An invitation to online information transmission
PRELIMINARY VERSION

Olivier Gossner[*]   Penélope Hernández[†]   Abraham Neyman[‡]

January 29, 2003

Abstract

We study a class of games that models costly information transmission in long term interactions. The signaling is restricted to the play of the game and thus have payoff consequences. The auxiliary game: A sequence of actions X_t is chosen by nature, and announced to a player, the prophet, but not to his teammate, the follower. A repeated game then proceeds in which the prophet and the follower choose actions, Y_t and Z_t at stage t. The stage payoff to the team at stage t depends on the triple (X_t, Y_t, Z_t).

For these games, we characterize the best payoff achievable to the team, both against a distribution over sequences, and against the worst possible case. The team faces a trade-off problem between stage payoff maximization and information transmission, and the optimal payoff is achieved for a corresponding optimal level of communication.

Our result has direct implications for repeated games played by finite automata, or by players with bounded recall.

1 Introduction

A common phenomenon in dynamic interactions is that information about future realized variables differs across agents. The importance of the impact of these informational asymmetries on the equilibrium outcomes of strategic situations is already acknowledged by a wide literature.

[*]THEMA, Université Paris 10 – Nanterre and CORE, Université Catholique de Louvain

[†]CORE, Université Catholique de Louvain

[‡]Institute of Mathematics and Center for the Study of Rationality Theory, Hebrew University of Jerusalem.

1

In particular, information asymmetries can be a source of inefficiencies if agents wish to coordinate their actions, which is likely if these agents form part of the same larger entity (such as a team, or a firm). Communication between some or all of the agents can impact the distribution of information among them, and can in some circumstances resolve inefficiencies due to (ex-ante) asymmetries. In some cases, communication between a group of agents can be beneficial for these agents, but not to society as a whole. For instance some firms could use communication to enforce collusion, or to choose a technological standard excluding the rest of the producers.

Communication may take place though a number of instruments which vary in their characterisitics. It may be free, or virtually free, or on the contrary be costly. It may also be limited in its quantity, by some bandwidth for instance, or be virtually unlimited. When society chooses to forbid communication between a group of agents, a regulator is in charge of – as much as it can – prevent such communication from taking place. Nevertheless, even under the strongest laws, it would be unrealistic to try to physically make impossible to decision makers to exchange information through phone lines or electronic means. The fact that such explicit means of transmitting data are forbidden makes it only more costly to the agents, through the risk they take of being caught. Even without using such explicit, and risky lines, decision makers may signal to each other part of their information though regular course of business, like quoting prices. After all, the agent's choices to exchange messages or not and that such possibilities can rarely be arbitrarily precluded *per se*.

We consider the problem faced by a team in which one player acquires more information than another about future realizations of some variable. The more informed agent, called the prophet, can transmit part of this extra information to his teammate, called the follower, through his choices of actions.

More precisely, a sequence of actions $(X_t)_t$ is chosen by nature, and announced to the prophet, but not to the follower. A repeated game then proceeds in which the follower and the prophet choose actions, Y_t and Z_t at stage t. The stage payoff to the team is given by $g(X_t, Y_t, Z_t)$.

Through his choice of actions, the prophet seeks to maximize stage payoffs, and to signal as much information as possible to the follower. This trade-off between payoff maximization and information is at the heart of our analysis.

For instance, in the most simple case in which g is independent of Z, the optimal strategy to the prophet is to choose the action Y_t that maximizes g against X_t, for each stage t.

Assume now that g depends on Z, but is independent of Y. In this case, the prophet's best strategy is to use his actions Y_t in such a way that the information they convey about future values of X is maximal.

For general games, we characterize the best payoff achievable to the team, and the corresponding rate of information transmission within the team. Two definitions of the best achievable payoff can be used:

In a first approach, one can assume that the team chooses its strategies in order to maximize expected payoff against a known distribution over sequences (X_t). The best payoff to the team then corresponds to the distribution against which the worst expected payoff can be obtained.

A complementary approach is the worst case analysis. Here, one estimates the payoff that the team can obtain using some strategies by the minimal payoff these strategies obtain over all possible sequences.

Our result shows that when the game is played for a sufficiently large number of stages, these two approaches yield the same value.

Information theory, initiated by the work of Shannon (1948), is concerned by the question maximizing information transmission for a given technology. Although our proofs rely heavily on these techniques, our model includes new cost considerations that are important to economics.

In section 2, we introduce the model of online communication between the team members. In section 3, we present the solution for a particular example of online matching pennies. Section 4 is devoted to the analysis of general games. Our result has direct implications for repeated games played by finite automata, or by players with bounded recall, these are presented in section 5.

2 General Model

2.1 The one-shot game

We consider a three player game G. The finite action sets for player 1, 2, and 3 are denoted I, J, and K. The (common) payoff function to players 2 and 3 is $g \colon I \times J \times K \to \mathbb{R}$. Player 2 is called the prophet, player 3 the follower, and players 2 and 3 form the team.

2.2 The repeated game

In the repeated game, the prophet has knowledge in advance of the actions of player 1. Hence, a (pure) strategy for the prophet can be represented by a mapping $Y \colon I^{\mathbb{N}} \to J^{\mathbb{N}}$ with coordinates $(Y_n)_n$. The follower observes the

past actions of player one and of the prophet. So, a (pure) strategy Z for the follower is a mapping $Z\colon I^{\mathbb{N}} \times J^{\mathbb{N}} \to Z^{\mathbb{N}}$ with coordinates $(Z_n)_n$ such that, Z_n depends on the actions of players 1 and 2 from stage 1 to $n-1$.

Given a sequence $X \in I^{\mathbb{N}}$ and strategies Y, Z for the team, the induced sequences of actions $(y_n)_n$ and $(z_n)_n$ of the prophet and the follower are given by the relations: $(y_n)_n = Y(X)$, $(z_n)_n = Z(X, Y)$. Any probability P on $I^{\mathbb{N}}$ together with strategies (Y, Z) induces a probability distribution $P_{Y,Z}$ on the set of sequences (x_t, y_t, z_t) in $(I \times J \times K)^{\mathbb{N}}$.

2.3 Questions

We address the following questions:

- What is the best expected payoff that the team can guarantee against a *fixed* distribution over sequences of player 1?

- What is the best payoff that the team can guarantee against *all* sequences of player 1?

We therefore introduce the following notations:
Given a probability ρ on $n \in \mathbb{N}$ and I^n, let

$$v_n(\rho) = \max_{Y,Z} \frac{1}{n} \mathbf{E}_{\rho_{Y,Z}} \sum_{t=1}^{n} g(x_t, y_t, z_t)$$

and let $v_n = \min_{\rho} v_n(\rho)$.

