2}$, and consider the following reduction from the $\mathcal{S}$-disjointness problem to binary-action games Γ_2^n. Let $S_1, S_2 \subset \mathcal{S} = \{0,1\}^{n/2}$. For each joint action $a \in \{0,1\}^n$ define $z = z(a) \in \{0,1\}^{n/2}$ by $x_i = a_{1,i} \oplus a_{2,i}$ for all $i \in \{1,\ldots,n/2\}$, and let the payoff functions be $u_{\ell,i}(a) = -d_H(z(a), S_\ell)$, for all $a \in \{0,1\}^n$, $i \in \{1,\ldots,n/2\}$, and $\ell \in \{1,2\}$. One can show that there exists a pure Nash equilibrium in this game iff $S_1 \cap S_2 \neq \emptyset$. However, improvement paths in these games are not necessarily finite.[18] Our potential game reduction will also use the Hamming distance to drive the joint action to a certain region, but will require a much more complex structure in order to guarantee that all improvement paths are finite.

We now present the potential game reduction. Let H be a Hamiltonian cycle in the hypercube $\{0,1\}^{n/2}$. For $x, y \in \{0,1\}^{n/2}$, when y immediately follows x in the cycle H we write $y = next(x)$ and $x = prev(y)$; also, let $r(x) \in \{1,\ldots,n/2\}$ denote the index of the unique bit in which x and $next(x)$ differ. Let

$$\begin{aligned} L &= \left\{xx : x \in \{0,1\}^{n/2}\right\} \\ &\cup \left\{yx : y, x \in \{0,1\}^{n/2},\ y = next(x)\right\}, \end{aligned}$$

where zw denotes the concatenation of the strings z and w. Clearly $L \subset \{0,1\}^n$ and $|L| = 2 \cdot 2^{n/2}$. In our games every joint action $a \in A$ will be mapped to some $z(a) \in \{0,1\}^n$, and the payoff of every player will increase as $z(a)$ approaches the set L. A pure Nash equilibrium a, if it exists, will always have $z(a) \in L$.

The players in T_1 have binary actions, i.e., $A_{1,i} = \{0,1\}$, whereas those in T_2 have four actions: $A_{2,i} = \{0,1\} \times \{0,1\}$. For an action $a_{2,i} = (c_{2,i}, d_{2,i}) \in A_{2,i}$, we will refer to $c_{2,i}$ as the *action bit* and to $d_{2,i}$ as the *done bit*. A joint action $a \in A$ can be written $a = (a_1, a_2)$, where $a_1 \in \prod_i A_{1,i}$ and $a_2 \in \prod_i A_{2,i}$ are the joint actions of T_1 and T_2, respectively. Given $a = (a_1, a_2)$, define

$$\begin{aligned} x_1 &\equiv x_1(a_1) = a_1 = (a_{1,1}, \ldots, a_{1,n/2}) \in \{0,1\}^{n/2}, \\ x_2 &\equiv x_2(a_2) = (c_{2,1}, \ldots, c_{2,n/2}) \in \{0,1\}^{n/2}, \\ d_2 &\equiv d_2(a_2) = (d_{2,1}, \ldots, d_{2,n/2}) \in \{0,1\}^{n/2}, \text{ and} \\ z &\equiv z(a) = x_1x_2 \in \{0,1\}^n \end{aligned}$$

(x_1x_2 is the concatenation of x_1 and x_2).

We will view L as a cycle that moves from each xx to yx, where $y = next(x)$, and then from yx to yy. As the joint action a changes, so does the resulting $z(a)$. To move $z(a)$ in L between xx and yx one player in T_1, namely $(1, r(x))$, must change his action; we call him the *forward active* 1-*player* at xx, and also the *backward active* 1-*player* at yx, and denote him by $r_1(xx) = r_1(yx) = (1, r(x))$. Similarly, the move between yx and yy is controlled by the action bit of one player in T_2, namely $r_2(yx) = r_2(yy) = (2, r(x))$, which we call the *forward active* 2-*player* at yx, and also the *backward active* 2-*player* at yy.

A high-level description of our reduction is as follows. Given two subsets S_1 and S_2 of $\mathcal{S} = \{0,1\}^{n/2}$, we define the payoff functions of the players such that: (1) all players want to reach L (i.e., have $z(a) \in L$) and stay in it; (2) when in L, only the active players have an incentive to change their actions; (3) if the joint action is xx and $x \in S_1 \cap S_2$ then no active player has an incentive to change his action, and we are at a pure Nash equilibrium; and (4) the payoff functions of the players in T_ℓ depend only on S_ℓ, for $\ell \in \{1,2\}$.

Specifically, for each player $(1,i)$ in T_1 we define his payoff function

$$u_{1,i}(a) = \begin{cases} -d, & \text{if } x_1x_2 \notin L, \\ 1, & \text{if } x_1x_2 \in L \text{ and } x_1 \neq x_2, \\ 2, & \text{if } x_1x_2 \in L, x_1 = x_2, x_1 \in S_1, \\ & \quad \text{and } d_{r_2(x_1x_2)} = 1, \\ 0, & \text{if } x_1x_2 \in L, x_1 = x_2, \\ & \quad x_1 \notin S_1 \text{ or } d_{r_2(x_1x_2)} = 0, \end{cases}$$

where $d = d_H(x_1x_2, L)$.

Thus, if $z = x_1x_2 \notin L$ then $u_{1,i}(a)$ is the negative of the Hamming distance from z to the set L (this provides the incentive always to move in the direction of L, and once L is reached not to leave it). If $x_1 = x_2$, $x_1 \in S_1$, and $d_{r_2(x_1x_2)} = 1$, then $u_{1,i}(a)$ has the maximal value of 2 (this is where the pure Nash equilibria will be, if at all); note that

[17] Besides reducibility and constructibility; see Section 3.1.

