Complexity and Effective Prediction

Abraham Neyman* and Joel Spencer†

April 28, 2008

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 player 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 were the second player is clairvoyant, i.e., would know the 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 $S(g)$.

*Institute of Mathematics and Center for the Study of Rationality, The Hebrew University of Jerusalem, Givat Ram, Jerusalem 91904, Israel (e-mail: aneyman@math.huji.ac.il).

†Department of Computer Science, and Department of Mathematics, New York University, USA (email: joelspencer@cims.nyu.edu).

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

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 [1] 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. [2] 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 $\liminf_{k\to\infty} \frac{\log m_k}{k} \geq \min(\log |I|, \log |J|)$.

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 [2] 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 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 [2]) 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 [2] 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 $\leq \max_{i\in I} \min_{j\in J} g(i, j) + \varepsilon$.

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 $(m_t)_{t\geq 1}$ 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

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

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 $\leq \eta$). 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 $\leq km$) 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* $v_* \leq$ *Val* $G(k,m) < v_* + \varepsilon$.

Proof. Obviously, Val $G(k,m) \geq v_*$. In order to prove the other inequality we assume without loss of generality that $|J| \leq |I|$. 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 $j_t^\ell \in J$, $1 \leq t \leq n$, $1 \leq \ell \leq \bar{\ell}$, and $t_1^\ell, t_2^\ell \in [n]$ with $1 \leq t_1^\ell < t_2^\ell \leq n$, we define an automaton $[(j_t^\ell)_{t,\ell}, (t_1^\ell, t_2^\ell)_\ell]$ with state set

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

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

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

The transition function is as follows.

$$\beta((\ell,t,r),i)=\begin{cases}(\ell,t+1,r) & \text{if } t<t_1^\ell\\ (\ell,t+1,r) & \text{if } t_1^\ell\le t<t_2^\ell \text{ and } g(i,j_t)>v_*\\ (\ell,t+1,r+1) & \text{if } t_1^\ell\le t<t_2^\ell \text{ and } g(i,j_t)\le v_*\\ (\ell,t_1^\ell+1,1) & \text{if } t=t_2^\ell, g(i,j_t)\le v_*, \text{ and}\\ & r\ge (t_2^\ell-t_1^\ell)(1-1/K)\\ (\ell+1,1,0) & \text{if } t=t_2^\ell \text{ and eithere } g(i,j_t)>v_* \text{ or}\\ & r<(t_2^\ell-t_1^\ell)(1-1/K)\end{cases}$$

If $\ell+1>2^{k(\log|J|-\delta)}$ we identify $\ell+1$ with 1.

The size of the automaton, $|M|$, is $\le 3k^2 2^{k(\log|J|-\delta)}$. Therefore,

$$\frac{\log|M|}{k}\le(\log|J|-\delta)+\frac{2\log k+\log 3}{k}.$$

Note that if in the play induced by a pure strategy of player 1 and the automaton, the states of the automaton of player 2 never visit the automaton state $(\bar{\ell},1,0)$, then the states of the automaton enter at some state $(\ell,t_1^\ell+1,1)$ a cycle of length $t_2^\ell-t_1^\ell$ and the average per-stage payoff over each cycle is $\le(1-1/K)v_*+\|g\|/K\le v_*+\varepsilon$.

We prove that if τ is the mixed strategy of player 2 defined as a random selection of the automaton $[(j_t^\ell)_{t,\ell},(t_1^\ell,t_2^\ell)_\ell]$, where j_t^ℓ are uniform iid ($1\le\ell\le\bar{\ell}$ and $1\le t\le n$) and $(j_t^\ell)_{t,\ell}$ independent of the uniform iid pairs $(t_1^\ell,t_2^\ell)_\ell$ (with $1\le t_1^\ell<t_2^\ell\le n$), then for every automaton of player 1 with $\le k$ states the probability of entering a cyclic play with average payoff per stage $\le v_*+\varepsilon$ is close to 1.

Let $j:I\to J$ be a function such that $g(i,j(i))\le v_*$. For every sequence $y=(j_1,\ldots,j_n)$ of actions of player 2 and every pure strategy σ of player 1 we denote by $S(\sigma,y)$ the set of all stages $1\le t\le n$ with $j_t\ne j(i_t)$ where $i_1,\ldots,i_n$ is the sequence of actions of player 1 when the pure strategy σ plays against the fixed sequence of actions $j_1,\ldots,j_n$ of player 2; $i_1=\sigma(\emptyset)$ and $i_t=\sigma(i_1,j_1,\ldots,i_{t-1},j_{t-1})$.

A subset S of $[n]$ is called (δ,K)-admissible if $|S|\le\frac{\delta n}{3K}$ and $|s-t|>2K$ for every two elements $t,s\in S$. Fix a pure strategy σ of player 1. First, let us count the number of sequences $y=(j_1,\ldots,j_n)\in J^n$ such that $S(\sigma,y)$ is (δ,K)-admissible.

