

How long to equilibrium? The communication complexity of uncoupled equilibrium procedures [☆]

Sergiu Hart ^{a,*}, Yishay Mansour ^b

^a Center for the Study of Rationality, Institute of Mathematics, and Department of Economics,
The Hebrew University of Jerusalem, Israel

^b School of Computer Science, Tel Aviv University, Israel

Received 11 December 2006

Available online 12 February 2008

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.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C7; D83

Keywords: Uncoupled dynamics; Nash equilibrium; Communication complexity; Correlated equilibrium; Speed of convergence

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., the books Fudenberg and Levine, 1998; Young, 2004). 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

[☆] Previous versions: April 2006, December 2006; *STOC 2007* (Extended Abstract: “The Communication Complexity of Uncoupled Nash Equilibrium Procedures”).

* Corresponding author at: The Hebrew University of Jerusalem, Center for the Study of Rationality, Feldman Building, Givat Ram, 91904 Jerusalem, Israel.

E-mail addresses: hart@huji.ac.il (S. Hart), mansour@tau.ac.il (Y. Mansour).

URLs: <http://www.ma.huji.ac.il/hart> (S. Hart), <http://www.cs.tau.ac.il/~mansour> (Y. Mansour).

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* (Aumann, 1974). 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 (Blum and Mansour, 2007; Cahn, 2004; Cesa-Bianchi and Lugosi, 2003, 2006; Foster and Vohra, 1997; Hart, 2005; Hart and Mas-Colell, 2000, 2001; Stoltz and Lugosi, 2005, 2007; Young, 2004) 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* (Hart and Mas-Colell, 2003): each player is assumed to know initially *only his own payoff function*, and not those of the other players. Parallel notions are “informationally decentralized” or “privacy-preserving” in economics, and “distributed” in computer science; in particular, the longstanding study of mechanism design in economics, starting with the work of Hurwicz (Hurwicz, 1960), has led to significant concepts and insights concerning the aggregation of information (see Jordan, in press, for a short survey).

The existing results on uncoupled dynamics and Nash equilibria are as follows. On the one hand, it has been shown that it is impossible for uncoupled dynamics that are deterministic and continuous² always to converge to a Nash equilibrium, even when it is unique (Hart and Mas-Colell, 2003). On the other hand, there are a number of uncoupled dynamics that converge to Nash equilibria in general games; these dynamics use various techniques such as hypothesis-testing, regret-testing, and other variants of exhaustive or stochastic search (Foster and Kakade, 2004; Foster and Young, 2003, 2006; Germano and Lugosi, 2007; Hart and Mas-Colell, 2006; Young, 2004).³ Since 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 this paper we will show that this is a general phenomenon rather than a deficiency of the existing literature: *the lower bound on the speed of convergence to Nash equilibria is exponential in the number of players*.

To make this precise, define a *Nash equilibrium procedure* as a dynamic process whereby the players reach a Nash equilibrium, whether *pure* or *mixed*. 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* (see Kushilevitz and Nisan, 1997) 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—needed by the players in order to reach a Nash equilibrium. This important connection was first observed in Conitzer and Sandholm (2004), 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).

We also exhibit simple procedures that yield upper bounds that are also exponential (Section 6).

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 Continuous with respect to both actions and time.

³ Dynamics of the “best-reply” variety have been studied in Bayesian setups where players possess certain probabilistic beliefs about the payoff functions of the other players (see, e.g., Jordan, 1991; Kalai and Lehrer, 1993); however, additional coordination between the players is needed to obtain convergence to Nash equilibria (cf. Section 4 in Jordan, 1991 and footnote 20 in Hart, 2005).

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 7 and Appendix B). 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.⁴

2. Preliminaries

2.1. Game-theoretic setting

The basic setting is as follows. There are $n \geq 2$ players, $i = 1, 2, \dots, n$. Each player i has a finite set of actions A_i with⁵ $|A_i| \geq 2$, 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., $u_i : A \rightarrow \mathbb{R}$. We extend u_i to Δ in a multilinear way, by defining $u_i(p_1, \dots, p_n) = \mathbf{E}[u_i(a_1, \dots, a_n)]$ for each $(p_1, \dots, p_n) \in \Delta$, where the expectation \mathbf{E} is taken with respect to the product distribution $p_1 \times \dots \times p_n$ on A . We denote this game by $G = (n, \{A_i\}_i, \{u_i\}_i)$.

For a joint action $a = (a_1, \dots, a_i, \dots, a_n) \in A$, let $a^{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ be the joint action of all players *except* player i . For each player i , the (pure) *best-reply correspondence* maps a joint action a^{-i} of the other players to the set $\text{BR}(a^{-i}; u_i) = \arg \max_{a_i \in A_i} u_i(a_i, a^{-i})$. A joint action $a \in A$ is a *pure Nash equilibrium* if $u_i(a) \geq u_i(b_i, a^{-i})$ for every player i and any action $b_i \in A_i$; or equivalently, $a_i \in \text{BR}(a^{-i}; u_i)$ 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 finite⁶; 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 $\text{CC}(\Pi, f, x, y)$. The communication complexity $\text{CC}(\Pi, f)$ 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., $\text{CC}(\Pi, f) = \max_{x, y \in \{0, 1\}^K} \text{CC}(\Pi, f, x, y)$. Finally, the *communication complexity* $\text{CC}(f)$ of computing a function f is the minimum over all protocols Π for computing f , i.e., $\text{CC}(f) = \min_{\Pi} \text{CC}(\Pi, f)$.

A well-studied function in communication complexity is the *disjointness* function. Let \mathcal{S} be a finite set; the \mathcal{S} -*disjointness function* $\text{DISJ}_{\mathcal{S}}$ is defined on the subsets of \mathcal{S} (i.e., on $\{0, 1\}^{\mathcal{S}} \times \{0, 1\}^{\mathcal{S}}$) by $\text{DISJ}_{\mathcal{S}}(S_1, S_2) = 1$ if the two inputs $S_1, S_2 \subset \mathcal{S}$ are disjoint sets (i.e., $S_1 \cap S_2 = \emptyset$), and $\text{DISJ}_{\mathcal{S}}(S_1, S_2) = 0$ otherwise. There is a large literature

⁴ See Hart and Mas-Colell (2006, Section 5(g)), particularly the last sentence there.

⁵ The number of elements of a finite set Z is denoted $|Z|$.

⁶ These are the “generalized ordinal potential games” (Monderer and Shapley, 1996).

⁷ We call them “agents” to avoid confusion with the players of the game.

on the communication complexity of the disjointness function (see Kushilevitz and Nisan, 1997). We state here one result that will be used to derive bounds in our setting (see Kushilevitz and Nisan, 1997, Section 1.3).

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

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, \dots, 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 (Hart and Mas-Colell, 2003, 2006).

We emphasize that we make no assumptions about the players' incentives, since we are interested in *lower bounds*, which give the *minimum* it takes to reach an equilibrium—*no matter what the incentives are*. Indeed, any form of strategic behavior when choosing the messages would be an additional restriction that can only increase the communication complexity (cf. Conitzer and Sandholm, 2004).