Hence, assuming the distribution over sequences $i.i.d$, and X ρ, $v_n(\rho)$ is the best expected payoff the team can obtain in the n-stage game.

Fixed strategies Y, Z of the team guarantee a payoff $\min_{(x_t)_t} \frac{1}{n} \sum_{t=1}^{n} g(x_t, y_t, z_t)$ against all sequences in the n-stage version. A prudent choice of strategies for the team would maximize this minimum payoff. Hence, given $n \in \mathbb{N}$, we let

$$w_n = \max_{Y,Z} \min_{(x_t)_t} \frac{1}{n} \sum_{t=1}^{n} g(x_t, y_t, z_t)$$

Since a strategy that guarantees w_n against all sequences of player 1 also guarantees w_n against any distribution of sequences, $w_n \le v_n(\rho)$ for any ρ and therefore $w_n \le v_n$.

Notice that the $\min\max$ theorem does not apply to this game since there are three players and correlated strategies are not available to the team.

3 Online matching pennies

3.1 The payoff function

In this section, we assume G is a 3-player game of matching pennies. Players 1, 2, and 3 choose i, j, and k in $\{0, 1\}$, and the payoff to players 2 and 3 is given by:

$$g(i, j, k) = \begin{cases} 1 & if \ \ i = j = k \\ 0 & \text{otherwise} \end{cases}$$

and can be represented by the payoff matrices:

1	0
0	0

0	0
0	1

where player 1 chooses the row, player 2 chooses the column, and player 3 chooses the matrix.

3.2 Candidate strategies

We present two strategies for the team and analyze their payoffs both against a fixed distribution ($\rho = (1/2, 1/2)$ i.i.d.) and for the worst case.

Example 1: Consider the strategies given by $y = x$ for the prophet and by an arbitrary sequence of actions $z = (z_1, \ldots, z_n)$ for the follower .

- Against ρ, the average expected payoff is 0.5.

- The worst possible sequence is $x = (1 - z_1, \ldots, 1 - z_n)$, and the corresponding payoff is 0.

Example 2: Assume the prophet plays on odd stages the next action of Player 1 and on even stages the follower and the prophet play the previous action of the the prophet. The follower plays an arbitrary sequence of actions on the odd stages.

The resulting sequences of actions is:

$$\begin{aligned} x &= (x_1, x_2, x_3, x_4, \ldots, x_n) \\ y &= (x_2, x_2, x_4, x_4, \ldots, x_n) \\ z &= (z_1, x_2, z_3, x_4, \ldots, x_n) \end{aligned}$$

- Against a sequence distributed according to ρ, the team obtains:

- 1 at even stages;
- an expected payoff of $\frac{1}{4}$ at odd stages;
- resulting in an average expected payoff is 0.625.

- Against the worst possible case, the payoffs are:
  - 1 at even stages;
  - 0 at odd stages;
  - resulting in an average payoff of 0.5.

The strategies in the example 1 do not involve any communication and this is the best payoff that the team in this case. Nevertheless, in the example 2, the communication at odd stages improves not only the expected payoff against ρ but the resulting payoff for the worst case .

3.3 The value

Throughout the paper, log denotes the logarithm function in basis 2.

Let v^* be the (unique) solution of the equation

$$H(x) + (1 - x) \log 3 = 1$$

where $H\colon [0,1] \to \mathbb{R}$ given by $H(x) = -x \log x - (1 - x) \log(1 - x)$ for $x \notin \{0,1\}$, $H(0) = H(1) = 0$ is the entropy function.

Whenever $n \in \mathbb{N}$ and μ is a distribution over a finite set I, $\mu^{\otimes n}$ denotes the distribution over I^n of a family $(x_i)_{1 \le i \le n}$ of i.i.d. random variables such that $x_i \sim \mu$.

Theorem 1 *The game of online matching pennies has value v^* in the following sense:*

1. *For $n \in \mathbb{N}$ and $\rho = (\frac{1}{2}, \frac{1}{2})^{\otimes n}$,*

$$v_n(\rho) \le v^*$$

2.

$$\lim_{n \to \infty} w_n = v^*$$

3.4 Notations and tools

We introduce the tools needed for the proof of theorem 1.

3.4.1 General notations

Given a finite set (or a measurable space) X we denote by $\Delta(X)$ the set of probability measures on X.

For $z \in \mathbb{R}$, we let $[z]$ and $\lceil z \rceil$ denote the integer part and the superior integer part of z respectively ($z - 1 < [z] \le z$ and $z \le \lceil z \rceil < z + 1$). Given a finite set Z, $|Z|$ denotes the cardinality of Z.

Given two sequences $a = (a_n)_n$ and $b = (a_n)_n$ of positive numbers, we write $a \doteq b$ whenever $\lim_{n \to \infty} \frac{\log a_n - \log b_n}{n} = 0$.

3.4.2 Entropy and conditional entropy

Let X be a random variable over a finite set Θ with distribution p. The entropy $H(X)$ of X is

$$H(X) = -\Sigma_{\theta \in \Theta} p(\theta) \log p(\theta) = -E_X \log p(X)$$

where $0 \log 0 = 0$ (by convention $\log$ is taken in basis 2). The entropy of a random variable depends on its distribution only. Thus, for $p \in \Delta(\Theta)$ we let $H(p) = H(X)$ for a random variable X with distribution p. By convention, if $p \in [0, 1]$, $H(p)$ also represents the entropy of a Bernoulli random variable of parameter p.

Given a pair of random variables (X_1, X_2) taking values in $\Theta_1 \times \Theta_2$ with joint distribution $p(\theta_1, \theta_2)$, we denote by $p(\theta_2 \mid \theta_1)$ the conditional probability that $X_2 = \theta_2$ given that $X_1 = \theta_1$. Define $h(X_2 \mid \theta_1) = -\Sigma_{\theta_2 \in \Theta_2} p(\theta_2 \mid \theta_1) \log p(\theta_2 \mid \theta_1)$. Thus $h(X_2 \mid \theta_1)$ is the entropy of X_2 when the realization $X_1 = \theta_1$ is known.

The conditional entropy $H(X_2 \mid X_1)$ of X_2 given X_1 is

$$H(X_2 \mid X_1) = E_{X_1}\left[h(X_2 \mid X_1)\right] = \sum_{\theta_1 \in \Theta_1} p(\theta_1) h(X_2 \mid \theta_1)$$

Direct computation shows that $H(X_1, X_2) = H(X_1) + H(X_2 \mid X_1)$. This extends to a family of random variables $(X_1, \ldots, X_n)$ to:

$$H(X_1, \ldots, X_n) = H(X_1) + \sum_{k=2}^{n} H(X_k \mid X_1, ..., X_k - 1)$$

3.4.3 Hamming distance

2^n stands for the sequences of zeroes and ones of length n. $\mathbb{I}$ stands for the indicator function. The Hamming distance between two sequences $x, y \in 2^n$ is denoted $d_H(x, y)$ $(= \sum_{t=1}^{n} \mathbb{I}(x_t \ne y_t))$.

