Excludability and Bounded Computational Capacity Strategies

Ehud Lehrer* and Eilon Solan†

September 9, 2003

Abstract

We study the notion of excludability in repeated games with vector payoffs, when one of the players is restricted to strategies with bounded computational capacity. We show that a closed set that does not contain a convex approachable set is excludable by player 2 when player 1 is restricted to use only bounded-recall strategies. We also show that when player 1 is restricted to use strategies that can be implemented by finite automata, player 2 can exclude him from any closed, non-convex set whose convex hull is minimal convex approachable.

*School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel; e-mail: lehrer@post.tau.ac.il

†School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel, and Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208-2001. e-mail: eilons@post.tau.ac.il, e-solan@kellogg.northwestern.edu

1 Introduction

In a seminal paper, Blackwell (1956) introduced and studied the notions of approachability and excludability in repeated games with vector payoffs, which are the analogues of the max-min level and the min-max level in standard repeated games with scalar payoffs.

In a repeated game with vector payoffs, a set is *approachable* by player 1 if player 1 has a strategy that guarantees, with arbitrarily high probability, that the long-run average payoff vector remains arbitrarily close to the set from some stage on, regardless of the strategy employed by player 2. A set is *excludable* by player 2 if player 2 has a strategy that guarantees, with arbitrarily high probability, that the long-run average payoff vector remains far from the set from some stage on, regardless of the strategy employed by player 1.

Blackwell (1956) provided a geometric condition that guarantees that a set is approachable, and proved that any convex set is either approachable by player 1, or excludable by player 2. Hou (1971) and Spinat (2002) fully characterized the family of approachable sets. Both proved that if a closed set is *minimal* (w.r.t. set inclusion) approachable then it satisfies Blackwell's geometric condition. Vieille (1992) studied the notions of weak-approachability and weak-excludability, which were also introduced by Blackwell (1956), and proved that any set is either weak-approachable or weak-excludable. For partial results on weak-approachability for 2-dimensional games were both players have 2 actions see Hou (1969).

We are interested in studying repeated games with vector payoffs when one of the players is restricted to strategies with bounded computational capacity. Two classes of such strategies that were extensively studied in the literature are strategies that can be implemented by finite automata (see, e.g., Neyman, 1985, Rubinstein, 1986, Kalai, 1990), and bounded-recall strategies, that is, strategies that may condition only on the last k pairs of actions played in the game, for some fixed k (see, e.g., Lehrer, 1988, Aumann and Sorin, 1989).

In a companion paper (Lehrer and Solan, 2003) we studied the notions of approachability with automata, and approachability with bounded-recall strategies. That is, which sets are approachable by player 1, if he can only use strategies that can be implemented by automata, or bounded-recall strategies (and player 2 is not restricted). We proved that the following three statements are equivalent for closed sets.

- F is approachable with automata.
- F is approachable with bounded-recall strategies.
- F contains a convex approachable set.

In the present paper we concentrate on the notions of excludability against automata, and excludability against bounded-recall strategies. A set is *excludable against automata* (resp. *against bounded-recall strategies)* by player 2 if player 2 has a strategy that ensures that when player 1 plays a finite automaton (resp. a bounded-recall strategy), the long-run average payoff is bounded away from the set.

We provide a complete characterization for the family of sets which are excludable against bounded-recall strategies: a set is excludable against bounded-recall strategies if and only if it is not approachable by bounded-recall strategies. Namely, if and only if it does not contain any convex approachable set. In the proof we analyze connectedness aspects of the set of short histories when one of the players is restricted to bounded-recall strategies.

Our result concerning excludability against automata is not as sharp. A set is *C-minimal approachable* if it is closed, approachable, and its convex hull does not strictly contain a convex approachable set. We show that a non-convex C-minimal approachable set is excludable against automata. Since any convex approachable set is approachable with automata, this implies that a C-minimal approachable set is either convex and approachable with automata, or non-convex and excludable against automata. For this result we study the Markov chain induced over the states of the automaton when the opponent plays a stationary strategy.