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 $G = (u_1, \dots, u_n)$ in the family \mathcal{G} . Initially, each player i has access only to his own “private” payoff function⁸ u_i . In each round $t = 1, 2, \dots$, 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 $(a_{1,t}, \dots, a_{n,t}) \in A$.

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 $(p_1, \dots, p_n) \in \Delta$ is a mixed Nash equilibrium of the game $G = (u_1, \dots, u_n)$ 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 $(a_1, \dots, a_n) \in A$ 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 $\text{CC}(\Pi, G)$ 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., $\text{CC}(\Pi, \mathcal{G}) = \max_{G \in \mathcal{G}} \text{CC}(\Pi, G)$. Next, $\text{CC}(\text{PURE}, \mathcal{G})$, 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., $\text{CC}(\text{PURE}, \mathcal{G}) = \min_{\Pi \in \text{PNEP}} \text{CC}(\Pi, \mathcal{G})$; similarly, $\text{CC}(\text{MIXED}, \mathcal{G}) = \min_{\Pi \in \text{MNEP}} \text{CC}(\Pi, \mathcal{G})$ is the *communication complexity of mixed Nash equilibrium procedures* for \mathcal{G} . Finally, when the games in the family \mathcal{G} are chosen according to a probability distribution \mathbf{P} , the *expected communication complexity of pure Nash equilibrium procedures* is $\min_{\Pi \in \text{PNEP}} \mathbf{E}[\text{CC}(\Pi, G)]$, where the expectation \mathbf{E} is taken with respect to \mathbf{P} ; we denote this by $\mathbf{E}[\text{CC}(\text{PURE}, \mathcal{G})]$.

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* $\text{tCC}(\Pi, G)$ 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

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 we assume without loss of generality that the communication proceeds through actions. Using any other set B_i instead of A_i will only affect the communication complexity by a constant factor (cf. Proposition 2). 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*; see Section 8.1 for *stochastic* procedures.

¹¹ Finite games always possess mixed Nash equilibria.

¹² Throughout this paper \log is always \log_2 .

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

$$\frac{1}{\log |A|} \text{CC} \leq \text{tCC} \leq \text{CC}. \tag{1}$$

(A similar connection for two-player games was observed in Conitzer and Sandholm, 2004.)

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 $s \geq 2$; see also Section 8.2.

3. Pure equilibria

In this section we derive exponential lower bounds on the communication complexity of *pure* Nash equilibrium procedures. Our result is¹³

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

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

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

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 specific 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 $s \geq 4$,*

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

Theorem 3 will be proved in Section 3.2 using a simple reduction from the disjointness problem (recall Theorem 1), whereas Theorem 4 will require a much more complex construction, which is relegated to Appendix A.

3.1. Reductions

We now show how to reduce the disjointness problem to the problem of finding pure Nash equilibria. Divide the player set $\{1, \dots, n\}$ into two sets T_1 and T_2 of size $n/2$ each (assume for simplicity that n is even), say $T_1 = \{1, \dots, n/2\}$ and $T_2 = \{n/2 + 1, \dots, n\}$. It will be convenient to rename the players such that the players in T_ℓ are (ℓ, i) for $i \in \{1, \dots, n/2\}$ 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).

¹³ Notation: $f(n) = \Omega(g(n))$ and $f(n) = O(g(n))$ mean that there exists a constant $C > 0$ such that $f(n) \geq Cg(n)$, respectively $f(n) \leq Cg(n)$, for all n .

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 $i \geq 3$, put¹⁴

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

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). \square

We can now prove Theorem 3.

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

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

¹⁴ Alternatively: put $u_{\ell,i}(a) = 0$ for all $a \in A$ and all $i \geq 3$.

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*

$$\mathbf{E}[\text{CC}(\text{PURE}, \Gamma_2^n)] = \Omega(2^n).$$

(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 Kushilevitz and Nisan, 1997, Sections 1.2 and 3.4).

Some further background from communication complexity is needed at this point. A *combinatorial rectangle* is $\mathcal{X} = X_1 \times \dots \times X_n$, 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, \dots, x_n , a combinatorial rectangle \mathcal{X} is called *monochromatic* if $f(x)$ has the same value for all $x = (x_1, \dots, x_n) \in \mathcal{X}$. A *minimal covering* of a function f using combinatorial rectangles is the minimum number of monochromatic combinatorial rectangles needed to represent f (i.e., the minimum 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 Kushilevitz and Nisan, 1997, Chapter 1).

In our setting, the combinatorial rectangles are $\mathcal{U} = U_1 \times \dots \times U_n$, 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 $(u_1, \dots, u_n) \in \mathcal{U}$), or (2) “no pure Nash equilibrium” (when no game $(u_1, \dots, u_n) \in \mathcal{U}$ 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 $\mathbf{P}[u_i: a_i \in \text{BR}(a^{-i}; u_i)] = \mathbf{P}[u_i: a_i \notin \text{BR}(a^{-i}; u_i)] = 1/2$.

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*

$$\mathbf{P}[(u_1, \dots, u_n): \text{the game } (u_1, \dots, u_n) \text{ has no pure Nash equilibrium}] \geq \alpha$$

for all $n \geq 2$.

Proof. First, we claim that for every $a \in \{0, 1\}^n$, the probability that a is a pure Nash equilibrium is 2^{-n} . Indeed, a is a pure Nash equilibrium if and only if for every player i the payoff function u_i satisfies $a_i \in \text{BR}(a^{-i}; u_i)$. This holds with probability 1/2, independently, for each player i , and so the probability that a is a pure Nash equilibrium is 2^{-n} .

Second, let N be the number of pure Nash equilibria; then $\mathbf{E}[N] = 1$, since there are 2^n joint actions $a \in \{0, 1\}^n$ and the probability of each one being a pure Nash equilibrium is 2^{-n} . Therefore $\mathbf{P}[N = 0] \geq \mathbf{P}[N \geq 2]$, since $1 = \mathbf{E}[N] = \sum_{k \geq 1} k\mathbf{P}[N = k] \geq 1\mathbf{P}[N = 1] + 2\mathbf{P}[N \geq 2] = (1 - \mathbf{P}[N = 0]) + \mathbf{P}[N \geq 2]$.

Third, let $Z = \{a \in \{0, 1\}^n: |\{i: a_i = 1\}| \text{ is even}\}$ be the set of joint actions with an even number of ones, and let ζ be the probability that in Z there are exactly 2 pure Nash equilibria. For any $a, b \in Z$ with $a \neq b$, the event that a is a pure Nash equilibrium is independent of the event that b is a pure Nash equilibrium (a and b differ in at least two

coordinates by the definition of Z ; hence for every player i we have $a^{-i} \neq b^{-i}$). Therefore the probability that exactly two elements of Z (whose size is $|Z| = 2^{n-1}$) are Nash equilibria is

$$\zeta = \binom{2^{n-1}}{2} (2^{-n})^2 (1 - 2^{-n})^{2^{n-1}-2},$$

which is positive for all $n \geq 2$. Moreover, $\zeta \rightarrow (1/8)e^{-1/2} > 0$ as $n \rightarrow \infty$, and so $\mathbf{P}[N = 0] \geq \mathbf{P}[N \geq 2] \geq \zeta$ is indeed bounded away from 0. \square

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

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

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

Proof. First, we claim that for every $a \in \{0, 1\}^n$ there exists a player i such that $a_i \notin \text{BR}(a^{-i}; u_i)$ for all $u_i \in U_i$. Indeed, otherwise, for every player i we would have $a_i \in \text{BR}(a^{-i}; u'_i)$ for some $u'_i \in U_i$. But this would imply that a is a pure Nash equilibrium for (u'_1, \dots, u'_n) , which belongs to \mathcal{U} since \mathcal{U} is a rectangle—in contradiction to the assumption that every game in \mathcal{U} has no pure Nash equilibrium.