We shall rely on the following bound on the size of a sphere of size i centered at $x \in 2^n$.

Proposition 1 *Given $i, n \in \mathbb{N}$, $i \leq n$ and $x \in 2^n$:*

$$|\{y, d_H(x, y) = i\}| = \binom{n}{i} \geq \frac{2^{nH(\frac{i}{n})}}{\sqrt{2n}}$$

Proof. The first equality is obvious, and the second follows directly from classical bounds. (See e.g. [22] p.309, 310.) ∎

3.5 Play against a distribution

We assume the distribution ρ of sequences of players 1 known to players 2 and 3, and obtain a bound on the best response payoff to the team. Then we derive a proof for part 1 of theorem 1.

We first provide a bound on the payoff that the team can obtain facing a sequence of per stage entropy h.

Proposition 2 *Let $\rho \in \Delta(2^n)$, and $h = \frac{1}{n}H(\rho)$*

$$h \leq H(v_n(\rho)) + (1 - v_n(\rho)) \log 3$$

Proof. Assume thus that $X = (X_1, \ldots, X_n)$ is drawn according ρ, and that $Y : 2^n \to 2^n$ and $Z : 2^{\mathbb{N}} \times 2^{\mathbb{N}} \to 2^{\mathbb{N}}$ are pure strategies of player 2 and player 3 respectively.

Let $\mathcal{F}_t$ denote the algebra of events spanned by the random variables $X_1, Y_1, \ldots, X_t, Y_t$. We define

$$g_t = E_\rho \left(\mathbb{I}(X_t = Z_t = Y_t) \mid \mathcal{F}_{t-1} \right)$$

Thus, g_t is the $\mathcal{F}_{t-1}$-measurable random variable that represents the expected payoff for the team at stage t given the past actions.

Note that Z_t being $\mathcal{F}_{t-1}$-measurable, the triple (X_t, Y_t, Z_t) may take only 4 values conditional to $\mathcal{F}_{t-1}$, as represented in the following tree:

8
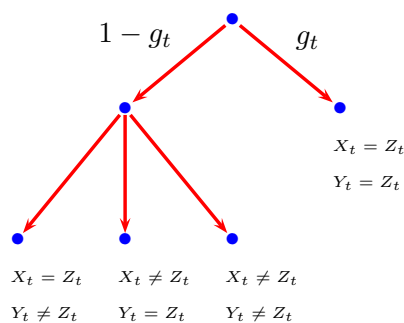
Hence, we deduce that:

$$h(X_t, Y_t \mid \mathcal{F}_{t-1}) \leq H(g_t) + (1 - g_t) \log 3$$

Taking expectations over histories in $\mathcal{F}_{t-1}$ yields:

$$H(X_t, Y_t \mid X_1, Y_1, \ldots, X_{t-1}, Y_{t-1}) \leq \mathbf{E}_\rho(H(g_t) + (1 - g_t) \log 3)$$

Summing over t now gives:

$$H(X_1, Y_1, \ldots, X_n, Y_n) \leq \mathbf{E}_\rho \sum_{t=1}^n (H(g_t) + (1 - g_t) \log 3)$$

Note also that:

$$
\begin{aligned}
H(X_1, Y_1, \ldots, X_n, Y_n) &= H(X_1, \ldots, X_n) + H(Y_1, \ldots, Y_n \mid X_1, \ldots, X_n) \\
&= H(X_1, \ldots, X_n) \\
&= nh
\end{aligned}
$$

since $(Y_1, \ldots, Y_n)$ is a function of $(X_1, \ldots, X_n)$. Hence,

$$h \leq \mathbf{E}_\rho \frac{1}{n} \sum_{t=1}^n (H(g_t) + (1 - g_t) \log 3)$$

We now apply twice Jensen's inequality to the concave mapping $x \mapsto H(x) + (1 - x) \log 3$ and obtain
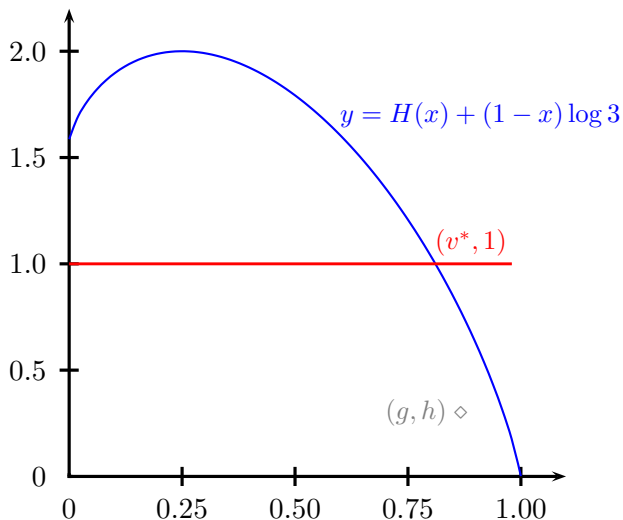
$$
\begin{aligned}
h &\leq \mathbf{E}_\rho(H(\frac{1}{n} \sum_{t=1}^n g_t) + (1 - \frac{1}{n} \sum_{t=1}^n g_t) \log 3) \\
&\leq H(\mathbf{E}_\rho \frac{1}{n} \sum_{t=1}^n g_t) + (1 - \mathbf{E}_\rho \frac{1}{n} \sum_{t=1}^n g_t) \log 3
\end{aligned}
$$

Hence, the expected payoff to the team $g = \mathbf{E}_\rho \frac{1}{n} \sum_{t=1}^n g_t$ verifies

$$h \leq H(g) + (1 - g) \log 3$$

The following figure shows the curve of equation $y = H(x) + (1-x) \log 3$. We see that the x-coordinate of the intersection between this curve and the straight line $y = h$ is minimal when $h = 1$ $(0 \leq h \leq 1)$, and equals v^*. From

9

this we deduce part 1 of theorem 1.

3.6 Play against all sequences

We now design (ε)-optimal strategies for the team against all sequences of player 1, and show that they can guarantee any payoff close to v^*. This will imply part 2 of theorem 1.

Let $x < v^*$, and $\eta = H(x) + (1-x)\log 3 - 1 > 0$. We now construct strategies for the team that (for sufficiently large n) guarantee x against *all* sequences.