We do not know whether any set which is not approachable with automata by player 1 is excludable against automata by player 2. Solving this problem will reveal whether the two notions related to strategies with bounded computational capacity are equivalent.

2 The Model and the Main Results

2.1 Repeated games with vector payoffs

In this section we define repeated games with vector payoffs.

A two-player *repeated game with vector payoffs* is a triplet (I, J, M), where I and J are finite sets of actions for the two players, and $M = (m_{i,j})_{i \in I, j \in J}$ is a d-dimensional vector payoff matrix, so that $m_{i,j} \in I\!R^d$ for every $i \in I$ and $j \in J$. We assume throughout that $\|M\|_\infty \leq 1$; that is, all payoffs are bounded by 1. We denote by $\mathcal{S}$ and $\mathcal{T}$ the sets of strategies in the repeated game of the players 1 and 2, respectively.

For every stage t denote by (i_t, j_t) the joint action played by players 1 and 2 at stage t. The average payoff vector up to stage n is $\bar{x}_n = \frac{\sum_{t=1}^n m_{i_t,j_t}}{n}$.

2.2 On bounded-capacity strategies

In this section we define two types of bounded computational capacity strategies: strategies with bounded recall, and strategies that can be implemented by automata.

Let $k \in I\!N$ be a natural number. A *k-bounded-recall strategy* of player 1 (resp. player 2) is a pair (m, σ) (resp. (m, τ)) where $m \in (I \times J)^k$ and $\sigma : (I \times J)^k \to \Delta(I)$ (resp. $\tau : (I \times J)^k \to \Delta(J)$). When playing a k-bounded-recall strategy (m, σ), at any stage player 1 plays $\sigma(x)$, where x is the string of the last k joint actions. He starts the game with the (virtual) memory m. Thus, at the first stage he plays the mixed action $\sigma(m)$, at the second stage he plays $\sigma(m', i_1, j_1)$, where m' are the last $k-1$ coordinates of m and (i_1, j_1) is the realized pair of actions of the two players at the first stage, and so on.

A *bounded recall* strategy is a strategy which is k-bounded recall, for some $k \in I\!N$. We denote by $\mathcal{S}_{BR}$ the set of all bounded-recall strategies of player 1.

A (non-deterministic) *automaton* A is given by (i) a finite set of states, (ii) a probability distribution over the set of states, according to which the initial state is chosen, (iii) a finite set of inputs, (iv) a finite set of outputs, (v) a function that assigns to every state a probability distribution over outputs, and (vi) a transition rule, that assigns to every state and every input a probability distribution over states.

When the set of outputs coincides with the set I of actions of player 1, and the set of inputs coincides with the set $I \times J$ of action pairs, an automaton implements a strategy for player 1 as follows. The initial state of the automaton is chosen according to the initial distribution given in (ii). At every stage, as a function of the current state an action of player 1 is chosen by the probability distribution given in (v), and a new state is chosen as a function of the pair of actions played (by both players), according to the probability distribution given in (vi). $\mathcal{S}_A$ denotes the set of player 1's strategies that can be implemented by an automaton.

Observe that every k-bounded-recall strategy can be implemented by an automaton with $|I \times J|^k$ states.

2.3 Excludability against bounded-capacity strategies

Let $d(x, y)$ denote the Euclidean distance between the points x and y in $\mathbb{R}^d$. For every set F in $\mathbb{R}^d$ and every $x \in \mathbb{R}^d$, let $d(x, F) = \inf_{y \in F} d(x, y)$ be the distance of x from F. For every $\delta > 0$, let $B(F, \delta) = \{x \in \mathbb{R}^d : d(x, F) \leq \delta\}$ be the set of all points which are δ-close to F.

Blackwell (1956) defined the notion of excludability in repeated games with vector payoffs. A set F is excludable if player 2 can guarantee with arbitrarily high probability that the long-run average payoff will never get close to F from some point on.