Second, as in the proof of Lemma 8, let Z be the set of all $a \in \{0, 1\}^n$ with an even number of ones. Define $Z_i = \{a \in Z : a_i \notin \text{BR}(a^{-i}; u_i) \text{ for all } u_i \in U_i\}$; then $U_i \subset \{u_i : a_i \notin \text{BR}(a^{-i}; u_i) \text{ for all } a \in Z_i\}$, and, as we saw just above, $Z = \bigcup_{i=1}^n Z_i$. Since the events $\{u_i : a_i \notin \text{BR}(a^{-i}; u_i)\}$ are independent for different $a \in Z_i$ (again, $a \neq b$ implies that $a^{-i} \neq b^{-i}$ for $a, b \in Z$) and each one has probability $1/2$, we get

$$\mathbf{P}[u_i \in U_i] \leq \prod_{a \in Z_i} \mathbf{P}[u_i : a_i \notin \text{BR}(a^{-i}; u_i)] = 2^{-|Z_i|}.$$

Therefore,

$$\mathbf{P}[(u_1, \dots, u_n) \in \mathcal{U}] = \prod_{i=1}^n \mathbf{P}[u_i \in U_i] \leq \prod_{i=1}^n 2^{-|Z_i|} \leq 2^{-2^{n-1}},$$

since $\sum_i |Z_i| \geq |Z| = 2^{n-1}$. \square

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 the discussion following the statement of Theorem 7, or Section 1.2 in Kushilevitz and Nisan, 1997). \square

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 $\text{mag}(\rho) = \log |M| + \log |K|$, 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 $\text{mag}(U_i) = \max_{u_i \in U_i, a \in A} \text{mag}(u_i(a))$. For each $a \in A$, the *encoding* of the payoff of player i at a is $\text{enc}(U_i, a) = \log |\{u_i(a) : u_i \in U_i\}|$; 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 $\text{enc}(U_i) = \max_{a \in A} \text{enc}(U_i, a)$. 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 $\text{enc}(U_i) = 1$ bit, whereas its magnitude is $\text{mag}(U_i) = \log M$ bits. Finally, if $\mathcal{U} = U_1 \times \dots \times U_n$ is a family of games, then $\text{enc}(\mathcal{U}) = \max_{1 \leq i \leq n} \text{enc}(U_i)$ and $\text{mag}(\mathcal{U}) = \max_{1 \leq i \leq n} \text{mag}(U_i)$.

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., $\text{enc}(\mathcal{U}^n) = 1$ and $\text{mag}(\mathcal{U}^n) = O(n)$), such that any mixed Nash equilibrium procedure over \mathcal{U}^n has communication complexity $\Omega(2^n)$, i.e.,*

$$\text{CC}(\text{MIXED}, \mathcal{U}^n) = \Omega(2^n).$$

Our construction is based on a generalization of Jordan’s game (Jordan, 1993) 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 $i \neq 2$ (where $a_0 \equiv a_n$) 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). 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 player¹⁵ $i - 1$.

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

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

(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

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

be the average of the values of f ; equivalently, this is the expected value of f when every player $i \geq 3$ randomizes uniformly, i.e., $p_i = 1/2$.

Lemma 11. *The modified Jordan game $J_n(f)$ has a unique Nash equilibrium (p_1, \dots, p_n) , where $p_i = 1/2$ for all players $i \neq 1$, and¹⁶*

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

¹⁵ This game has a unique Nash equilibrium $(1/2, \dots, 1/2)$ (this also follows from Lemma 11 below).

¹⁶ Recall that p_i stands for the probability of action 1, i.e., $p_i = \mathbf{P}[a_i = 1]$.

Proof. Let (p_1, \dots, p_n) be a Nash equilibrium of $J_n(f)$. Assume that $p_i \neq 1/2$ for some player $i \geq 2$; without loss of generality, $p_i > 1/2$. Then player $i + 1$ must be playing purely the action 1, i.e., $p_{i+1} = 1$. Repeating this argument implies that $p_{i+2} = \dots = p_n = p_1 = 1$. Therefore player 2, who wants to mismatch player 1, must be playing the action 0 (here we use the fact that $0 \leq f < 1$), i.e., $p_2 = 0$. Hence player 3 must be playing 0, and so on; i.e., $p_3 = p_4 = \dots = p_i = 0$, contradicting $p_i > 1/2$.

So we have $p_i = 1/2$ for all $i \neq 1$. As for player 1, the probability p_1 is determined by the condition that player 2 must be indifferent between his two actions ($p_2 = 1/2$ and so player 2 plays both actions with positive probability), i.e.,

$$u_2^f(p_1, 0, 1/2, \dots, 1/2) = u_2^f(p_1, 1, 1/2, \dots, 1/2),$$

or, recalling the definitions of u_2^f and $\mu(f)$,

$$p_1 = p_1\mu(f) + 1 - p_1.$$

Solving for p_1 yields (3) and completes the proof. \square

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 $x = (x_1, \dots, x_{n-2})$ 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 $h : \{0, 1\}^{n-2} \rightarrow \{0, 1\}$. For every $h \in \mathcal{H}$, define a function f_h on $\{0, 1\}^{n-2}$ by

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

for each $x \in \{0, 1\}^{n-2}$, where $\text{prime}(k)$ is the k -th prime, starting for convenience with $\text{prime}(0) = 2$ (thus $\text{prime}(1) = 3$, $\text{prime}(2) = 5$, 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 $\mu(f_{h_1}) \neq \mu(f_{h_2})$.*

Proof. Let r_1, \dots, r_m be m distinct prime numbers, and let $\rho = \sum_{i=1}^m 1/r_i$. Express ρ as a reduced fraction $\rho = M/K$, and let $R = \prod_{i=1}^m r_i$. We claim that $K = R$. Indeed, on the one hand the common denominator in the sum ρ is R , so K divides R . On the other hand, multiplying ρ by $K \prod_{i=2}^m r_i$ gives

$$M \prod_{i=2}^m r_i = \frac{K \prod_{i=2}^m r_i}{r_1} + \sum_{i=2}^m K \prod_{j \neq 1, i} r_j;$$

therefore $K \prod_{i=2}^m r_i / r_1$ is an integer, and since the r_i are distinct primes, it follows that r_1 divides K . The same holds for all i , and so $\prod_{i=1}^m r_i = R$ divides K . Altogether $K = R$ (since each one divides the other), as claimed.

