Games and Economic Behavior 69 (2010) 165–168

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Complexity and effective prediction ☆

Abraham Neyman [a,*], Joel Spencer [b,c]

[a] Institute of Mathematics and Center for the Study of Rationality, The Hebrew University of Jerusalem, Givat Ram, Jerusalem 91904, Israel
[b] Department of Computer Science, New York University, USA
[c] Department of Mathematics, New York University, USA

ARTICLE INFO

Article history:
Received 10 December 2007
Available online 27 May 2009

JEL classification:
C73
D83
C44

ABSTRACT

Let $G = \langle I, J, g\rangle$ be a two-person zero-sum game. We examine the two-person zero-sum repeated game $G(k,m)$ in which players 1 and 2 place down finite state automata with k, m states respectively and the payoff is the average per-stage payoff when the two automata face off.

We are interested in the cases in which player 1 is "smart" in the sense that k is large but player 2 is "much smarter" in the sense that $m \gg k$. Let $S(g)$ be the value of G where the second player is clairvoyant, i.e., would know player 1's move in advance.

The threshold for clairvoyance is shown to occur for m near $\min(|I|, |J|)^k$. For m of roughly that size, in the exponential scale, the value is close to $S(g)$. For m significantly smaller (for some stage payoffs g) the value does not approach $S(g)$.

1. Introduction

Let $G = \langle I, J, g\rangle$ be a two-person zero-sum game; I and J are the finite action sets of player 1 and player 2 respectively, and $g : I \times J \to \mathbb{R}$ is the payoff function to player 1. The repeated game where player 1's, respectively player 2's, possible strategies are those implementable by automata of size k, respectively size m, and the payoff is the average per-stage payoff, is denoted $G(k,m)$. Ben-Porath (1993) proves that the value of $G(k,m)$ converges to the value of the stage game G as k goes to infinity and $\frac{\log m}{k} + \frac{\log k}{m}$ goes to 0.

It follows that in order to have an asymptotic nonvanishing advantage in the repeated game with finite state automata an exponentially larger automata size is needed. Neyman (1997) proves that if $\liminf_{k\to\infty} \frac{\log m_k}{k} > \log |J|$ then the value of $G(k, m_k)$ converges, as k goes to infinity, to the max min of the stage game where player 1 maximizes over his pure stage actions $i \in I$ and player 2 minimizes over his pure stage actions $j \in J$. Applying this result to the special case where for some function $r : I \to J$ the stage-payoff function is $g(i, j) = -1$ if $j = r(i)$ and $g(i, j) = 0$ if $j \neq r(i)$, we obtain the following: if $\liminf_{k\to\infty} \frac{\log m_k}{k} > \log |r(I)|$, then, for sufficiently large k, player 2 has a strategy (in $G(k, m_k)$) such that for every strategy of player 1 the expected empirical distribution of the action pairs (i, j) is essentially supported on the set of action pairs of the form $(i, r(i))$.

The main result of the present paper is a complete characterization of the asymptotic relation between k and m_k such that player 2 can effectively predict the moves of player 1 in $G(k, m_k)$, namely, no matter what the stage-game payoff function g, the values of the games $G(k, m_k)$ converge to the max min $(\max_{i\in I} \min_{j\in J} g(i, j))$. This asymptotic relation is $\liminf_{k\to\infty} \frac{\log m_k}{k} \geqslant \min(\log |I|, \log |J|)$.

☆ This research was supported in part by Israel Science Foundation grants 263/03 and 1123/06. The authors enjoyed the hospitality, generosity and stimulating atmosphere of Microsoft Research (Redmond) in the initial phases of this research.

* Corresponding author.
E-mail addresses: aneyman@math.huji.ac.il (A. Neyman), joelspencer@cims.nyu.edu (J. Spencer).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2009.05.007

The “matching pennies” game provides a good example. Here $I = J = \{0, 1\}$ and $g(x, y) = -1$ when $x = y$ and $g(x, y) = +1$ when $x \neq y$. From (Neyman, 1997) when $m > (2.001)^k$ then $G(k, m) \to -1$. The second player is so much “smarter” than the first player that the second player can effectively predict the first player’s move. Our new result says essentially that there is not a phase transition at 2^k. When $m \sim (1.999)^k$ then the value of $G(k, m)$ is strictly more than but quite close to -1. We note that our proof for general games G was first done for the matching pennies game and that, indeed, the basic ideas of the general proof can be derived from this particular example.

An open problem (see Neyman, 1997) is the quantification of the feasible “level of prediction” when the limit of $\frac{\log m_k}{k}$ equals θ and $0 < \theta < \min(\log|I|, \log|J|)$. It is conjectured in (Neyman, 1997) that for every, but at most one, value of $\theta > 0$, if $\frac{\log m_k}{k} \to_{k\to\infty} \theta$ then the value of $G(k, m_k)$ converges as k goes to infinity. More explicitly, there is a nonincreasing function $v : (0, \infty) \to \mathbb{R}$ and $\theta_0 > 0$ such that v is continuous at all $\theta \neq \theta_0$ and such that if $\frac{\log m_k}{k} \to_{k\to\infty} \theta > 0$ and $\theta \neq \theta_0$ then the value of $G(k, m_k)$ converges to $v(\theta)$ as k goes to infinity.

Our result shows (in particular) that such a discontinuity cannot happen at $\theta_0 = \min(\log|I|, \log|J|)$. More explicitly, we prove that for every $\varepsilon > 0$ there is $\delta > 0$ such that if k is sufficiently large and $\log m > k(\log|J| - \delta)$ then the value of $G(k, m)$ is $\leqslant \max_{i\in I} \min_{j\in J} g(i, j) + \varepsilon$.