Definition 2.1 (Blackwell, 1956) *A set F is* excludable *by player 2 if there exists a strategy $\tau \in \mathcal{T}$ such that*

$$\exists \varepsilon > 0, \forall \eta > 0, \exists N \in \mathbb{N}, \forall \sigma \in \mathcal{S}, \qquad \mathbf{P}_{\sigma,\tau}(\inf_{n \geq N} d(\bar{x}_n, F) < \varepsilon) < \eta.$$

Here, $\mathbf{P}_{\sigma,\tau}$ is the probability distribution over the space of infinite plays induces by the strategy pair (σ, τ).

We are interested in studying when a given set is excludable by player 2, provided player 1 is restricted to use bounded-recall strategies, and strategies that can be implemented by automata.

Definition 2.2 *A set F is* excludable against bounded-recall strategies *by player 2 if there exists a strategy $\tau \in \mathcal{T}$ such that*

$$\exists \varepsilon > 0, \forall \eta > 0, \forall \sigma \in \mathcal{S}_{BR}, \exists N \in \mathbb{N}, \qquad \mathbf{P}_{\sigma,\tau}(\inf_{n \geq N} d(\bar{x}_n, F) < \varepsilon) < \eta.$$

The set is excludable against automata *if a similar condition holds, when $\mathcal{S}_{BR}$ is replaced by $\mathcal{S}_A$.*

Observe that in this definition, N depends on the strategy used by player 1, whereas in Definition 1, it does not. If N is required to be independent of the strategy employed by player 1, excludability against bounded-recall strategies (or against automata) turns out to be equivalent to excludability. Nevertheless, it is desirable that N depends only on the *size* of the memory of the strategy (or on the size of the automaton), and not on the strategy itself. Studying excludability against bounded computational capacity strategies under this stronger definition is left for future research.

2.4 C-Minimal Approachable Sets and Sharp Points

To present our results, we define the dual notion of approachability.

Definition 2.3 (Blackwell, 1956) *A set F is* approachable *by player 1 if there exists a strategy $\sigma \in \mathcal{S}$ such that*

$$\forall \varepsilon > 0, \forall \eta > 0, \exists N, \forall \tau \in \mathcal{T}, \qquad \mathbf{P}_{\sigma,\tau}(\sup_{n \geq N} d(\bar{x}_n, F) \geq \varepsilon) < \eta.$$

In this case we say that σ approaches *F.*

A set F is approachable if player 1 can guarantee with arbitrarily high probability that the long-run average payoff will be arbitrarily close to F. A closed set F is *minimal approachable* if there is no proper closed subset of F which is approachable.

For every pair (p, q) of mixed actions (i.e., distributions over the respective action sets), denote by $m_{p,q} = \sum_{i,j} p_i m_{i,j} q_j$ the expected vector payoff. This is the expected stage-payoff when player 1 plays the mixed action p and player 2 plays the mixed action q.

Definition 2.4 (Spinat, 2002) *A closed set F is a B*-set *if for every $z \in \mathbb{R}^d$ there is a mixed action $p_F(z)$ of player 1, and a point $y \in A$, such that (i) $d(z, y) = d(z, F)$, and (ii) for every mixed action q of player 2, $\langle z - y, m_{p_F(z),q} - y \rangle \leq 0$.*[1]

Blackwell (1956) proved that every B-set is approachable, and Hou (1969, Theorem 3) and Spinat (2002, Theorem 4) proved that every minimal approachable set is a B-set.

We are going to study closed approachable sets whose convex hull does not strictly contain a convex approachable set. Denote by $\mathrm{conv}(F)$ the convex hull of F.

Definition 2.5 *A closed approachable set F is* C-minimal approachable *if there is no proper closed and convex subset of* $\mathrm{conv}(F)$ *which is approachable.*

[1] For $z, y \in \mathbb{R}^d$, $\langle z, y \rangle = \sum_{i=1}^{d} z_i y_i$ is the standard inner product.