Now, for $\ell = 1, 2$,

$$2^{n-2} \mu(f_{h_\ell}) = \sum_{x \in \{0, 1\}^{n-2}} f_{h_\ell}(x) = \sum_{x: h_\ell(x)=1} \frac{1}{\text{prime}([x]_2)} = \frac{M_\ell}{K_\ell},$$

where M_ℓ/K_ℓ is a reduced fraction. If $\mu(f_{h_1}) = \mu(f_{h_2})$ then $K_1 = K_2$, or

$$\prod_{x:h_1(x)=1} \text{prime}([x]_2) = \prod_{x:h_2(x)=1} \text{prime}([x]_2),$$

which implies that the binary functions h_1 and h_2 are identical. \square

Next, the magnitude of $\mathcal{F}_{\mathcal{H}}$ is $O(n)$ bits, since $\text{prime}(k) = O(k \log k)$ by the Prime Number Theorem, and so $\log(\text{prime}([x]_2)) \leq \log(\text{prime}(2^{n-2})) = O(n)$; whereas the encoding of $\mathcal{F}_{\mathcal{H}}$ is just 1 bit, since $f_h(x)$ has only two possible values, $1/\text{prime}([x]_2)$ 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 $\text{enc}(\mathcal{U}) = 1$ and $\text{mag}(\mathcal{U}) = O(n)$.*

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. \square

6. Upper bounds

In this section we will show that two rather trivial procedures achieve near-optimal communication complexity. This demonstrates the power of our lower bounds, showing that even trivial procedures achieve a near-optimal performance.

We start with a pure Nash equilibrium procedure on binary-action games Γ_2^n .

Enumeration: Let Π_{ENUM} be the following procedure. All 2^n joint actions are examined in turn, with each player communicating whether or not he is best-replying there. Formally, at every period $t = 0, 1, \dots, 2^n - 1$, let $b \in \{0, 1\}^n$ be such that $[b]_2 = t$ (where $[b]_2 = \sum_{i=1}^n b_i 2^{n-i}$ is the integer corresponding to the binary string b); then player i broadcasts $a_{i,t} = 1$ if b_i is a best reply to b^{-i} (i.e., $u_i(b) \geq u_i(1 - b_i, b^{-i})$), and broadcasts $a_{i,t} = 0$ otherwise. If all players broadcasted 1 at a certain period t , then the procedure terminates at that point with the corresponding b (i.e., $[b]_2 = t$) as a pure Nash equilibrium of the game. Otherwise, the procedure ends at period 2^n with “no pure Nash equilibrium.”

Since every one of the n players communicates 1 bit in each one of the 2^n periods, we have

Proposition 15. *For every $n \geq 2$,*

$$\text{CC}(\text{PURE}, \Gamma_2^n) \leq \text{CC}(\Pi_{\text{ENUM}}, \Gamma_2^n) = n2^n.$$

For the next procedure, consider first the family of binary-action games that also have binary payoffs, i.e., $u_i(a) \in \{0, 1\}$ for all $a \in A = \{0, 1\}^n$. For every game G let $\bar{x}(G)$ be a Nash equilibrium of G . Think of \bar{x} as a predetermined and agreed-upon rule that selects a Nash equilibrium for each game (for example, among pure Nash equilibria select that one a that minimizes $[a]_2$; among mixed Nash equilibria, select that one (p_1, \dots, p_n) that minimizes the p_i lexicographically). The procedure is as follows.

Broadcast: Let Π_{BCAST} be the following procedure. The values of the payoff functions at each one of the 2^n joint actions are broadcast in turn, by all players. Formally, at every period $t = 0, 1, \dots, 2^n - 1$, let $b \in \{0, 1\}^n$ be such that $[b]_2 = t$; then player i communicates $a_{i,t} = u_i(b)$. At period $t = 2^n$ every player can reconstruct the payoff function of every other player: $u_i(b) = a_{i,[b]_2}$ for each $b \in \{0, 1\}^n$, and so they all know the game G , and thus the agreed-upon Nash equilibrium $\bar{x}(G)$.

The communication complexity of this procedure is again $n2^n$. Clearly, the same holds for any family of games $\mathcal{U}_1^n \subset \Gamma_2^n$ whose encoding is 1 bit (recall Section 5; the magnitude of \mathcal{U}_1^n does not matter and may be arbitrarily large),

whereas if the encoding is r bits, then the number of periods becomes $r2^n$ (every payoff entry takes r periods to transmit). So we have

Proposition 16. *For every $n \geq 2$ let $\mathcal{U}_r^n \subset \Gamma_2^n$ be a family of binary-action games whose encoding is at most r bits, i.e., $\text{enc}(\mathcal{U}_r^n) \leq r$. Then,*

$$\text{CC}(\text{MIXED}, \mathcal{U}_r^n) \leq \text{CC}(\Pi_{\text{BCAST}}, \mathcal{U}_r^n) \leq rn2^n.$$

7. Correlated equilibria

In this section we study the communication complexity of reaching a correlated equilibrium, and prove that it is polynomial rather than exponential in the number of players. 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 (2005) 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 $\mathcal{U}_u^n \subset \Gamma_2^n$ be the family of n -person binary-action games with integer payoffs of magnitude at most u bits, i.e., $\max_{1 \leq i \leq n} \text{mag}(u_i) \leq u$; our correlated equilibrium procedure will have a communication complexity that is polynomial in the number of players n and the magnitude of the payoffs u (for simplicity we again consider only binary-action games; otherwise, it would be polynomial in n, u , and $\max_{1 \leq i \leq n} |A_i|$).

We start by recalling the definition of a correlated equilibrium; see Aumann (1974). 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 $a = (a_1, \dots, a_n) \in A$ 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 $G = (u_1, \dots, u_n)$ that was given as input.¹⁷ Let CEP be the collection of correlated equilibrium procedures. Similarly to $\text{CC}(\text{MIXED}, \mathcal{G})$ and $\text{CC}(\text{PURE}, \mathcal{G})$, we define the *communication complexity of correlated equilibrium procedures* for a family of games \mathcal{G} as $\text{CC}(\text{CORRELATED}, \mathcal{G}) = \min_{\Pi \in \text{CEP}} \text{CC}(\Pi, \mathcal{G}) = \min_{\Pi \in \text{CEP}} \max_{G \in \mathcal{G}} \text{CC}(\Pi, G)$.

We come now to the construction of Papadimitriou (2005), which consists of running an ellipsoid algorithm in the Hart–Schmeidler setup (Hart and Schmeidler, 1989). 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 Papadimitriou (2005).

We define the procedure Π_{CORR} as follows. All players simulate the algorithm of Papadimitriou (2005). At each step of the ellipsoid algorithm, an n -tuple of mixed strategies $p = (p_1, \dots, p_n) \in \Delta = \prod_{i=1}^n \Delta_i$ 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 (2005) proves, first, that a correlated equilibrium is reached in a number of steps that is bounded by a polynomial in n and u ; and second, that the n -tuples of mixed strategies $p \in \Delta$ generated at every step have a magnitude $\text{mag}(p) = O(nu)$ bits. Therefore, when the payoffs $u_i(a)$ for all $a \in A$ are integers of at most u bits, the expected payoff $u_i(p)$ for $p \in \Delta$ requires at most $O(n \text{mag}(p) + u + n) = O(n^2u)$ bits (since it is a weighted sum of

¹⁷ Finite games always possess correlated equilibria.

