האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

ONLINE MATCHING PENNIES

by

OLIVIER GOSSNER, PENÉLOPE HERNÁNDEZ and ABRAHAM NEYMAN

Discussion Paper # 316 May 2003

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

Online matching pennies

Olivier Gossner* Penélope Hernández† Abraham Neyman‡

May 6, 2003

Abstract

We study a repeated game in which one player, the prophet, acquires more information than another player, the follower, about the play that is going to be played. We characterize the optimal amount of information that can be transmitted online by the prophet to the follower, and provide applications to repeated games played by finite automata, and by players with bounded recall.

1 Introduction

The classical paradigm of game theory assumes the full rationality of interactive agents. In particular, it often assumes unlimited computational power. However, there are many games where it is impossible to assume that the players can implement all strategies. In fact, the number of strategies of the repeated game grows at a double exponential rate in the number of repetitions, and many of the strategies are not implementable by reasonable-sized computing agents.

A central question that arises is how does the outcome of a given interaction depend on the agents' computational power. Over the last two decades many papers have developed a theory that enables one to quantify the impact on the outcome of such limits on implementable strategies. The general framework of this theory considers repeated games where each player is restricted to a subclass of strategies that are implementable by

*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.
This research was supported in part by The Israel Science Foundation grant 382/98 and by the Zvi Hermann Shapira Research Fund.

machines with prescribed bounds on its complexity. For instance, Ben Porath [2], Kalai [5], Kalai and Stanford [6], Megiddo and Widgerson [10], Neyman [11] [13], Neyman and Okada [14], Papadimitriou and Yannakakis [15] study repeated games with bounded automata, Lehrer [7], [8] study repeated games with bounded recall, and Gossner [4] studies repeated games played by polynomial time Turing machines.

The study of equilibria of the resulting strategic dynamic interactions leads to phenomena that are absent in the (unconstrained) repeated game. For instance, consider a 3-player repeated game in which players 2 and 3 form a team, player 1 is restricted to choosing a periodic sequence of period m and players 2 and 3 are restricted to automata of sizes m_2 and m_3 respectively. Results due to Ben Porath [2] and Neyman [12] show that, whenever $m_2 \geq \exp(Km)$ for K sufficiently large, and $m_3 = o(\log m)$, complexity bounds imply the sequence can be fully decoded by player 2, whereas it remains completely unpredictable for player 3. Hence, it may be profitable for the team that player 2 shares part of its knowledge of the sequence with player 3. The only way that player 2 can send information to player 3 is by the particular choice of action during the course of the game. This is what we refer to as *online communication* through actions, opposed to *offline communication* through some outside communication channel. We investigate the best means for the team to achieve this communication, which leads to the study of online communication by boundedly rational agents.

In order to single out the problem of efficient online communication, we consider an auxiliary 3-player game, called the 3-player online matching pennies. Its analysis is of independent interest and it enables to derive results regarding repeated games with bounded automata and/or bounded recall.

The 3-player matching pennies is a zero-sum game in which each one of the 3 players chooses an action in $\{0, 1\}$, and the payoff to players 2 and 3 is 1 if the 3 actions coincide and 0 otherwise. In the n-stage version of online matching pennies, player 1 first chooses a sequence of actions in $\{0, 1\}^n$, which is then announced to player 2. The repeated game then proceeds, in which player 1 plays the chosen sequence, player 2 plays according to its knowledge of the sequence and its observation of player 3's actions, and player 3 plays only conditional to past actions. Clearly, the problem faced by the team in the game of online matching pennies is the same as in the previously discussed game played by finite automata, namely to efficiently use online communication from player 2 to player 3.

The performance of a strategy for the team can be measured under different assumptions on how the sequence is chosen. In particular, the sequence may be random i.i.d. $(\frac{1}{2}, \frac{1}{2})$, or it may be the worst possible against this given strategy. In both cases, we study the optimization problem of the team. Our main result shows, somewhat surprisingly, that the value, i.e., the maximal payoff that the team can guarantee, in both cases coincide. Namely, we prove the existence of a value v^* such that there exist pure strategies of the team that guarantee v^* in the long run against all sequences, and furthermore, no strategy of the team can obtain more than v^* against an i.i.d. $(\frac{1}{2}, \frac{1}{2})$ sequence. We also give an analytical formula for v^*, and design ε-optimal strategies.