[18] For example, take S_1, S_2 such that $S_1 \cap S_2 \neq \emptyset$ and there are two vectors w and v with $d_H(w,S_1) < d_H(w,S_2)$, $d_H(v,S_2) < d_H(v,S_1)$, and $d_H(w,v) = 1$. Let i be the index where w and v differ; players $(1,i)$ and $(2,i)$ can then alternate indefinitely in performing improvement steps.

players in T_1 can test $d_{r_2(x_1x_2)} = 1$ since the identity of the active 2-player $r_2(x_1x_2) = r_2(z(a))$ is just a function of the joint action a. If $x_1 \neq x_2$ then $u_{1,i}(a) = 1$, and otherwise $u_{1,i}(a) = 0$ (this will cause the players in T_1 to prefer to move from $x_1 = x_2$ to $x_1 \neq x_2$, unless both $x_1 \in S_1$ and $d_{r_2(x_1x_2)} = 1$).

For each player $(2, i)$ in T_2, we first define an auxiliary function $GoodDone_{2,i}$

$$GoodDone_{2,i}(a) = \begin{cases} 0, & \text{if } x_1 = x_2, (2,i) = r_2(x_1x_2), \\ & \quad \text{and } \ d_{2,i} \neq \mathbf{1}_{\{x_2 \in S_2\}}, \\ 1, & \text{otherwise,} \end{cases}$$

and then the payoff function

$$u_{2,i}(a) = \begin{cases} -d_H(x_1x_2, L), & \text{if } x_1x_2 \notin L, \\ 0, & \text{if } x_1x_2 \in L \\ & \quad \text{and } x_1 \neq x_2, \\ 2 \cdot GoodDone_{2,i}(a), & \text{if } x_1x_2 \in L, x_1 = x_2, \\ & \quad \text{and } x_2 \in S_2, \\ GoodDone_{2,i}(a), & \text{if } x_1x_2 \in L, x_1 = x_2, \\ & \quad \text{and } x_2 \notin S_2. \end{cases}$$

The idea is that when $x_1 = x_2$ the active 2-player $(2, i) = r_2(x_1x_2)$ should "signal" through his done bit whether or not $x_2 \in S_2$ (this is needed to let the players in T_1 know when a Nash equilibrium has been reached); if he does not signal correctly he is "penalized" by having $GoodDone_{2,i} = 0$ instead of 1, which decreases his payoff.

Claim 16. *The constructibility and reducibility properties hold for the potential game reduction.*

Proof. By definition of the reduction, the payoffs of the players in T_ℓ depend only on S_ℓ, and so the constructibility property holds. It remains to show that the reducibility property holds.

We will distinguish five types of joint actions a in A and analyze each in turn.

(1) a such that $z(a) = x_1x_2 \in L, x_1 = x_2, x_1 \in S_1$, $d_{r_2(x_1x_2)} = 1$ and $x_2 \in S_2$—thus $x_1 = x_2 \in S_1 \cap S_2$—is a pure Nash equilibrium, since all players get their maximal payoff of 2 (we have $GoodDone_{2,i}(a) = 1$ for all players $(2, i)$ in T_2). Such an a is obtained from $x = x_1 = x_2 \in S_1 \cap S_2$ by putting $a_{1,i} = x_{(i)}$ (= the i's coordinate of x) for each player $(1, i)$ in T_1, and $a_{2,i} = (c_{2,i}, d_{2,i})$ with action bit $c_{2,i} = x_{(i)}$ and arbitrary done bit $d_{2,i}$ for each player $(2, i)$ in T_2, except for the active 2-player $r_2(x_1x_2)$, whose done bit is $d_{r_2(x_1x_2)} = 1$.

(2) a such that $z(a) = x_1x_2 \notin L$ cannot be a Nash equilibrium since at least one player (ℓ, i), by changing his action, can bring the new $z(a')$ closer to L and thus increase his payoff by 1.

(3) a such that $z(a) = x_1x_2 \in L, x_1 \neq x_2$ cannot be a Nash equilibrium, since the (forward) active 2-player, by changing his action bit and also setting his done bit correctly (to $d_{r_2(x_1x_2)} = \mathbf{1}_{\{x_2 \in S_2\}}$), can increase his payoff from 0 to either 1 or 2.

(4) a such that $z(a) = x_1x_2 \in L, x_1 = x_2$, and either $x_1 \notin S_1$ or $d_{r_2(x_1x_2)} = 0$ cannot be a Nash equilibrium since the active 1-player can increase his payoff from 0 to 1 by changing his action.

(5) a such that $z(a) = x_1x_2 \in L, x_1 = x_2, x_1 \in S_1$, $d_{r_2(x_1x_2)} = 1$ and $x_2 \notin S_2$ cannot be a Nash equilibrium since $GoodDone_{r_2(x_1x_2)}(a) = 0$ and so the active 2-player $r_2(x_1\, x_2)$ can increase his payoff from 0 to 1 by changing his done bit to $d_{r_2(x_1x_2)} = 0$.

Now (1)–(5) cover all possibilities, and we have shown that if $S_1 \cap S_2 \neq \emptyset$ then there is a pure Nash equilibrium (case (1)), whereas if $S_1 \cap S_2 = \emptyset$ then there is no pure Nash equilibrium. ∎

In the complete version of the paper we show that when $S_1 \cap S_2 \neq \emptyset$ all the improvement paths are finite, which establishes Theorem 4.