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 u , and we have shown

Theorem 17. *For every $n \geq 2$ and $u \geq 1$, let $\mathcal{U}_u^n \subset \Gamma_2^n$ be the family of n -person binary-action games with integer payoffs of magnitude at most u , i.e., $\max_{1 \leq i \leq n} \text{mag}(u_i) \leq u$. Then there exists a correlated equilibrium procedure Π_{CORR} whose communication complexity over \mathcal{U}_u^n is polynomial in n and u , i.e.,*

$$\text{CC}(\text{CORRELATED}, \mathcal{U}_u^n) \leq \text{CC}(\Pi_{\text{CORR}}, \mathcal{U}_u^n) \leq \text{polynomial}(n, u).$$

In Appendix B we will 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; moreover, it is attained by simple procedures. We will also analyze the communication complexity of reaching correlated *approximate* equilibria.

8. Extensions

8.1. Stochastic procedures

For simplicity we have discussed only *deterministic* procedures; however, our results carry over to “stochastic” (or “randomized”) procedures. In an *equilibrium stochastic procedure*, the players may randomize when choosing their actions (i.e., $a_{i,t} \in A_i$ is chosen according to a probability distribution, which may depend on the history and on player i 's payoff function), and the procedure terminates in an appropriate equilibrium (pure, mixed, or correlated, as the case may be). The analysis of stochastic procedures is in terms of the *expected* communication complexity (i.e., the expectation of the number of bits transmitted until termination), which we denote CCS (where S stands for “stochastic”).

In general, the communication complexity of deterministic and stochastic protocols are polynomially related (see Kushilevitz and Nisan, 1997, Section 2.3), and so we expect the analysis of stochastic procedures also to lead to exponential bounds. In particular, the lower bound on the communication complexity of the \mathcal{S} -disjointness problem $\text{DISJ}_{\mathcal{S}}$ remains $\Omega(|\mathcal{S}|)$ also for stochastic protocols (see Kushilevitz and Nisan, 1997, Section 3.4), so the result of Theorem 3 holds also for pure Nash equilibrium *stochastic* procedures. As for Theorem 10, the counting argument of Section 5 applies to any procedure, whether deterministic or stochastic. See Appendix B for the analysis of some correlated equilibrium stochastic procedures.

8.2. Larger action spaces

We have mainly considered games in Γ_2^n where every player has 2 actions. What happens when the number of actions increases to $s > 2$? This has two effects. First, the space of joint actions is larger: s^n rather than 2^n ; and second, the number of communication bits per message is now $\log s$ rather than 1. Of course, as we noted at the end of Section 2.3, any lower bound for Γ_2^n is also a lower bound for Γ_s^n . However, most of our techniques turn out to be easily extendable to Γ_s^n , and to yield better bounds when $s > 2$. For instance, the matching pennies reduction (Section 3.2) gives a lower bound of $\Omega(s^n)$ when applied to Γ_s^n (recall Theorem 1: the disjointness problem for subsets of a set of size s^n has communication complexity s^n). As for the upper bounds, the enumeration and broadcast procedures Π_{ENUM} and Π_{BCAST} (Section 6) lead to a communication complexity of $O(ns^n \log s)$ and $O(rns^n)$, respectively.

8.3. 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 $p = (p_1, \dots, p_n) \in \prod_{i=1}^n \Delta_i = \Delta$ 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 would be interesting to study the communication complexity of reaching Nash approximate equilibria, and determine whether or not it is also exponential in the number of players.

Acknowledgments

We thank Fabrizio Germano, Adam Kalai, Eyal Kushilevitz, Andreu Mas-Colell, Noam Nisan, and the anonymous referees for useful discussions and suggestions, and the Institute for 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).

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 property¹⁸ 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 $i \in \{1, \dots, n/2\}$ 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 $i \in \{1, \dots, n/2\}$, and let the payoff functions be $u_{\ell,i}(a) = -d_H(z(a), S_\ell)$, for all $a \in \{0, 1\}^n$, $i \in \{1, \dots, n/2\}$, 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.¹⁹ 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 $y = \text{next}(x)$ and $x = \text{prev}(y)$; also, let $r(x) \in \{1, \dots, n/2\}$ denote the index of the unique bit in which x and $\text{next}(x)$ differ. Let

$$L = \{xx : x \in \{0, 1\}^{n/2}\} \cup \{yx : y, x \in \{0, 1\}^{n/2}, y = \text{next}(x)\},$$

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

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

($x_1 x_2$ is the concatenation of x_1 and x_2).

We will view L as a cycle that moves from each xx to yx , where $y = \text{next}(x)$, 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 $(1, r(x))$,

¹⁸ Besides reducibility and constructibility; see Section 3.1.

¹⁹ 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.

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

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

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 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 $GoodDone_{2,i}$

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

and then the payoff function

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

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 18. *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 $\text{GoodDone}_{r_2(x_1x_2)}(a) = 0$ and so the active 2-player $r_2(x_1x_2)$ 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. \square

Next, we will show that when $S_1 \cap S_2 \neq \emptyset$ all the improvement paths are finite.

Lemma 19. *Consider an improvement step from a to a' . If $z(a) = x_1x_2 \in L$, then $z(a') \in L$ and the improving player is either the active 1-player $r_1(x_1x_2)$ or the active 2-player $r_2(x_1x_2)$.*

Proof. Improvement steps cannot lead outside L : at a all payoffs are non-negative since $z(a) \in L$, whereas if $z(a') \notin L$ then all payoffs become negative. Only an active player can change his action bit such that $z(a') \in L$; as for the done bit of a player in T_2 , it affects his own payoff only when he is the active 2-player. \square

Lemma 20. *Consider an improvement step from a to a' . If $z(a) = x_1x_2 \in L$ and the improving player is the backward active 2-player, then he modifies only his done bit.*

Proof. A backward active 2-player exists only in the case where $x_1 = x_2 = x$; changing his action leads from xx to xy where $y = \text{prev}(x)$, and thus makes his payoff 0—which cannot be an improvement. \square

Lemma 21. *Consider an improvement path that starts with an improvement step by a player $(2, j)$ in T_2 from a with $z(a) \in L$ to a' . Then from a' until the next improvement step by a player from T_2 , if any, we have $\text{GoodDone}_{2,i} = 1$ for all players $(2, i) \in T_2$.*

Proof. We distinguish two cases according to the first improvement step. If $(2, j)$ modified only his done bit, then $z(a)$ was of the form xx , and $\text{GoodDone}_{2,j}$ changed from 0 to 1. So we are now at $z(a') = xx$ with $\text{GoodDone}_{2,i}(a') = 1$ for all players $(2, i) \in T_2$ (for player $(2, i)$, where $i \neq j$, this follows from the definition of $\text{GoodDone}_{2,i}$), and that will remain so as long as no player in T_2 moves (since an improvement step by a 1-player leads to yx where, again by definition, $\text{GoodDone}_{2,i} = 1$).