The number of (δ, K)-admissible subsets $S \subset [n]$ is at least the number of subsets of no more than $\frac{\delta n}{3K}$ elements of a set with $[n - \frac{\delta n}{3K}(2K+1)] \geq k$ elements, and thus at least $\binom{k}{[\frac{\delta n}{3K}]}$, which is $\geq 2^{H(\frac{\delta}{3K})k}$ for sufficiently large k. If $|S| \leq [\frac{\delta n}{3K}]$ then the number of sequences $y = (j_1, \ldots, j_n)$ such that $S(\sigma, y) = S$ equals $(|J|-1)^{|S|}$. We deduce that for every pure strategy σ of player 1 the number of sequences $y = (j_1, \ldots, j_n) \in J^n$ such that $S(\sigma, y)$ is (δ, K)-admissible is $\geq (|J|-1)^{[\frac{\delta n}{3K}]} 2^{H(\frac{\delta}{3K})k}$.

For every subset S of $[n]$ we denote by $S(K)$ the set of all $1 \leq t \leq n$ with $|t - s| \leq K$ for some $s \in S$. Note that $|S(K)| \leq |S|(2K+1)$. Therefore, if $|S| \leq \frac{\delta n}{3K}$ then $n - |S(K)| \geq k$.

Therefore, if A is an automaton of player 1 of size $\leq k$, and $m_1, \ldots, m_n$ are the sequence of states of the automaton A when playing against the sequence $y = (j_1, \ldots, j_n)$, and $|S(\sigma^A, y)| \leq \frac{\delta n}{3K}$, then there are two distinct stages $t_1 < t_2$ in $[n] \setminus S(\sigma^A, y)(K)$ with $m_{t_1} = m_{t_2}$.

Therefore, the probability that the random play defined by the pair of strategies σ^A of player 1 and τ of player 2 will enter at the (random) stage t_1^1 a cycle of length $t_2^1 - t_1^1$ is

$$\geq \frac{1}{2k^2} 2^{H(\frac{\delta}{3K})k} 2^{-\log|J|k(1+\delta)} \geq 2^{k(2\delta - \log|J|)}$$

where the last inequality (uses the inequality $H(\frac{\delta}{3K}) > 3\delta + \delta \log |J|$ and) holds for k sufficiently large so that $\frac{1}{2k^2} > 2^{-\delta k}$.

The same computation shows in fact that if $t(\ell)$ is the first stage where the automaton of player 2 reaches state $(\ell, 1, 0)$ then conditional on $t(\ell) < \infty$ and conditional on $(j_t^{\ell'})_{\ell' < \ell}$ and $(t_1^{\ell'}, t_2^{\ell'})_{\ell' < \ell}$ the probability that the random play defined by the pair of strategies σ^A of player 1 and τ of player 2 will enter at the (random) stage $t(\ell) + t_1^\ell - 1$ a cycle of length $t_2^\ell - t_1^\ell$ is

$$\geq \frac{1}{2k^2} 2^{H(\frac{\delta}{3K})k} 2^{-\log|J|k(1+\delta)} \geq 2^{k(2\delta - \log|J|)}$$

where the last inequality holds for k sufficiently large so that $\frac{1}{2k^2} > 2^{-\delta k}$.

Therefore, if $\bar{\ell} = [2^{k(\log|J| - \delta)}]$, then the probability that $t(\bar{\ell}) < \infty$ is

$$\leq (1 - 2^{k(2\delta - \log|J|)})^{\bar{\ell} - 1} \to 0.$$

This completes the proof. $\square$

Theorem 2 *Let I and J be two finite sets. There exists a stage-game payoff function $g : I \times J \to \{0, -1\}$ with $\max_{i \in I} \min_{j \in J} g(i, j) = -1$ such that $\forall \varepsilon > 0\ \exists \delta > 0$ and k_0 sufficiently large such that if $k > k_0$ and $\log m < k \min(\log |I|, \log |J|) - k\varepsilon$ then the value of $G(k, m)$ (where $G = \langle I, J, g \rangle$) is $\geq -1 + \delta$.*

Proof. Let I^* be a subset of I with $|I^*| = \min(|I|, |J|)$ and $r : I^* \to J$ a 1-1 function. Define the stage-game payoff function g by $g(i, j) = 0$ for $i \in I^*$ and $j \neq r(i)$ and $g(i, j) = -1$ otherwise. Let p be the uniform distribution on I^*, and let $\mathcal{Q}(p, \theta)$ be the set of all distributions $Q \in \Delta(I \times J)$ with marginal p on I (i.e., $Q_I = p$) and $H(Q_I) + H(Q_J) - H(Q) \leq \theta$. Note that if $\theta < \log |I^*|$ then for every $Q \in \mathcal{Q}(p, \theta)$ we have $E_Q g(i, j) > -1$. As the set $\mathcal{Q}(p, \theta)$ is closed we deduce that $\min_{Q \in \mathcal{Q}(p,\theta)} E_Q g(i, j) > -1$. It follows from [4] that if $\log m_k \leq \theta < \log |I^*|$ then $\liminf_{k \to \infty}$ Val $G(k, m_k) \geq \min_{Q \in \mathcal{Q}(p,\theta)} E_Q g(i, j) > -1$. □