For any mixed action q of player 2, denote $H(q) = \{m_{p,q}\colon p \text{ is a mixed action of player 1}\}$. When player 2 plays the mixed action q, the expected stage-payoff is always in $H(q)$, regardless of the mixed action chosen by player 2.

Definition 2.6 *A point x is a set F is* sharp *if for every $\varepsilon > 0$ there is some mixed action q of player 2 such that $H(q) \cap F \subseteq B(x, \varepsilon)$.*

Definition 2.7 *A point x in a convex set $F \subseteq I\!R^d$ is* exposed *if x is a unique point at which a linear functional attains its maximum over F. That is, x is exposed if there is $y \in I\!R^d$ such that $\langle x, y\rangle > \langle x', y\rangle$ for every $x' \in F \setminus \{x\}$.*

Example 2.8 *Consider the following game, where both players have two actions, payoffs are two dimensional, and the payoffs matrix is*

$-1, 1$	$1, -1$
$2, 2$	$-1, -1$

For every mixed action $q = (q_L, 1 - q_L)$ of player 2, where q_L is the probability that player 2 plays the left column,

$$H(q) = \text{conv}\{(1 - 2q_L, -1 + 2q_L), (-1 + 3q_L, -1 + 3q_L)\}.$$

Denote $F = [(0, 0), (2, 2)] \cup [(0, 0), (1, -1)]$. This is the union of two intervals. Thus, $\text{conv}(F)$ *is the triangle with vertices $(0, 0)$, $(1, -1)$ and $(2, 2)$.*

The points $(0, 0)$, $(1, -1)$ and $(2, 2)$ are the exposed points of $\text{conv}(F)$. The points $(1, -1)$ and $(2, 2)$ are also sharp in $\text{conv}(F)$, since $H((0, 1)) \cap \text{conv}(F) = (1, -1)$ and $H((1, 0)) \cap \text{conv}(F) = (2, 2)$. The point $(0, 0)$ is not sharp in $\text{conv}(F)$. Indeed, when player 1 plays the mixed action $p = (\frac{3}{5}, \frac{2}{5})$ he ensures that for every mixed action q of player 2, the first coordinate of $m_{p,q}$ is $\frac{1}{5}$, so that $H(q) \cap \text{conv}(F) \not\subseteq B((0, 0), \varepsilon)$ for every mixed action q, provided $\varepsilon < \frac{1}{5}$. For that reason all points in a sufficiently small neighborhood of $(0, 0)$ are not sharp in $\text{conv}(F)$ as well. In particular, $\text{conv}(F)$ does not contain the convex hull of its sharp points.

One can verify that Blackwell's geometric condition for approachable sets implies that F is minimal approachable.

We now argue that F is not C-minimal approachable. Define z_β to be the point $(-1, \beta)$, $-1 \leq \beta \leq 1$. For every $-1 \leq \beta \leq 1$ the closest point to z_β in F is the origin. For $\beta \neq 0$, the graph of $g_\beta(x) = \frac{1-\beta^2}{\beta(\beta+5)} - \frac{x}{\beta}$ is perpendicular to z_β. Player

1 has a mixed action $p_\beta = (\frac{3\beta+3}{\beta+5}, \frac{2-2\beta}{\beta+5})$ *such that for every mixed action* q *of player 2, the point* $m_{p_\beta,q}$ *lies on the right side of the graph of* $g_\beta(x)$, $\beta \neq 0$. *When* $\beta < 0$ *the graph of* g_β *meets the interval* $[(0,0),(2,2)]$ *at the point* $(\frac{1-\beta}{\beta+5}, \frac{1-\beta}{\beta+5})$ *and when* $\beta > 0$ g_β *meets the interval* $[(0,0),(1,-1)]$ *at the point* $(\frac{1+\beta}{\beta+5}, \frac{1+\beta}{\beta+5})$. *Since* $(\frac{1-\beta}{\beta+5}, \frac{1-\beta}{\beta+5})$ *and* $(\frac{1+\beta}{\beta+5}, \frac{1+\beta}{\beta+5})$ *are bounded away from the origin, there is a small* $\lambda > 0$ *such that* $\mathrm{conv}(F) \setminus \mathrm{conv}\{(0,0),(\lambda,\lambda),(\lambda,-\lambda)\}$ *is still a convex approachable set. Therefore,* F *is not C-minimal.*

The following lemma shows that Example 1 is not coincidental.