If the first improvement step by $(2, j)$ involved the action bit, then Lemma 20 implies that $(2, j)$ was the forward active 2-player, and so $z(a) = yx$ and $z(a') = yy$ (where $y = \text{next}(x)$). But unless $(2, j)$ also set the done bit correctly (to $\mathbf{1}_{\{x \in S_2\}}$) it could not have been an improvement step, and so $\text{GoodDone}_{2,j}(a') = 1$. From here on the argument is identical to the first case. \square

Corollary 22. *Consider an improvement path that starts at a with $z(a) \in L$. At each step after the first improvement step performed by some player from T_2 , all the players $(2, i)$ in T_2 have $\text{GoodDone}_{2,i} = 1$.*

Lemma 23. *Consider an improvement path that starts at a with $z(a) \in L$. Then there are no consecutive improvement steps performed by players from the same set T_ℓ , where $\ell \in \{1, 2\}$.*

Proof. By Lemma 19 only the active players can perform improvement steps, and the path never leaves L .

By way of contradiction, assume that two players in T_1 perform consecutive improvement steps. It must therefore be the same player moving back and forth between xx and yx (where $y = \text{next}(x)$), but then his payoff after the two consecutive steps does not change—a contradiction.

If two players in T_2 perform consecutive improvement steps, then, again, it must be the same player, and so his payoff must have first increased from 0 to 1, and then from 1 to 2. But then that second step was from an $xx \in L$ with $x \notin S_2$ (payoff of 1) to a $yy \in L$ with $y \in S_2$ (payoff of 2), which is impossible (since it involves two players modifying their actions). \square

Proposition 24. *Consider an improvement path. Then after at most $n + 2$ initial steps: (I) the path reaches L and never leaves it; (II) all the players $(2, i)$ in T_2 have $\text{GoodDone}_{2,i} = 1$; and either (IIIa) all improvement steps are performed by forward active players or (IIIb) a pure Nash equilibrium is reached.*

Proof. Every improvement step outside L decreases the distance to L , so after at most n steps we must reach L , and never leave it (Lemma 19). While in L the improvement steps alternate between T_1 and T_2 (Lemmata 19 and 23), so in at most 2 more steps (in total, no more than $n + 2$ steps from the start), an active player from T_2 performs an improvement step. By Corollary 22, from that time on $GoodDone_{2,i} = 1$ for all the players $(2, i)$ in T_2 . Therefore, in particular, no improvement step can change only the done bit, and so the players in T_2 can make only forward moves (Lemma 20). As for T_1 , if there is a backward move, then it goes from yx to xx and must increase the payoff of the active 1-player from 1 to 2. Therefore $x \in S_1$ and $d_{r_2(xx)} = 1$, which together with $GoodDone_{r_2(xx)} = 1$ implies that $x \in S_2$, and so the payoffs of all players equal 2, the maximal payoff—and a pure Nash equilibrium has been reached. \square

Proposition 25. *If $S_1 \cap S_2 \neq \emptyset$ then the game G generated by the potential game reduction has finite improvement paths.*

Proof. Assume by contradiction that an infinite improvement path exists. Proposition 24 implies that from some point on (IIIa) always holds, and so the improvement steps are all performed by active forward players. Therefore all states in L are traversed in turn, in particular xx for $x \in S_1 \cap S_2$. At that point the payoff of every player $(2, i)$ in T_2 is 2 (since $GoodDone_{2,i} = 1$ by (II) of Proposition 24); also, $d_{r_2(xx)} = \mathbf{1}_{\{x \in S_2\}} = 1$ (since $GoodDone_{r_2(xx)} = 1$), and so the payoff of every player in T_1 is 2 too, and no further improvement is possible. \square

We can now complete the proof of our result.

Proof of Theorem 4. Let $\mathcal{G} \subset \Gamma_4^n$ be the family of games that are obtained in our construction for all $S_1, S_2 \subset \mathcal{S} = \{0, 1\}^{n/2}$; Claims 18 and 5 imply that $CC(\text{PURE}, \mathcal{G}) = \Omega(|\mathcal{S}|) = \Omega(2^{n/2})$. By Proposition 25, every game $G \in \mathcal{G}$ either has the finite improvement path property, i.e., $G \in \mathcal{FIP}_4^n$ (when $S_1 \cap S_2 \neq \emptyset$) or G has no pure Nash equilibrium (when $S_1 \cap S_2 = \emptyset$). Now $CC(\text{PURE}, \mathcal{FIP}_4^n) \geq CC(\text{PURE}, \mathcal{G} \cap \mathcal{FIP}_4^n)$ (the communication complexity is defined as the worst case over all games in the family), and we claim that $CC(\text{PURE}, \mathcal{G} \cap \mathcal{FIP}_4^n) = CC(\text{PURE}, \mathcal{G})$. Indeed, take any pure Nash equilibrium procedure Π over $\mathcal{G} \cap \mathcal{FIP}_4^n$, and let its communication complexity be K bits. We can then use Π over the whole family \mathcal{G} , stopping it once K bits have been transmitted without a pure Nash equilibrium having been reached. Since all games in \mathcal{G} having pure Nash equilibrium lie in \mathcal{FIP}_4^n by Proposition 25, this is indeed a pure Nash equilibrium procedure for \mathcal{G} , and its communication complexity is also K . So

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

This establishes Theorem 4. \square

Appendix B. Correlated equilibria

In this appendix we show that, in the settings where we have obtained lower bounds on the communication complexity of pure and mixed Nash equilibrium procedures, there are simple procedures reaching a correlated equilibrium, whose communication complexity is low (polynomial, or even zero). Moreover, the support of the correlated equilibria that are reached is also low (the *support* of a correlated equilibrium Q is the number of joint actions $a \in A$ for which $Q(a) > 0$).

Finally, we study the communication complexity of reaching correlated approximate equilibria, and show that regret-minimization techniques yield simple polynomial upper bounds.

B.1. Modified Jordan games

We will exhibit a simple distribution that turns out to be a correlated equilibrium for *all* modified Jordan games, as defined in Section 5. For each $n \geq 3$, consider the following four joint actions in A :

$$\begin{aligned} z_{00} &= (0, 1, 0, 0, 0, \dots, 0), \\ z_{01} &= (0, 1, 1, 0, 0, \dots, 0), \\ z_{10} &= (1, 0, 0, 1, 1, \dots, 1), \end{aligned}$$

$$z_{11} = (1, 0, 1, 1, 1, \dots, 1)$$

(i.e., $z_{bc} = (b, 1 - b, c, b, \dots, b)$ for $b, c \in \{0, 1\}$), and let Q_n be the distribution that assigns equal probability of $1/4$ to each one of these four joint actions.