Let $p = \frac{1-x}{1+2x}$ and $q = \frac{2}{3}(1-x)$. The strategies are defined over blocks of length n in such a way that in any block after the first, and for any sequence X, the proportion of stages for which $Z_t \neq X_t$ is close to q, and the proportion of stages for which $Y_t \neq X_t$ conditional on $Z_t = X_t$ is close to p. The proportion of stages in which $Z_t = Y_t = X_t$ is then close to $(1-p)(1-q) = x$.

During each block (after the first), the follower has to interpret the message sent by the prophet during the previous block in order to choose a sequence of actions $\tilde{Z}$. This sequence of actions should be such that it matches $\lceil (1-q)n \rceil$ times the sequence $\tilde{X}$ of player 1. We call each sequence $\tilde{Z}$ that may be chosen by the follower an *action plan*.

During a block, assuming that the sequence of actions of the follower $\tilde{Z}$

10

matches $\lceil(1-q)n\rceil$ times the sequence $\tilde{X}$ of player 1, the prophet chooses a sequence of actions $\tilde{Y}$ such that:

- Among the $\lceil(1-q)n\rceil$ stages in which $\tilde{Z}$ and $\tilde{X}$ match, $\tilde{Y}$ matches $\tilde{X}$ exactly $\lceil(1-p)(1-q)n\rceil = \lceil xn\rceil$ times (and mismatches $\tilde{X}$ about $p(1-q)n = \frac{1-x}{3}n$ times).

- Among the $\lceil qn\rceil$ stages in which $\tilde{Z}$ and $\tilde{X}$ do not match, $\tilde{Y}$ matches $\tilde{X}$ exactly $\lceil\frac{q}{2}n\rceil = \lceil\frac{1-x}{3}n\rceil$ times (and mismatches $\tilde{X}$ about the same number of times).

Let $M(\tilde{X}, \tilde{Z})$ be the set of sequences $\tilde{Y}$ that satisfy these frequency requirements. We design strategies in such a way that, by the choice of a particular $\tilde{Y}$, the prophet indicates the follower which sequence of actions to play during the next block. Hence, we call each such sequence $\tilde{Y}$ a *message*.

In order for the above construction to work, we must find a set of action plans A such that there exists a one to one mapping from messages to action plans, and such that for each $\tilde{X}$, there exists $\tilde{Z} \in A$ that matches $\lceil(1-q)n\rceil$ times the sequence $\tilde{X}$. To check the existence of such a mapping, we need to estimate the size of the set of messages, and the minimal size of a set A having the required property.

A message is given by the choice of $\lceil(1-p)(1-q)n\rceil$ stages among $\lceil(1-q)n\rceil$, and of $\lceil\frac{q}{2}n\rceil$ among $\lfloor qn\rfloor$. Hence, the size of the set of messages is:

$$\binom{\lceil n(1-q)\rceil}{\lceil(1-p)(1-q)n\rceil}\binom{\lfloor qn\rfloor}{\lceil\frac{q}{2}n\rceil} \doteq 2^{(H(p)(1-q)+q)n}$$

Since $(1-p)(1-q) = x$ and $p(1-q) = \frac{q}{2} = \frac{1-x}{3}$, the following trees are equivalent:

11
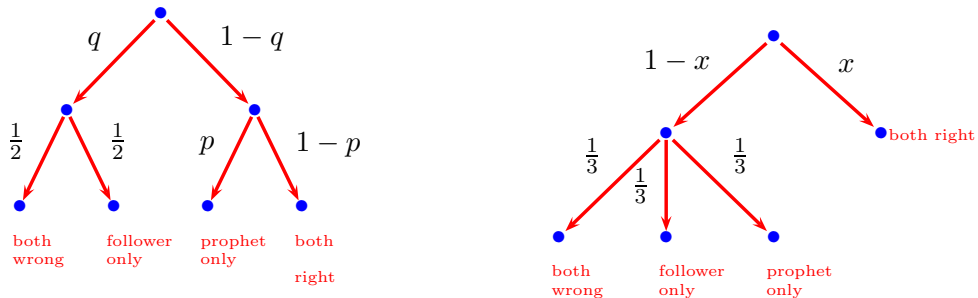
Hence, $H(q) + q + (1-q)H(p) = H(x) + (1-x) \log 3 = 1 - \eta$. Therefore,

$$H(p)(1-q) + q = 1 - H(q) + \eta$$

and the set of messages has size $\doteq 2^{n(1-H(q)+\eta)}$.

3.6.1 On the minimal size of a set of action plans

We prove that there exists a set of action plans $A = A(n)$ of size $|A(n)| \doteq 2^{n(1-H(q))}$, such that for every $\tilde{x} \in 2^n$, there exists $\tilde{z} \in A$ that matches $\tilde{X}$ exactly $\lceil n(1-q) \rceil$ times.

Lemma 1 *There exists a sequence of sets $(A(n))_n \in 2^n$ such that*

- *For every $\tilde{X} \in 2^n$, there exists $\tilde{Z} \in A$ that matches $\tilde{X}$ exactly $\lceil n(1-q) \rceil$ times,*

- $|A(n)| \doteq 2^{n(1-H(q))}$.

Proof. We prove this lemma following a probabilistic method (see for instance [2]). More precisely, we consider a random subset of 2^n composed of $|A(n)| \doteq 2^{n(1-H(q))}$ independently and uniformly distributed points, and prove that the probability that $A(n)$ satisfies the first condition of the lemma is positive. This will imply the existence of a realization $A(n)$ that fits it.

Let then $q'_n = \frac{\lfloor qn \rfloor}{n}$, and $\alpha_n = \lceil n\sqrt{2n} 2^{n(1-H(q'_n))} \rceil \doteq 2^{n(1-H(q))}$. Take a family $(Z_i)_{1 \le i \le \alpha_n}$ of i.i.d. uniformly drawn points in 2^n, and $A(n) = \{Z_i, 1 \le i \le \alpha_n\}$. Prob denotes the probability induced by the Z_i's.