Lemma 2.9 *Suppose that* $\mathrm{conv}(F)$ *is approachable. Then* F *is C-minimal approachable if and only if every exposed point of* $\mathrm{conv}(F)$ *is sharp.*

Proof. We start with the "only if" direction. If F is a singleton the lemma trivially holds. We therefore may assume that F is not a singleton. Let x be an exposed point of $\mathrm{conv}(F)$, so that there exists $y \in I\!R^d$ such that $\langle x', y\rangle < \langle x, y\rangle$, for every $x' \in F \setminus \{x\}$. We have to prove that x is sharp.

Denote, for every $\delta > 0$,

$$M_\delta(x) = \{x' \in F \colon \langle x', y\rangle > \langle x, y\rangle - \delta\}.$$

Since F is not a singleton, for $\delta > 0$ sufficiently small $\mathrm{conv}(F \setminus M_\delta(x))$ is non-empty, closed, convex and a strict subset of $\mathrm{conv}(F)$. Since $\mathrm{conv}(F)$ is C-minimal approachable, $\mathrm{conv}(F \setminus M_\delta(x))$ is not approachable. By Blackwell (1956) every B-set is approachable, and thus $\mathrm{conv}(F \setminus M_\delta(x))$ is not a B-set. Therefore, for every[2] $\delta > 0$ there exist $z_\delta \in I\!R^d$ and $y_\delta \in \mathrm{conv}(F \setminus M_\delta(x))$ that satisfy $d(z_\delta, y_\delta) = d(z_\delta, \mathrm{conv}(F \setminus M_\delta(x)))$ such that for every $p \in \Delta(I)$ there is $q \in \Delta(J)$ with $\langle z_\delta - y_\delta, m_{p,q} - y_\delta\rangle > 0$. By the Min-Max Theorem, for every $\delta > 0$ there exists $q_\delta \in \Delta(J)$ such that

$$\langle z_\delta - y_\delta, m_{p,q_\delta} - y_\delta\rangle > 0 \qquad \forall p \in \Delta(I). \tag{1}$$

Since y_δ is the closest point to x_δ in $\mathrm{conv}(F \setminus M_\delta(x))$,

$$\langle z_\delta - y_\delta, a - y_\delta\rangle \leq 0 \qquad \forall a \in \mathrm{conv}(F \setminus M_\delta(x)). \tag{2}$$

Eqs. (1) and (2) imply that $H(q_\delta) \cap \mathrm{conv}(F \setminus M_\delta(x)) = \emptyset$. Thus, $H(q_\delta) \cap \mathrm{conv}(F) \subseteq M_\delta(x)$. However, since x is exposed, $\cap_{\delta>0} M_\delta(x) = \{x\}$, and therefore for every $\varepsilon > 0$

[2] Here and below, when we say "for every $\delta > 0$" we mean "for every $\delta > 0$ such that $F \setminus M_\delta(x)$ is not empty."

there is $\delta > 0$ such that $M_\delta(x) \subset B(x, \varepsilon)$. Thus, for every $\varepsilon > 0$ there is q_δ such that $H(q_\delta) \cap F \subseteq H(q_\delta) \cap \mathrm{conv}(F) \subseteq B(x, \varepsilon)$, so that x is sharp.

As for the "if" direction, assume that $\mathrm{conv}(F)$ is approachable and all its exposed points are sharp. Suppose to the contrary that there is a proper closed and convex subset G of $\mathrm{conv}(F)$ which is approachable. Thus, not all extreme points of $\mathrm{conv}(F)$ are in G.