Lemma 26. *The distribution Q_n is a correlated equilibrium of the modified Jordan game $J_n(f)$ for every $n \geq 3$ and every $f : \{0, 1\}^{n-2} \rightarrow [0, 1)$.*

Proof. For each a in the support of Q_n , i.e., $a \in \{z_{00}, z_{01}, z_{10}, z_{11}\}$, we have the following: every player $i \neq 2, 3, 4$ always matches his predecessor, and so he gets his highest payoff of 1 and will not deviate; also, player 2 always mismatches player 1, and so he will not deviate either (we use here $0 \leq f < 1$). As for players $i = 3, 4$, when $a_i = 0$ it is equally likely (according to Q_n) that $a_{i-1} = 0$ and $a_{i-1} = 1$, and so the payoff of i is $1/2$ and a deviation to $b_i = 1$ does not change his payoff; the same holds when $a_i = 1$. \square

Since there is a correlated equilibrium that does not depend on the specific modification f , no communication is needed to reach it. Therefore we have established the following:

Theorem 27. *For every $n \geq 3$, the communication complexity of reaching a correlated equilibrium over the family of binary-action games $\mathcal{U}^n \subset \Gamma_2^n$ of Theorem 10 is zero, i.e.,*

$$CC(\text{CORRELATED}, \mathcal{U}^n) = 0.$$

B.2. Bayesian setting

Recall the Bayesian setting of Section 4. The probability distribution \mathbf{P} over the family Γ_2^n of binary-action games is obtained by putting $u_i(0, a^{-i}) = 0$ and $u_i(1, a^{-i}) = 1$ with probability $1/2$, and $u_i(0, a^{-i}) = 1$ and $u_i(1, a^{-i}) = 0$ with probability $1/2$, independently over all i and a^{-i} . For each player i and action $b = 0, 1$, let $B_{i,b} := \{a \in A : a_i = b, u_i(a) = 1\}$ be the set of joint actions where player i plays the action b , and that is his best-reply action.

We start with a preparatory lemma. Let $m = 4n^2$. After the payoff functions $\{u_i\}_i$ have been chosen (according to \mathbf{P}) and every player i has been informed of his own u_i , let player i select a random subset of size m of $B_{i,0}$, denote it $S_{i,0}$, and a random subset of size m of $B_{i,1}$, denote it $S_{i,1}$, with all joint actions in each $B_{i,b}$ equally likely to be selected. Put $S_i = S_{i,0} \cup S_{i,1}$ and $S = \bigcup_i S_i$; then $S \subset A$ is a random set containing at most $2mn$ joint actions (the same joint action may be selected by different players). Let ζ_S denote the uniform distribution on S (i.e., $\zeta_S(s) = 1/|S|$ for every $s \in S$).

Lemma 28. *ζ_S is a correlated equilibrium in the game (u_1, \dots, u_n) with probability at least $1 - 2ne^{-n}$.*

Note that there are two randomizations here: first, the payoff functions u_i (according to \mathbf{P}), and second, the selections S_i (conducted after each player i knows his u_i); we will write \mathbf{Pr} for the resulting joint probability on both payoffs and selections.

Proof. In order for ζ_S to be a correlated equilibrium, the $2n$ inequalities

$$\sum_{a \in S: a_i=b} [u_i(b, a^{-i}) - u_i(1 - b, a^{-i})] \geq 0, \tag{B.1}$$

for all i and $b = 0, 1$, must be satisfied. Fix a player i and an action $a_i = b$ in $\{0, 1\}$, and let $T = \{a \in S : a_i = b\}$. For each $a \in A$ such that $a_i = b$, let $X_a = u_i(b, a^{-i}) - u_i(1 - b, a^{-i})$; then $X_a \in \{1, -1\}$, and we want to obtain an upper bound on the probability that $\sum_{a \in T} X_a < 0$, i.e., the corresponding inequality (B.1) is violated.

If $a \in T$ was selected by player i , i.e., $a \in S_i$, then necessarily $a \in S_{i,b}$ and so $X_a = 1$ (since $S_{i,b} \subset B_{i,b}$). If $a \notin S_i$, then a was selected by some other player $j \neq i$, and then it is equally likely that $X_a = 1$ and $X_a = -1$ (since player i 's payoff and player j 's selection are independent). Let K be the number of elements in $T \setminus S_i$; then $0 \leq K \leq 2m(n - 1)$ and $\sum_{a \in T} X_a = m + \sum_{a \in T \setminus S_i} X_a = m + \sum_{k=1}^K Y_k$, where the Y_k are independent random variables with $\mathbf{Pr}[Y_k = 1] = \mathbf{Pr}[Y_k = -1] = 1/2$. Applying Hoeffding's inequality (Hoeffding, 1963) yields $\mathbf{Pr}[\sum_{a \in T} X_a <$

$0 | K] = \Pr[\sum_{k=1}^K Y_k < -m] \leq e^{-m^2/(2K)} \leq e^{-n}$ for every $K \leq 2m(n - 1)$ (recall that we took $m = 4n^2$), and so $\Pr[\sum_{a \in T} X_a < 0] \leq e^{-n}$.

The same computation applies to all i and $b = 0, 1$, and so the probability that at least one of the $2n$ inequalities (B.1) is violated is at most $2ne^{-n}$. \square

Consider now the following stochastic procedure Π_{BAYES} . In the first stage, each player i selects a random set $S_i = S_{i,0} \cup S_{i,1}$ of joint actions as above, and then broadcasts it. In the second stage, each player i verifies whether the resulting $S = \bigcup_{i=1}^n S_i$ satisfies his two inequalities (B.1) (only u_i matters here), and then broadcasts the result, using one bit. If all inequalities are satisfied, then ζ_S is a correlated equilibrium and the procedure terminates. If not, then in the third stage all players broadcast their complete payoff functions, and then they each compute a correlated equilibrium using the same algorithm, so as to get the same result. Since Π_{BAYES} is a stochastic procedure, for each game G we consider the expected communication complexity $\text{CCS}(\Pi_{\text{BAYES}}, G)$ (see Section 8.1).

Theorem 29. *The expected communication complexity of the correlated equilibrium stochastic procedure Π_{BAYES} is $O(n^4)$, i.e.,*

$$\mathbf{E}[\text{CCS}(\Pi_{\text{BAYES}}, \mathcal{G})] = O(n^4),$$

where the expectation \mathbf{E} is according to the probability distribution \mathbf{P} of Theorem 7.

Proof. Each one of the n sets S_i has $2m = 8n^2$ elements of n bits each, and so the first stage has a communication complexity of $8n^4$ bits. Then each player sends a single bit, which adds n bits of communication. With probability at most $2ne^{-n}$ the resulting ζ_S is not a correlated equilibrium, and then there is an additional communication complexity of $n2^n$ of broadcasting the payoff functions. Altogether, $\mathbf{E}[\text{CCS}(\Pi_{\text{BAYES}}, \mathcal{G})] = 8n^4 + n + 2ne^{-n}n2^n \leq 8n^4 + n + 2n^2 = O(n^4)$. \square

B.3. Correlated approximate equilibria

We will now show that “regret-minimization” procedures that reach correlated approximate equilibria have a low communication complexity. Let $\varepsilon > 0$; a *correlated ε -equilibrium* is defined in the same way as a correlated equilibrium, except that now all the inequalities must hold only within ε .