For $X \in 2^n$, let $S(X; \lfloor qn \rfloor)$ be the sphere centered at X of radius $\lfloor qn \rfloor$ w.r.t. the Hamming distance. Lemma 1 implies:

$$|S(X; \lfloor qn \rfloor)| \ge \frac{2^{nH(q'_n)}}{\sqrt{2n}}$$

For $X \in 2^n$ and Z uniformly drawn in 2^n, we then have:

$$\text{Prob}\left(d_H(X, Z) \ne \lfloor qn \rfloor\right) \le \left(2^n - \frac{2^{nH(q'_n)}}{\sqrt{2n}}\right) \frac{1}{2^n} = 1 - \frac{2^{n(H(q'_n)-1)}}{\sqrt{2n}}$$

Then:

$$
\begin{aligned}
\operatorname{Prob}\left(\nexists Z \in A_n,\ d_H(X,Z) = \lfloor qn \rfloor\right) \ &\leq\ \left(1 - \frac{2^{n(H(q_n')-1)}}{\sqrt{2n}}\right)^{\alpha_n} \\
&\leq\ \left(\exp(-\frac{2^{n(H(q_n')-1)}}{\sqrt{2n}})\right)^{\alpha_n} \\
&\leq\ \exp(-n)
\end{aligned}
$$

Thus, for A_n randomly chosen, the expected number of points $X \in 2^n$ such that $\nexists Z \in A_n,\ d_H(X,Z) = \lfloor qn \rfloor$ is less than $2^n e^{-n} < 1$. Therefore, there exists a realization, i.e. a subset $A(n)$ of 2^n, such that this number of points X is zero, which means that for every $X \in 2^n$ there is $Z \in A(n)$ for which $d_H(X,Z) = \lfloor qn \rfloor$. ■

3.6.2 Construction of the optimal strategies

From the property of $A(n)$ and the fact that for n large enough $|M(\tilde{X}, \tilde{Z})| > A(n)$, we deduce the existence of families of message maps $(m_{\tilde{X},\tilde{Z}})_{\tilde{X},\tilde{Z} \in 2^n}$ and action maps $(a_{\tilde{X},\tilde{Z}})_{\tilde{X},\tilde{Z} \in 2^n}$ with $m_{\tilde{X},\tilde{Z}} \colon 2^n \to M(\tilde{X}, \tilde{Z})$, $a_{\tilde{X},\tilde{Z}} \colon M(\tilde{X}, \tilde{Z}) \to A(n)$ and such that

$$
\forall \tilde{X}' \tilde{X}' \in 2^n,\ \ d_H(a_{\tilde{X},\tilde{Z}}(m_{\tilde{X},\tilde{Z}}(\tilde{X}')), \tilde{X}') = \lfloor qn \rfloor \tag{1}
$$

We shall construct strategies for the team over blocks of length n. In these strategies, $m_{\tilde{X}_k,\tilde{Z}_k}$ is used by the prophet to choose a sequence of actions $\tilde{Y}_k$ in the k-th block as a function of the sequence $\tilde{X}_{k+1}$ of player 1 in the $k+1$-th block (knowing also the sequences $\tilde{X}_k$, $\tilde{Z}_k$ of players 1 and 3 in the k-th block). The follower then uses $a_{\tilde{X}_k,\tilde{Z}_k}$ to choose a sequence of actions in the $k+1$-th block as a function of $\tilde{Y}_k$. This is summarized by the following diagram:

$$
\begin{array}{ccccc}
2^n & \xrightarrow{m_{\tilde{X}_k,\tilde{Z}_k}} & M(\tilde{X}_k, \tilde{Z}_k) & \xrightarrow{a_{\tilde{X},\tilde{Z}}} & A(n) \\
\tilde{X}_{k+1} & \longrightarrow & \tilde{Y}_k & \longrightarrow & \tilde{Z}_{k+1}
\end{array}
$$

Property (1) then ensures that:

$$
d_H(\tilde{X}_{k+1}, \tilde{Z}_{k+1}) = \lfloor qn \rfloor
$$

We now define formally the strategies σ, τ for the prophet and the follower over blocks of length n. For $k \in \mathbb{N}$ let $\tilde{X}_k = (X_{kn+1}, \ldots, X_{(k+1)n})$, $\tilde{Y}_k =$

13

$(Y_{kn+1}, \ldots, Y_{(k+1)n})$ and $\tilde{Z}_k = (Z_{kn+1}, \ldots, Z_{(k+1)n})$ denote the actions of players 1, 2 and 3 during the k-th block, and let σ_k, τ_k represent the strategies of the prophet and the follower during the k-th block.

The first block: $k = 1$

During the first block, the prophet plays the actions of the sequence of the second block, while the follower plays a constant sequence of 1's:

$\tau_1(X, Y) = (1, \ldots, 1)$ and $\sigma_1(X) = \tilde{X}_2$.

The second block: $k = 2$

During the second block, the follower plays a sequence in $A(n)$ which is at a Hamming distance $\lfloor qn \rfloor$ of $\tilde{X}_2$, and the prophet tells the prophet what to play during block 3:

$\tau_2(X, Y) = \tilde{Z}_2$ such that $d_H(\tilde{Y}_1, \tilde{Z}_2) = \lfloor qn \rfloor$,

$\sigma_2(X) = m_{\tilde{X}_2, \tilde{Z}_2}(\tilde{X}_3)$.

Subsequent blocks: $k > 2$

In each subsequent block k, the follower interprets the previous message of the prophet in order to play a sequence of Hamming distance $\lfloor qn \rfloor$ of $\tilde{X}_k$, and the prophet signals the follower which sequence to play during the next block:

$\tau_k(X, Y) = a_{\tilde{X}_{k-1}, \tilde{Z}_{k-1}}(\tilde{Y}_{k-1})$,

$\sigma_k(X) = m_{\tilde{X}_k, \tilde{Z}_k}(\tilde{X}_{k+1})$.

4 General games

4.1 The main result

For μ a distribution on I, let $\mathcal{Q}(\mu)$ represent the set of distributions on $I \times J \times K$ with marginal μ on I, and set

$$u(\mu) = \max\{g(Q) \mid Q \in \mathcal{Q}(\mu) \text{ s.t. } H_Q(j, i \mid k) \geq H(\mu)\}.$$

Theorem 2 *The repeated game has value* $\min_\mu u(\mu)$ *in the following sense:*

1. For $\mu \in \Delta(I)$:
$$v_n(\mu^{\otimes n}) \leq u(\mu).$$

2.
$$\lim w_n = \min_\mu u(\mu)$$

14

4.2   Notations and tools

We introduce the tools and technical lemmas needed for the proof of theorem 2.

The type $\rho(a)$ of a finite sequence $a = (a_1, \ldots, a_n) \in A^n$, over a finite alphabet A is the empirical distribution of this sequence (i.e. $\rho(a)(b) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}_{a_i = b}$, for $b \in A$). Given $\mu \in \Delta(A)$, the type set of μ is $T^\mu(n) = \{a \in A^n, \rho(a) = \mu\}$. Finally, the set of types is $\mathbb{T}_n(A) = \{\mu \in \Delta(A), T_n(\mu) \neq \emptyset\}$. The following estimates the size of a type set $\mu \in \mathbb{T}_n(A)$ (see for instance theorem 12.1.3 page 282 of Cover and Thomas [11]):

$$\frac{2^{nH(\mu)}}{(n+1)^{|A|}} \leq |T^\mu(n)| \leq 2^{nH(\mu)} \tag{2}$$

We shall rely on this property of concavity of conditional entropies.