We introduce formally the model of online matching pennies in Section 2, and illustrate it with examples of strategies in Section 3. Section 4 presents the main result, provides the formula for v^*, and introduces tools for the proofs. In Section 5, we bound the payoff that the team can obtain against a random sequence by v^*. Moreover, we construct ε-optimal strategies that guarantee $v^* - \varepsilon$ against all sequences in long games in Section 6. Finally, we conclude in Section 7 with applications to repeated games played by finite automata, or by players with bounded recall.

2 Model: Online matching pennies

2.1 The one-shot game

We consider 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. The payoff to player 1 is $-g(i,j,k)$.

Player 2 is called the prophet, and player 3 the follower. Since they have a common payoff function, players 2 and 3 form a team.

2.2 The repeated game

A (pure) strategy for the prophet is a mapping $Y\colon \{0,1\}^{\mathbb{N}} \to \{0,1\}^{\mathbb{N}}$ with coordinates $(Y_n)_n$. A (pure) strategy Z for the follower is a mapping $Z\colon \{0,1\}^{\mathbb{N}} \times \{0,1\}^{\mathbb{N}} \to \{0,1\}^{\mathbb{N}}$ with coordinates $(Z_n)_n$ such that Z_n depends on the actions of players 1 and 2 from stages 1 to $n-1$.

Hence, player 1 chooses a sequence of actions, and does not react to the team's actions. Player 2 has knowledge in advance of the actions of player 1, and can play accordingly. Player 3's actions depend on the past actions of players 1 and 2.

Given a sequence $X \in \{0,1\}^{\mathbb{N}}$ 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 $\{0,1\}^{\mathbb{N}}$ together with strategies (Y,Z) induces a probability distribution $P_{Y,Z}$ on the set of sequences (x_t, y_t, z_t) in $(\{0,1\} \times \{0,1\} \times \{0,1\})^{\mathbb{N}}$.

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 $\{0,1\}^n$, 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 $X \sim \rho$, $v_n(\rho)$ is the best expected payoff the team can obtain in the n-stage game, and v_n is the best expected payoff against the worst distribution ρ.

The minimal payoff that strategies Y, Z of the team guarantee against all sequences is $\min_{(x_t)_t} \frac{1}{n} \sum_{t=1}^{n} g(x_t, y_t, z_t)$ 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, $w_n \leq v_n(\rho)$ for any ρ and therefore $w_n \leq v_n$.

3 Candidate strategies

We present two strategies for the team and analyze their payoffs both against the distribution $\rho = (1/2, 1/2)^{\otimes n}$, 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 prophet. The follower plays an arbitrary sequence of actions on the odd stages.

The resulting sequences of actions are:

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

- 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 of 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 Example 1 do not involve any communication. Strategies of Example 2, which involves communication at odd stages, improve both in terms of expected payoff against ρ and against the worst case.

4 Main result

We let v^* be the (unique) solution of the equation

$$H(x) + (1-x)\log_2 3 = 1$$

where H is the entropy function $H(x) = -x\log_2 x - (1-x)\log_2(1-x)$ for $x \in (0,1)$.

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

1. *With $P = (\frac{1}{2}, \frac{1}{2})^{\otimes \mathbb{N}}$, for all strategies Y, Z of the team and for all $n \in \mathbb{N}$,*
$$\frac{1}{n}\sum_{t=1}^{n} \mathbf{E}_{P_{Y,Z}} g(x_t, y_t, z_t) \leq v^*.$$
2. *There exist strategies Y, Z of the team such that: for all $\varepsilon > 0$ there exists N such that for all $n \geq N$ and for all P,*
$$\frac{1}{n}\sum_{t=1}^{n} \mathbf{E}_{P_{Y,Z}} g(x_t, y_t, z_t) \geq v^* - \varepsilon.$$

4.1 Notations and tools

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

4.1.1 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 ($z - 1 < [z] \leq z$ and $z \leq \lceil z \rceil < z + 1$). 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 $\lim_{n\to\infty} \frac{\log a_n - \log b_n}{n} = 0$.

4.1.2 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

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

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

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

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:

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

4.1.3 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 $d_H(x,y)$ $(= \sum_{t=1}^{n} \mathbb{I}(x_t \neq y_t))$.