Straszewicz Theorem (see Rockafellar, 1970, p. 167) states that the set of exposed points is dense in the set of extreme points. Therefore, there is an exposed point x of $\mathrm{conv}(F)$ which is not in G. Furthermore, since G is closed, there is $\varepsilon > 0$ such that $B(x, 2\varepsilon) \cap G = \emptyset$. By assumption this exposed point is sharp in $\mathrm{conv}(F)$. It means that there is a mixed action of player 2, q_ε, such that $H(q_\varepsilon) \cap G = \emptyset$. However, this implies that G is not approachable by player 1. Indeed, if player 2 plays the mixed action q_ε at every stage, the long-run average payoff is in $H(q_\varepsilon)$, which is bounded away from G. This contradicts the assumption that G is approachable. ■

2.5 The Main Results

Our main result concerning excludability against automata is the following.

Proposition 2.10 *A closed set F that does not contain the convex hull of its sharp points is excludable against automata.*

Corollary 2.11 *Let F be a non-convex closed C-minimal approachable set. Then, F is excludable against automata.*

Proof. By Straszewicz Theorem the closure of the convex hull of the set of exposed points of $\mathrm{conv}(F)$ is $\mathrm{conv}(F)$. Lemma 1 implies that the exposed points of $\mathrm{conv}(F)$ are sharp, and therefore the closure of the convex hull of the sharp points in F is $\mathrm{conv}(F)$. Since F is non-convex and closed, we conclude that F does not contain the convex hull of its sharp points. The corollary follows from Proposition 2. ■

By Lehrer and Solan (2003, Proposition 1), any convex approachable set is approachable with automata. This result, along with Corollary 1, implies the following.

Corollary 2.12 *A closed C-minimal approachable set is either convex and approachable with automata, or non-convex and excludable against automata.*

Our main result concerning excludability against bounded-recall strategies is the following.

Proposition 2.13 *A closed set that does not contain a convex approachable set is excludable against bounded-recall strategies.*

The dual notion to excludability against bounded-recall strategies is approachability with bounded-recall strategies.

Definition 2.14 (Lehrer and Solan, 2003) *A set F is* approachable with bounded-recall strategies *by player 1 if for every $\delta > 0$ there exists $k \in \mathbb{N}$ and a k-bounded-recall strategy σ that approaches $B(F, \delta)$.*

Lehrer and Solan (2003, Proposition 2) proved that any convex approachable set is approachable with bounded-recall strategies. This, together with Proposition 3, implies that any set that is not approachable with bounded-recall strategies is excludable against bounded-recall strategies.

To make the discussion complete, we briefly mention the notion of excludability with a bounded computational capacity strategy. A set is *excludable with a bounded-recall strategy* (resp. *with automata*) by player 1 if player 1 has a bounded-recall strategy (resp. a strategy that can be implemented by an automaton) that ensures the long-run average payoff remains away from the set. Theorem 1 in Lehrer and Solan (2003) implies the following.

Proposition 2.15 *A set is excludable with bounded-recall strategies (or with automata) by player 1 if and only if the interior of the complement of the set contains a convex approachable set.*

Before proving the results, we explain the main ideas using Example 1.

Example 1 - continued *As argued before, the set F is approachable, but not C-minimal approachable. By Proposition 2 F is excludable against automata. We now explain how to construct, for this example, a strategy that ensures that the average payoff remains far from F. Thus, we assume that player 1 employs a strategy that can be implemented by an automaton.*

Suppose that player 2 plays the stationary strategy $(0,1)$; that is, he always plays the right column. This means that the evolution of the state of the automaton that implements the strategy of player 1 follows a Markov chain, so that the average payoff *converges to a limit. Since $H((0,1)) \cap F = (1,-1)$, this limit is either out of F, or equal to $(1,-1)$.*

Similarly, when player 2 plays the stationary strategy $(1,0)$, since $H((1,0)) \cap F = (2,2)$, the limit of the average payoff is either out of F, or equal to $(2,2)$.