Theorem 30. *For every $n \geq 2$ and $u \geq 1$, let $\mathcal{U}_u^n \subset \Gamma_2^n$ be the family of n -person binary-action games with payoffs of magnitude at most u , i.e., $\max_{1 \leq i \leq n} \text{mag}(u_i) \leq u$. Then for every $\varepsilon > 0$ there exists a correlated ε -equilibrium stochastic procedure $\Pi_{\varepsilon\text{-CORR}}$ whose expected communication complexity over \mathcal{U}_u^n is polynomial in n and u , i.e.,*

$$\text{CCS}(\Pi_{\varepsilon\text{-CORR}}, \mathcal{U}_u^n) = O\left(\frac{n^2}{\varepsilon^2} + nu\right).$$

Proof. For every $\varepsilon, \delta > 0$ there exists a regret-minimization procedure that guarantees that, after $T = C \log(n/\delta)/\varepsilon^2$ periods, there is a probability of at least $1 - \delta$ that the time average of the played joint actions constitutes a correlated ε -equilibrium (where C is an appropriate constant); e.g., see Cesa-Bianchi and Lugosi (2006, Section 7.4). Let Π_{RM} be such a procedure for $\delta = 2^{-n}$; its communication complexity is $nT = O(n^2/\varepsilon^2)$, since in each one of the T periods every player’s action corresponds to one bit of communication.

The procedure $\Pi_{\varepsilon\text{-CORR}}$ starts by running Π_{RM} ; let $a^t = (a_{1,t}, \dots, a_{n,t}) \in A$ be the joint action at time t . When Π_{RM} terminates, after T periods, each player i computes his average regrets and checks whether they are both at most ε (i.e., $(1/T) \sum_{t=1}^T [u_i(b, a^{-i,t}) - u_i(a^t)] \leq \varepsilon$ for²⁰ $b = 0, 1$); he then broadcasts the result using one bit. With probability at least $1 - 2^{-n}$ all average regrets are at most ε , and then the uniform distribution on (a^1, \dots, a^T) is a correlated ε -equilibrium and $\Pi_{\varepsilon\text{-CORR}}$ terminates.

Otherwise, all players broadcast their complete payoff functions, after which they all compute a correlated equilibrium using the same algorithm. The communication at this stage requires $O(n2^n u)$ bits, but it only happens with

²⁰ For binary-action games, conditional (“internal”) and unconditional (“external”) regrets are the same.

probability at most 2^{-n} . The total expected communication complexity is therefore $O(n^2/\varepsilon^2) + n + 2^{-n}O(n2^n u) = O(n^2/\varepsilon^2 + nu)$. \square

References

- Aumann, R.J., 1974. Subjectivity and correlation in randomized strategies. *J. Math. Econ.* 1, 67–96.
- Blum, A., Mansour, Y., 2007. From external to internal regret. *J. Machine Learning Res.* 8, 1307–1324.
- Cahn, A., 2004. General procedures leading to correlated equilibria. *Int. J. Game Theory* 33, 21–40.
- Cesa-Bianchi, N., Lugosi, G., 2003. Potential-based algorithms in on-line prediction and game theory. *Machine Learning* 51, 239–261.
- Cesa-Bianchi, N., Lugosi, G., 2006. *Prediction, Learning and Games*. Cambridge Univ. Press.
- Conitzer, V., Sandholm, T., 2004. Communication complexity as a lower bound for learning in games. In: Brody, C.E. (Ed.), *Machine Learning, Proceedings of the Twenty-First International Conference (ICML 2004)*. ACM, pp. 185–192.
- Foster, D., Kakade, S.M., 2004. Deterministic calibration and Nash equilibrium. In: Shawe-Taylor, J., Singer, Y. (Eds.), *Learning Theory, 17th Annual Conference on Learning Theory (COLT 2004)*. Springer, pp. 33–48.
- Foster, D., Vohra, R., 1997. Calibrated learning and correlated equilibrium. *Games Econ. Behav.* 21, 40–55.
- Foster, D., Young, H.P., 2003. Learning, hypothesis testing, and Nash equilibrium. *Games Econ. Behav.* 45, 73–96.
- Foster, D., Young, H.P., 2006. Regret testing: Learning to play Nash equilibrium without knowing you have an opponent. *Theoret. Econ.* 1, 341–367.
- Fudenberg, D., Levine, D.K., 1998. *The Theory of Learning in Games*. MIT Press.
- Germano, F., Lugosi, G., 2007. Global Nash convergence of Foster and Young’s regret testing. *Games Econ. Behav.* 60, 135–154.
- Hart, S., 2005. Adaptive heuristics. *Econometrica* 73, 1401–1430.
- Hart, S., Mas-Colell, A., 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68, 1127–1150.
- Hart, S., Mas-Colell, A., 2001. A general class of adaptive strategies. *J. Econ. Theory* 98, 26–54.
- Hart, S., Mas-Colell, A., 2003. Uncoupled dynamics do not lead to Nash equilibrium. *Amer. Econ. Rev.* 93, 1830–1836.
- Hart, S., Mas-Colell, A., 2006. Stochastic uncoupled dynamics and Nash equilibrium. *Games Econ. Behav.* 57, 286–303.
- Hart, S., Schmeidler, D., 1989. Existence of correlated equilibria. *Math. Operations Res.* 14, 18–25.
- Hoeffding, W., 1963. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58, 13–30.
- Hurwicz, L., 1960. Optimality and informational efficiency in resource allocation processes. In: Arrow, K.J., Karlin, S., Suppes, P. (Eds.), *Mathematical Models in the Social Sciences*. Stanford Univ. Press.
- Jordan, J., 1991. Bayesian learning in normal form games. *Games Econ. Behav.* 3, 60–81.
- Jordan, J., 1993. Three problems in learning mixed equilibria. *Games Econ. Behav.* 5, 368–386.
- Jordan, J., in press. Information aggregation and prices. In: Durlauf, S., Blume, L. (Eds.), *The New Palgrave Dictionary of Economics*, second ed. Palgrave Macmillan.
- Kalai, E., Lehrer, E., 1993. Rational learning leads to Nash equilibrium. *Econometrica* 61, 1019–1045.
- Kushilevitz, E., Nisan, N., 1997. *Communication Complexity*. Cambridge Univ. Press.
- Monderer, D., Shapley, L.S., 1996. Potential games. *Games Econ. Behav.* 14, 124–143.
- Papadimitriou, C.H., 2005. Computing correlated equilibria in multi-player games. In: Gabow, H.N., Fagin, R. (Eds.), *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*. ACM, pp. 49–56.
- Stoltz, G., Lugosi, G., 2005. Internal regret in on-line portfolio selection. *Machine Learning* 59, 125–159.
- Stoltz, G., Lugosi, G., 2007. Learning correlated equilibria in games with compact sets of strategies. *Games Econ. Behav.* 59, 187–208.
- Young, H.P., 2004. *Strategic Learning and its Limits*. Oxford Univ. Press.