2. Preliminaries

Fix a two-person zero-sum stage game $G = \langle I, J, g\rangle$. A finite history of the repeated game G^* is an element of $(I \times J)^*$, i.e., all finite strings $(i_1, j_1, \ldots, i_t, j_t)$ (including the empty string $\emptyset$). A pure strategy σ of player 1 is a function $\sigma : (I \times J)^* \to I$ and a pure strategy τ of player 2 is a function $\sigma : (I \times J)^* \to J$. A pair of pure strategies (σ, τ) induces a play $(i_1, j_1, \ldots)$, defined inductively as follows: $i_1 = \sigma(\emptyset)$, $j_1 = \tau(\emptyset)$, $i_{t+1} = \sigma(i_1, j_1, \ldots, i_t, j_t)$, and $j_{t+1} = \tau(i_1, j_1, \ldots, i_t, j_t)$. The average payoff per stage is defined as $\lim_{T\to\infty} \sum_{t=1}^{T} g(i_t, j_t)$ whenever the limit exists, and if the play is induced by the strategy pair (σ, τ) we denote this average per-stage payoff by $g(\sigma, \tau)$.

An *automaton* of player 2 consists of a set of *states* M, an *action function* $\alpha : M \to J$, a *transition function* $\beta : M \times I \to M$, and an *initial state* $m^* \in M$. The *size* of an automaton $A = \langle M, m^*, \alpha, \beta\rangle$ is the number $|M|$ of states.

An *automaton* $A = \langle M, m^*, \alpha, \beta\rangle$ for player 2 defines a strategy $\tau = \tau^A$ as follows. Define the sequence of states $(m_t)_{t\geqslant 1}$ by $m_1 = m^*$ and $m_{t+1} = \beta(m_t, i_t)$. Note that m_t is a function of $i_1, j_1, \ldots, i_{t-1}, j_{t-1}$. Define

$$\tau\big(s_1 = (i_1, j_1), \ldots, s_{t-1} = (i_{t-1}, j_{t-1})\big) = \alpha(m_t).$$

Analogously one defines an automaton for player 1.

The set of all automata of size m of a player, as well as the set of all his/her strategies implementable by automata of size m, are denoted $\mathcal{A}(m)$.

We denote by $[k]$ the set $\{1, \ldots, k\}$ if k is a positive integer and $[\eta]$ also denotes the integer part of η (the largest integer $\leqslant \eta$). No confusion should result.

3. The main result

The deterministic play induced by a pure strategy $\sigma \in \mathcal{A}(k)$ of player 1 and a pure strategy $\tau \in \mathcal{A}(m)$ of player 2 enter a cycle (of length $\leqslant km$) and therefore the average payoff per stage is well defined and is denoted $g(\sigma, \tau)$. A mixed strategy $\sigma \in \Delta(\mathcal{A}(k))$ of player 1 and a mixed strategy $\tau \in \Delta(\mathcal{A}(m))$ induce a random play, which is a mixture of periodic plays, and therefore the expected average payoff per stage is well defined and denoted $g(\sigma, \tau)$.

Theorem 1. *Fix a two-person zero-sum stage game $G = \langle I, J, g\rangle$ and set $v_* = \max_{i\in I} \min_{j\in J} g(i, j)$. Then, $\forall\varepsilon > 0\ \exists\delta > 0$ and k_0 such that if $k > k_0$ and $\log m > k(\min(\log|I|, \log|J|) - \delta)$ then $v_* \leqslant \operatorname{Val} G(k, m) < v_* + \varepsilon$.*

Proof. Obviously, $\operatorname{Val} G(k, m) \geqslant v_*$. In order to prove the other inequality we assume without loss of generality that $|J| \leqslant |I|$. Fix $\varepsilon > 0$. Let $K > 0$ be a sufficiently large constant so that $2\|g\|/K < \varepsilon$ where $\|g\| = \max_{i,j} |g(i, j)|$. For $0 < x < 1$ we denote by $H(x)$ the entropy of the probability vector $(x, 1 - x)$, namely, $H(x) = -x\log_2 x - (1 - x)\log_2(1 - x)$. The following properties of the entropy function are used in the sequel: $H(x)/x \to \infty$ as $x \to 0+$, H is strictly increasing on $(0, 1/2)$, and if $m_k/k \to_{k\to\infty} x \in (0, 1)$ then $\frac{1}{k}\log\binom{k}{m_k} \to_{k\to\infty} H(x)$.

As $H(\frac{\delta}{3K})/\delta \to_{\delta\to 0+} \infty$ we can choose $\delta > 0$ sufficiently small so that $H(\frac{\delta}{3K}) > 3\delta + \delta\log|J|$.

Let $n = [(1 + \delta)k]$ and $\bar{\ell} = [2^{k(\log|J|-\delta)}]$. For every list $(j_t^\ell)_{t,\ell}, (t_1^\ell, t_2^\ell)_\ell$, where $j_t^\ell \in J$, $1 \leqslant t \leqslant n$, $1 \leqslant \ell \leqslant \bar{\ell}$, and $t_1^\ell, t_2^\ell \in [n]$ with $1 \leqslant t_1^\ell < t_2^\ell \leqslant n$, we define an automaton $[(j_t^\ell)_{t,\ell}, (t_1^\ell, t_2^\ell)_\ell]$ with state set

$$M = \big\{(\ell, t, r) \mid 1 \leqslant \ell < 2^{k(\log|J|-\delta)},\ t \leqslant t_2^\ell,\ 0 \leqslant r < n\big\},$$

initial state $(1, 1, 0)$, and action function

$$\alpha(\ell, t, r) = j_t^\ell.$$