Lemma 2 *The function $Q \mapsto H_Q(\ell \mid \mathbf{i}))$ is concave on the (convex) set $\mathcal{Q}$ of probability measures on $I \times L$.*

Proof. Follows from the concavity of entropy. Indeed, let P be a distribution (with finite support) over $\mathcal{Q}$ and set $\bar{Q} = \sum_Q P(Q)Q$. Then

$$
\begin{aligned}
E_P(H_Q(\ell \mid \mathbf{i})) &= E_P\left(\sum_i Q(i)H(Q(\cdot \mid i))\right) \\
&\leq \sum_i \bar{Q}(i)h_{\bar{Q}}(\ell \mid i) \\
&= H_{\bar{Q}}(\ell \mid \mathbf{i})
\end{aligned}
$$

which completes the proof of the lemma. ■

Given a distribution Q on $I \times J \times K$, its marginal on e.g., $I \times J$ is denoted $Q_{I \times J}$. We define $T^Q(n)$ as the Q-strong-typical sequences of length n.

Given $x = (x_1, \ldots, x_n) \in I^n$, $y = (y_1, \ldots, y_n) \in J^n$, and $z = (z_1, \ldots, z_n) \in K^n$ we identify (x, y, z) (respectively, (x, z)) with the n-sequence whose i-th element is (x_i, y_i, z_i) (respectively (x_i, z_i)).

5   Proof of the main result

5.1   Play against distributions

Proof of part 1 of thm. 2.   Assume that $X = (X_1, \ldots, X_n)$ is a sequence of i.i.d. I-valued r.v.s where X_t has distribution μ on I. Let Y and Z be fixed 'pure strategies' of players 2 and 3.

The sequence $(X_1, Y_1, Z_1 \ldots, X_n, Y_n, Z_n)$ is a function of the sequence $(X_1, \ldots, X_n)$. Therefore, its entropy $H(X_1, Y_1, Z_1, \ldots, X_n, Y_n, Z_n)$ equals the entropy of $(X_1, \ldots, X_n)$ which equals $nH(\mu)$.

Denote by $\mathcal{F}_t$ the algebra of events spanned by the random variables $X_1, Y_1, \ldots, X_t, Y_t$. Denote by Q_t the conditional distribution of (X_t, Y_t, Z_t) given $\mathcal{F}_{t-1}$. Set

$$g_t = E_\mu \left(g(X_t, Z_t, Y_t) \mid \mathcal{F}_{t-1} \right).$$

Then, g_t is $\mathcal{F}_{t-1}$-measurable, and

$$E_{Q_t}(g(X_t, Y_t, Z_t)) = g_t.$$

Recall that Z is a pure strategy of the follower. Therefore, Z_t is a function of the past actions $(X_1, Y_1, \ldots, X_{t-1}, Y_{t-1})$ of players 1 and 2. In particular,

$$H(X_t, Y_t, Z_t \mid X_1, \ldots, Y_{t-1}, X_{t-1}) = H(X_t, Y_t \mid X_1, \ldots, Y_{t-1}, X_{t-1}) \quad (3)$$

Note also that by definition of Q_t:

$$H(X_t, Y_t, Z_t \mid \mathcal{F}_{t-1}) = E_\rho \left(h(X_t, Y_t, Z_t \mid \mathcal{F}_{t-1}) \right) = E_\rho H(Q_t) \quad (4)$$

The fact that Z_t is $\mathcal{F}_{t-1}$-measurable implies that the marginal distribution of Q_t on K is a Dirac measure. Then, $H_{Q_t}(k) = 0$ and:

$$H_{Q_t}(i, j \mid k) = H_{Q_t}(i, j) \quad (5)$$

Now, let $\bar{Q}_t = E_\rho Q_t$ and $Q = \frac{1}{n} \sum_{t=1}^n \bar{Q}_t$. We have:

$$
\begin{aligned}
nH(\mu) &= H(X_1, Y_1, \ldots, X_n, Y_n) \\
&= \sum_{t=1}^n H(X_t, Y_t \mid X_1, Y_1, \ldots, X_{t-1}, Y_{t-1}) \\
&= \sum_{t=1}^n E_\rho H(Q_t) \\
&= \sum_{t=1}^n E_\rho (H_{Q_t}(j, i \mid k)) \\
&\leq \sum_{t=1}^n H_{\bar{Q}_t}(j, i \mid k) \\
&\leq nH_Q(j, i \mid k)
\end{aligned}
$$

where the third equality follows from (4), next one from (5), and both inequalities from lemma 2. Therefore,

$$g(\rho, Y, Z) = E_Q(g(i, j, k))$$

where Q is a distribution on $I \times J \times K$ whose marginal on I equals μ and such that $H_Q(j, i \mid k) \geq H(\mu)$. Therefore $E_Q(g(i, j, k) \leq u(\mu)$, and the result follows. ∎

5.2 Play against all sequences

Proof of part 2 of theorem 2. We generalize the proof of theorem 1 to all games.

For every fixed $\varepsilon > 0$ we will specify a sufficiently large n. The play will be divided into consecutive blocks, each of length $f(n) + n$ where $f(n) = O(\log n)$. The I-type of a block is defined as the empirical distribution of the actions of player 1 in the last n stages of the block. The strategy Y will signal in the first $f(n)$ stages of each block I-type of the block. There are at most $n^{|I|}$ distinct I-types and therefore $O(\log n)$ stages suffice to specify the I-type.

The information transmitted by the prophet in the last n-stages of a block will be used by the follower in the next future block with identical I-type.

We will fix a finite family of distributions $\mathcal{Q}$ such that for every possible I-type (i.e., n-empirical distribution μ of a possible sequence $x \in I^n$) there is $Q = Q(\mu) \in \mathcal{Q}$ s.t. $Q_I = \mu$ and such that $H_Q(i, j \mid k) > H_Q(i)$ and $g(Q) > u(\mu) - \varepsilon$.

The strategy pair Y, Z of the prophet and the follower will ensure that excluding the first block of each n-empirical distribution type μ the joint n-empirical distribution equals $Q(\mu)$.

The details follow.