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. [9], p. 309, 310.) ■

5 The information constraint

We assume the distribution ρ of sequences of player 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 distribution over the sequences of length n.

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 $X = (X_1, \dots, X_n)$ 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 players 2 and 3 respectively.

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

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

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:

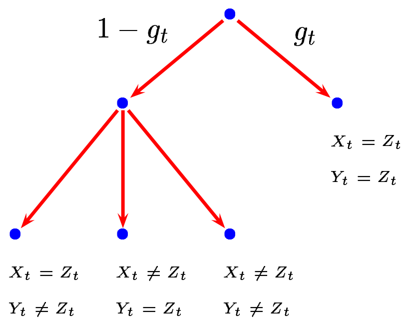

Hence, we deduce that:

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

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

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

Summing over t now gives:

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

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,

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

We now apply twice Jensen's inequality to the concave mapping $x \mapsto H(x) + (1 - x) \log 3$ and obtain

$$\begin{aligned} h &\le \mathbf{E}_\rho (H(\frac{1}{n} \sum_{t=1}^{n} g_t) + (1 - \frac{1}{n} \sum_{t=1}^{n} g_t) \log 3) \\ &\le H(\mathbf{E}_\rho \frac{1}{n} \sum_{t=1}^{n} g_t) + (1 - \mathbf{E}_\rho \frac{1}{n} \sum_{t=1}^{n} g_t) \log 3. \end{aligned}$$

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

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

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 $h = 1$ $(0 \le h \le 1)$, and equals v^*. From this we deduce Part 1 of Theorem 1.

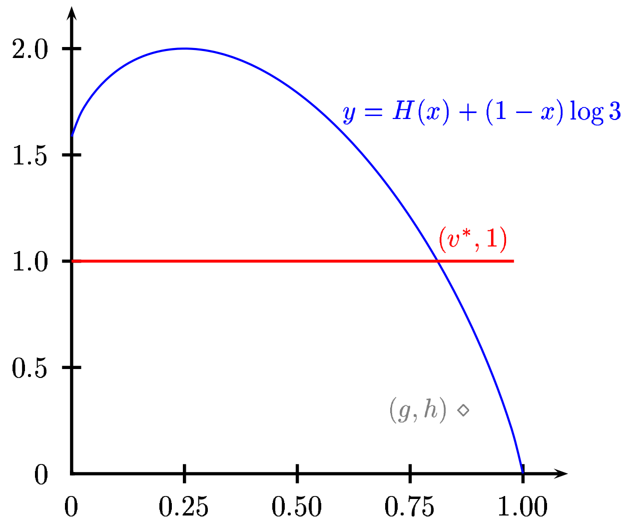

■

6 Design of ε-optimal strategies

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 $\tilde{Z}$ 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 to 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 mapping from the set of messages onto the set of 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:

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

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

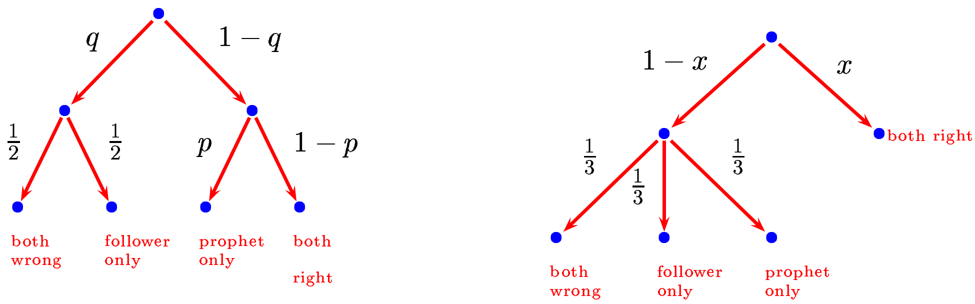

Hence, $H(q) + q + (1-q)H(p) = H(x) + (1-x)\log 3 = 1 - \eta$. Therefore,

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

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

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 [1]). 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:

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

Then:

$$\begin{aligned} \text{Prob}\,(\nexists Z \in A(n),\ d_H(X,Z) = \lfloor qn \rfloor) &\le \left(1 - \frac{2^{n(H(q'_n)-1)}}{\sqrt{2n}}\right)^{\alpha_n} \\ &\le \left(\exp(-\frac{2^{n(H(q'_n)-1)}}{\sqrt{2n}})\right)^{\alpha_n} \\ &\le \exp(-n). \end{aligned}$$

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 $d_H(X,Z) = \lfloor qn \rfloor$. ■

Remark that the set of messages of the prophet depends on the actions $\tilde{X}$ and $\tilde{Z}$ of the sequence and the follower. This is so because the set of messages is defined in terms of statistics for $(\tilde{X}, \tilde{Y}, \tilde{Z})$. On the other hand, the set A of action plans for the follower does not depend on $\tilde{X}$ and $\tilde{Y}$.