Theorem 1 and Theorem 2 imply

Corollary 1 *Let I and J be the the stage-strategy sets of player 1 and 2 respectively. A necessary and sufficient condition so that, for every stage-game payoff function $g : I \times J \to \mathbb{R}$, the values of $G(k, m_k)$ converge to $\max_{i \in I} \min_{j \in J} g(i, j)$ as k goes to infinity is $\liminf_{k \to \infty} \frac{\log m_k}{k} \geq \min(\log |I|, \log |J|)$.*

Obviously, for a given stage-game payoff function g, a weaker asymptotic relation may suffice for the values of $G(k, m_k)$ to converge to $\max_{i \in I} \min_{j \in J} g(i, j)$. The characterization of the necessary and sufficient asymptotic relation corresponding to a given stage-game payoff function is left open.

4 Applications and open problems

Let $G = \langle I, J, g^1, g^2 \rangle$ be a two-person non-zero-sum game. The set $F = F(G)$ of feasible payoffs in the repeated game equals the convex hull of the set of single-stage payoffs $\{(g^1(i, j), g^2(i, j)) : i \in I,\ j \in J\}$. The set of feasible payoffs of the game $G(k, m)$ is a subset of F that converges to F as k and m go to infinity.

Let G_1 and G_2 be the two-person zero-sum games $\langle I, J, g^1 \rangle$ and $\langle I, J, -g^2 \rangle$ respectively. Assume that for a given $\theta \geq 0$, for all sequences $m_k \geq k$ with $\lim_{k \to \infty} \frac{\log m_k}{k} \to_{k \to \infty} \theta$ the values of the games $G_1(k, m_k)$ and $G_2(k, m_k)$

converge to $v^1(\theta)$ and $-v^2(\theta)$. Note that [1] proves that this assumption holds for $\theta = 0$, [3] proves that it holds for $\theta > \log|J|$ (and thus also for $\theta > \min(\log|J|, \log|I|)$, and the present paper proves (in particular) that it holds for $\theta \geq \min(\log|J|, \log|I|)$.

In [2] it is conjectured that the assumption holds for all $\theta \geq 0$ with the possible exception of one value $\theta = \theta_0 > 0$. If the assumption holds for *some* $\theta \geq 0$ we obtain the following folk theorem: the set of equilibrium points of the game $G(k, m_k)$ converges, as $m_k \geq k \to \infty$ and $\frac{\log m_k}{k} \to_{k\to\infty} \theta$, to the set of all points $x = (x_1, x_2)$ in F with $x_1 \geq v^1(\theta)$ and $x_2 \geq v^2(\theta)$.

The folk theorem that corresponds to $\theta \geq \min(\log|J|, \log|I|)$ asserts that the set of equilibrium points of the game $G(k, m_k)$ converges, as $k \to \infty$ and $\liminf_{k\to\infty} \frac{\log m_k}{k} \geq \min(\log|J|, \log|I|)$, to the set of all points $x = (x_1, x_2)$ in F with $x_1 \geq \max_{i\in I} \min_{j\in J} g^1(i,j)$ and $x_2 \geq \min_{i\in I} \max_{j\in J} g^2(i,j)$.

The result of the present paper is insufficient for the complete asymptotic characterization of the equilibrium points of the finitely repeated (T-stage) game $G^{T_k}(k, m_k)$ when $\frac{\log m_k}{k} \to \log|J|$. This is because we have not quantified how long it takes the smarter player to outguess the actions of the other player. Such quantification can be formulated by the properties of the value of the game $G^{T_k}(k, m_k)$. It is known that the value of $G^{T_k}(k, m_k)$ (where G is a two-person zero-sum game $\langle I, J, g\rangle$) converges to $\max_{i\in I} \min_{j\in J} g(i,j)$ as $k \to \infty$, whenever $k \log k = o(T_k)$ and m_k is sufficiently large, e.g., $m_k = \infty$ or $m_k > |I|^{T_k}$ (see [5]), and it is conjectured in [2] that it converges to the value of the stage game ($= \max_{p\in\Delta(I)} \min_{j\in J} \sum_{i\in I} p(i)g(i,j)$) whenever $T_k = o(k \log k)$ and $m_k = \infty$. A positive answer to this conjecture, together with [3], results in a complete asymptotic characterization of the equilibrium payoffs of the game $G^{T_k}(k, m_k)$, when $m_k \geq k \geq (1+\varepsilon)T_k$.

References

[1] Ben-Porath, E. (1993). Repeated games with finite automata, *Journal of Economic Theory*, **59**, 17–32. Based on Ban-Porath's Master thesis (1986).

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

[3] Neyman, A. (1998). Finitely repeated games with finite automata, *Mathematics of Operations Research*, **23**, 513–552.

[4] Neyman, A. (2006). The value of memory (tentative title). Forthcoming.

[5] Neyman, A. and D. Okada (2000). Two-person repeated games with finite Automata, *International Journal of Game Theory*, **29**, 309–325.