We now define a set of action plans $A_Q(n)$ for the follower such that for every sequence of actions $X \in T^{Q_I}(n)$ of player 1, there exists an action plan Z such that $(X, Z) \in T^{Q_{I \times K}}(n)$:

Lemma 3 *Let $Q_{I \times K} \in \Delta(I \times K)$. There exists a sequence of subsets $A_Q(n)$ of $T^{Q_K}(n)$ of size $\doteq 2^{(H(\mathbf{i}) - H(\mathbf{i}|\mathbf{k}))n} = 2^{I(\mathbf{i}:\mathbf{k})n}$ such that whenever $Q_{I \times K} \in \mathbb{T}_n(I \times K)$, for every $X \in T^{Q_I}(n)$ there exists $Z \in A_Q(n)$ such that $(X, Z) \in T^{Q_{I \times K}}(n)$.*

Proof. Assume $Q_{I \times K} \in \mathbb{T}_n(I \times K)$. We follow the same plan as for the proof of lemma 1 and consider a random set $\mathbf{A}_Q(n)$ of size $\alpha(n) = \lceil H(\mathbf{i})(n +$

$1)^{|I \times K|+1} 2^{nI(\mathbf{i};\mathbf{k})}\rceil$ that is the set of realizations of a family $\{\mathbf{Z}_k, 1 \le k \le \alpha(n)\}$ of i.i.d. random variables in K^n, where $\mathbf{Z}_k = (\mathbf{z}_{k,1}, \ldots, \mathbf{z}_{k,n})$ is i.i.d. according to Q_K. Let P be the induced probability over realizations of the $\mathbf{Z}_k$'s. An application of Sanov's Theorem (see for instance [11], thm. 12.4.1 p. 292) shows that for any $X \in T^Q(n)$ and $1 \le k \le \alpha(n)$:

$$P((X, \mathbf{Z}_k) \in T^{Q_{I \times K}}(n)) \ge \frac{1}{(n+1)^{|I \times K|}} 2^{nI(\mathbf{i};\mathbf{k})}$$

From this, we deduce that:

$$
\begin{aligned}
P(\nexists Z \in \mathbf{A}_Q(n), (X, Z) \in T^{Q_{I \times K}}(n)) &\le (1 - \frac{1}{(n+1)^{|I \times K|}} 2^{nI(\mathbf{i};\mathbf{k})})^{\alpha(n)} \\
&\le \exp(-nH(\mathbf{i}))
\end{aligned}
$$

Hence, the expected number of $X \in T^Q(n)$ such that $\nexists Z \in \mathbf{A}_Q(n), (X, Z) \in T^{Q_{I \times K}}(n)$ is at most $\exp(-nH(\mathbf{i}))|T^{Q_K}(n)| \le e^{-nH(\mathbf{i})} 2^{nH(\mathbf{i})} < 1$ by equation (2). Therefore, there exists a realization $A_Q(n)$ that verifies the condition. $\blacksquare$

Now, we consider the set of messages of the prophet. He communicates to the follower, sending a sequence $Y \in M(X, Z)$ such that $Y \in T^{Q_J}$ and for every $(x, z) \in T^{Q_{I \times K}}(n)$, $(x, y, z) \in T^Q(n)$.

The cardinality of $M(X, Z)$ is:

$$
\begin{aligned}
|M(X, Z)| &= \frac{|\{(x, y, z) \in T^Q(n)\}|}{|\{(x, z) \in T^{Q_{I \times K}}(n)\}|} \\
&\ge \frac{2^{H_Q}}{(n+1)^{|I \times J \times K|}} \frac{1}{2^{H_{Q_{I \times K}}}} \\
&\ge \frac{2^{H_Q(j|i,k)}}{(n+1)^{|I \times J \times K|}}
\end{aligned}
$$

Lemma 4 *For every $(X, Z) \in T^{Q_{I \times K}}(n)$ there is a subset $M(X, Z)$ of $T^{Q_J}(n)$ of size $\doteq 2^{H(\mathbf{j}|\mathbf{i},\mathbf{k})n}$ such that for every $Y \in M(X, Z)$ we have $(X, Y, Z) \in T^Q(n)$.*

As the cardinality of $M(X, Z)$ is greater than $2^{I(\mathbf{i};\mathbf{k})}$ we can construct mapping from $M(X, Z)$ to $A_Q(n)$.

Finally, if Q is a distribution on $I \times J \times K$ such that $H_Q(i, j \mid k) > H_Q(i)$, or equivalently, $H_Q(j \mid i, k) > H_Q(i) - H_Q(i \mid k)$, then we deduce that for every $(X, Z) \in T^{Q_{I \times K}}$ the message set $(M(X, Z))$ have more elements that

18

the target set of actions plans $(A_Q(n))$ and therefore we can implement (as in the online matching pennies) strategies Y, Z such that on all blocks (we leave $O(\ln n)$ between blocks to signal the empirical distribution of the sequence) where the empirical distribution of the sequence is Q_I the empirical distribution of the triples 'equals' (approximately) Q.

As we discussed, we now have to construct a finite family of distributions $\mathcal{Q}$ such that for every empirical distribution μ of a possible sequence $x \in I^n$ there is $Q = Q(\mu) \in \mathcal{Q}$ s.t. $Q_I = \mu$ and such that $H_Q(i, j \mid k) > H_Q(i)$ and $g(Q) > u(\mu) - \varepsilon$.

Using the strategy pair Y, Z that signals in the first $f(n) = O(\ln n)$ stages each block of size $n + f(n)$ stages the empirical distribution μ of the last n elements of the sequence in the block, and thereafter (excluding possible the first block of each distribution μ) 'implements' $Q(\mu)$, we guarantee that in the $k(f(n) + n)$ stages the payoff to the team will be at least

$$\min_{\mu} u_n(\mu) - \varepsilon - \|g\|(\ln n/n - \|\mathcal{Q}\|/k)$$

This completes the proof. ∎

6 Applications to repeated games played by boundedly rational players

The corollaries of the present section address repeated games played by finite automata or by players with bounded recall. We follow the notation of the previous sections. In particular the three player stage game G has actions sets I, J, and K, and payoff function g to the team of players 2 (the prophet) and 3 (the follower).

It should be pointed out that for any ε-approximate optimal strategies of the prophet and the follower, the complexity of the strategies of the follower is bounded. E.g., it is implementable by a strategy of bounded recall and thus also by a finite automaton.

In what follows we wish to deduce the foresight of the prophet from his computational superiority. For that we bound the complexity of the sequence and provide sufficient lower bounds for the size of automata or the length of recall needed to generate the foresight and to implement the strategy.

We assume that the sequence (of player 1) is an m_1-periodic sequence. This will be the case if it is generated by a non-interactive automaton with m_1 states.