6.2 Construction of the optimal strategies

From the property of A and the fact that for n large enough $|M(\tilde{X}, \tilde{Z})| > A$, 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 $m_{\tilde{X},\tilde{Z}}\colon 2^n \to M(\tilde{X},\tilde{Z})$, $a_{\tilde{X},\tilde{Z}}\colon M(\tilde{X},\tilde{Z}) \to A$ and such that

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

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:

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

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 $\tilde{X}_k = (X_{(k-1)n+1}, \dots, X_{kn})$, $\tilde{Y}_k = (Y_{(k-1)n+1}, \dots, Y_{kn})$ and $\tilde{Z}_k = (Z_{(k-1)n+1}, \dots, Z_{kn})$ 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.

- **During the first block** ($k=1$), the prophet plays the actions of the sequence of the second block, while the follower plays a constant sequence of 1's. Formally:

$$\begin{cases} \sigma_1(X) & = \tilde{X}_2 \\ \tau_1(X,Y) & = (1,\dots,1). \end{cases}$$

- **During the second block** ($k=2$), the follower plays a sequence in A 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. Formally:

$$\begin{cases} \tau_2(X,Y) & = \tilde{Z}_2 \text{ such that } d_H(\tilde{Y}_1,\tilde{Z}_2) = \lfloor qn \rfloor, \\ \sigma_2(X) & = m_{\tilde{X}_2,\tilde{Z}_2}(\tilde{X}_3)\,. \end{cases}$$

- **In each subsequent block** ($k>2$), 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 to the follower which sequence to play during the next block:

$$\begin{cases} \tau_k(X,Y) &= a_{\tilde{X}_{k-1},\tilde{Z}_{k-1}}(\tilde{Y}_{k-1}), \\ \sigma_k(X) &= m_{\tilde{X}_k,\tilde{Z}_k}(\tilde{X}_{k+1}). \end{cases}$$

7 Applications to repeated games played by boundedly rational players

The corollaries of the present section address 3-players repeated matching pennies played by finite automata or by players with bounded recall. We follow the notation of the previous sections.

Notice that in the previous construction of ε-optimal strategies for the team, the complexity of the strategies of the follower is bounded. Indeed, it is implementable by a strategy of bounded recall (of memory $2n$, where n is the size of a block) 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 m_1-periodic. This will be the case if it is generated by a non-interactive automaton with m_1 states, or by an oblivious bounded recall strategy of memory $\log_2 m_1$.

In order for player 2 to be able to record and keep track of the sequence and the stage within the sequence he needs an automaton of size $2^{m_1}m_1$. (2^{m_1} states suffice to record the periodic cycle, and m_1 states suffice 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 $m_2 > |2^{m_1}|m_1$ 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 $X = (X_1, \dots)$ we have*

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

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

A minor modification of the strategy of the prophet is needed when we restrict ourselves to bounded recall strategies. This modification calls 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 of size m s.t. for every infinite m_1-periodic sequence $X = (X_1, \dots)$ we have*

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

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

References

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

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

[3] T. M. Cover and J. A. Thomas. *Elements of Information Theory.* Wiley Series in Telecomunications. John Wiley & Sons, New York 1991.

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

[5] 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, 1990.

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

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

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

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

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

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

[12] A. Neyman. Cooperation, repetition, and automata, S. Hart, A. Mas Colell, eds., *Cooperation: Game-Theoretic Approaches*, NATO ASI Series F, 155. Springer-Verlag, Berlin, 233-255, 1997.

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

[14] A. Neyman and D. Okada. Repeated games with bounded entropy, *Games and Economic Behavior* **30**, 228-247, 2000.

[15] C. Papadimitriou and M. Yannakakis. On complexity as bounded rationality. *Proceeding of the 26th ACM Symposium on Theory of Computation*, ACM Press, New York, 726-733, 1994.