In order for player 2 to be able to record and keep track of the sequence and the stage within the sequence he need an automaton of size $I^{m_1} m_1$. (The first term $|I^{m_1}|$ is to enable recording the periodic cycle, and the second factor containing m_1 elements enables to keep track of the time within the cycle.) Therefore we have,

Corollary 1 *For every $\varepsilon > 0$ there is m sufficiently large s.t. for every m_1 and m_2 with $m_2 > |I^{m_1}| m_1$ there are pure strategies σ (of player 2) implementable by an automaton of size m_2 and a strategy τ (of player 3) implementable by an automaton of size m s.t. for every infinite m_1-periodic sequence $X = (X_1, \ldots)$ we have*

$$\sum_{t=1}^{n} g(X_t, Y_t, Z_t) \geq v^* n - \varepsilon n - m$$

where $Y_t = \sigma(X_1, \ldots, X_{t-1})$, and $Z_t = \tau(Y_1, \ldots, Y_{t-1})$.

A minor modification of the strategy of the prophet is needed when we restrict ourselves to bounded recall strategies. This modification call on player 2 to mark the start and end of the cycle, by strings of his own actions that will not appear elsewhere. We skip the details. We thus have,

Corollary 2 *For every $\varepsilon > 0$ there is m sufficiently large s.t. for every m_1 and m_2 with $m_2 > m_1$ there are pure strategies σ (of player 2) of recall of size m_2 and a strategy τ (of player 3) of recall size m s.t. for every infinite m_1-periodic sequence $X = (X_1, \ldots)$ we have*

$$\sum_{t=1}^{n} g(X_t, Y_t, Z_t) \geq v^* n - \varepsilon n - m$$

where $Y_t = \sigma(X_{t-m_2}, Y_{t-m_2}, \ldots, X_{t-1}, Y_{t-1})$, and $Z_t = \tau(Y_{t-m_2}, \ldots, Y_{t-1})$.

References

[1] Abreu, D. and Rubinstein, A. (1988) The structure of Nash equilibrium in repeated games with finite automata, *Econometrica* **56**, 1259–1282.

[2] N. Alon and J. Spencer. The Probabilistic Method. John Wiley & Sons, New York, 1991.

[3] Anderlini, L. (1989) Some notes on Church's Thesis and common interest games, *Theory and Decisions* **29**, 19–52.

[4] Anderlini, L. and Sabourian, H. (1995) Cooperation and effective computability, *Econometrica* **63**, 1337–1369.

[5] Aumann, R.J. (1981) Survey of repeated games, in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern* 11-42, Mannheim, Wien, Zurich: Wissenschaftsverlag, Bibliographisches Institut.

[6] Aumann, R.J. and Sorin, S. (1989) Cooperation and bounded recall, *Games and Economic Behavior* **1**, 5–39.

[7] E. Ben-Porath. Repeated games with finite automata, *Journal of Economic Theory* **59**, 17-32, 1993.

[8] Binmore, K.G. (1987) Modeling Rational Players, I, *Econ. Phil.* **3**, 179–214.

[9] Binmore, K.G. (1988) Modeling Rational Players, II, *Econ. Phil.* **4**, 9–55.

[10] Cho, I.-K (1994) Bounded rationality, neural networks and folk theorem in repeated games with discounting, *Economic Theory* **4** , 935–957.

[11] T. M. Cover and J. A. Thomas. *Elements of information theory.* Wiley Series in Telecommunications. Wiley, 1991.

[12] Gilboa, I. and Samet, D. (1989) Bounded Versus Unbounded Rationality: The Tyranny of the Weak, *Games and Economic Behavior* **1, 3**, 213–221.

[13] O. Gossner. Sharing a long secret in a few public words, *DP* 2000-15, THEMA, 2000.

[14] Gossner, O. and Hernández, P. (2001) On the complexity of coordination, Core DP 01/47. *to appear in Mathematics of Operations Research.*

[15] Gossner, O. and Vieille, N. (1999) How to play with a biased coin, THEMA DP 99-31. To appear in *Games and Economic Behavior.*

[16] Gossner, O. and Tomala, T. (2001) Entropy in repeated games with imperfect monitoring, mimeo.

[17] Hernández, P. and Urbano, A. (2001) Pseudorandom processes: Entropy and Automata, *WP-AD* 2001-22, IVIE.

[18] E. Kalai. Bounded rationality and strategic complexity in repeated games, T. Ichiishi, A. Neyman, Y. Tauman, eds., *Game Theory and Applications*, Academic Press, New York, 131-157.

[19] E. Kalai and W. Standford. Finite rationality and interpersonal complexity in repeated games, *Econometrica* **56, 2**, 397-410, 1988.

[20] E. Lehrer. Repeated games with stationary bounded recall strategies, *Journal of Economic Theory* **46**, 130-144, 1988.

[21] E. Lehrer. Finitely many players with bounded recall in infinitely repeated games, *Games and Economic Behavior* **7**, 390-405, 1994.

[22] Mac Williams F. J. and Sloane N.J.A., (1977) *The Theory of Error Correcting Codes*, North-Holland.

[23] N. Meggido and R. Widgerson. On computable beliefs of rational machines, *Games and Economics Behavior* **1**, 144-169, 1989.

[24] A. Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoner's dilemma, *Economics Letters* **19**, 227-229, 1985.

[25] Neyman, A. (1997) Cooperation, Repetition, and Automata, in *Cooperation: Game-Theoretic Approaches*, edited by S. Hart and A. Mas Colell. NATO ASI Series F, **155**. Springer-Verlag, 233–255.

[26] A. Neyman. Finitely repeated games with finite automata. *Mathematics of Operations Research* **23**, 513-552, 1998.

[27] Neyman, A. and Okada, D. (1999) Strategic Entropy and Complexity in Repeated Games, *Games and Economic Behavior* **29**, 191–223.

[28] Neyman, A. and Okada, D. (2000) Repeated Games with Bounded Entropy, *Games and Economic Behavior* **30**, 228–247.

[29] Neyman, A. and Okada, D. (2000) Two-Person Repeated Games with Finite Automata, *International Journal of Game Theory* **29**, 309–325.

[30] C. Papadimitriou and M. Yannakakis. On complexity as bounded rationality. *Proceeding of the 26th ACM Symposium on Theory of Computation*, 726-733.

[31] Papadimitriou, C. H. and Yannakakis M. (1998) On bounded rationality and computation complexity, *mimeo*, 1998.

[32] Piccione, M. (1989) Finite automata equilibria with discounting and unessential modifications of the stage game, *Journal of Economic Theory* **56**, 180–193.

[33] Piccione, M. and Rubinstein, A. (2002) Modeling economic interaction of agents with diverse abilities to recognize equilibrium patterns, *mimeo*.

[34] Rubinstein, A. (1986) Finite Automata play the repeated prisoners dilemma, *Journal of Economic Theory* **39**, 83–96.

[35] Simon, H. (1955) A behavioral model of rational choice, *Q. J. Econ.* **64**, 99–118.

[36] Zemel, E. (1989) Small talk and cooperation: A note on bounded rationality, *Journal of Economic Theory* **49**, 1–9.