To ensure that the average payoff remains out of F, player 2 plays in blocks. In some blocks he plays the stationary strategy $(0,1)$, and in some the stationary strategy $(1,0)$. Player 2 tries to balance the number of stages in which each stationary strategy is used, so that whenever a new block begins, he chooses that stationary strategy that was used less frequently in previous blocks. The block ends once average payoff gets close to $(1,-1)$ or $(2,2)$ respectively. Since the evolution of the state of the automaton follows a Markov chain, if a block is finite, one can bound its length. If some block never ends, average payoff remains far from F. If all blocks are finite, then the frequency of stages in which each stationary strategy is used converges to $\frac{1}{2}$, so that the average payoff is close to $\frac{1}{2}(1,-1) + \frac{1}{2}(2,2) = (\frac{3}{2}, \frac{1}{2})$, which is out of F.

A small caveat is that player 2 does not know the strategy of player 1, so that he does not know how fast the average payoff converges to its limit. Therefore the length of block k increases with k.

3 Excludability

3.1 Excludability against automata

Here we prove Proposition 2, which states that any closed set that does not contain the convex hull of its sharp points is excludable against automata.

The following estimate on Markov chains with finite state space will be useful. Chebychev's inequality (see Shiryayev, 1984, p. 121) and the decomposition theorem of irreducible chains (see Feller, 1968, p. 405) imply that for every finite irreducible Markov chain[3] there is a constant $C > 0$ such that for every two states s, s', and every $\lambda > 0$,

$$\mathbf{P}\left(\left|\frac{\mathbb{1}_{s_1=s} + \ldots + \mathbb{1}_{s_n=s}}{n} - \pi(s)\right| > \lambda \mid \text{ the initial state is } s'\right) < \frac{C}{n\lambda^2}, \qquad (3)$$

where $\mathbb{1}_{s_t=s} = 1$ if the state at stage t is s and 0 otherwise, and $\pi = (\pi(s))_s$ is the (unique) invariant distribution.

[3] A finite Markov chain is *irreducible* if any state can be reached from any other state.

Proof of Proposition 2: Let F be a closed set that does not contain the convex hull of its sharp points. Thus, there are sharp points $x_1, ..., x_L$ and non-negative numbers, $\lambda_1, ..., \lambda_L$, such that $\sum_{l=1}^{L} \lambda_l = 1$ and $z := \sum_{l=1}^{L} \lambda_l x_l \notin F$. Choose $\delta > 0$ such that $d(z, F) > 2\delta$. Thus, $d(\sum_{l=1}^{L} \lambda_l B(x_l, \delta), F) > \delta$.[4]

Since $(x_l)_{l=1}^{L}$ are sharp, there are mixed actions $q_1, ..., q_L \in \Delta(J)$ such that $H(q_l) \cap \text{conv}(F) \subseteq B(x_l, \frac{\delta}{2})$, for $l = 1, ..., L$.

We now define a strategy τ that plays in blocks whose lengths are history dependent. Below we prove that this strategy guarantees that the average payoff remains far from F.

- There are L types of blocks, one for each $l = 1, ..., L$. Block of type l is referred to as an l-block.
- For every l, at every stage of an l-block τ plays the mixed action q_l.
- Suppose that the k-th block is an l-block. The block terminates when two conditions are simultaneously satisfied: (i) the average payoff within this block is in $B(x_l, \delta)$, and (ii) the length of the block is at least $k^4 + k$. If the average payoff along any prefix of the block longer than $k^4 + k$ is never in the respective ball, the block never terminates.
- Suppose that at stage n the k-th block terminates. Let m_l be the overall number of stages up to stage n spent in l-blocks (in particular, $\sum_{l=1}^{L} m_l = n$). The next block is a j-block if j is the minimal index in $\text{argmax}_{l=1,...,L}(\lambda_l - \frac{m_l}{n})$. Thus, if the gap between λ_j and the relative frequency of past stages spent in j-blocks is maximal, it means that j-blocks are under-played. To correct that, a j-block is played, and thereby the gap is narrowed.

To prove that F is excludable against automata, we fix an automaton A and $\eta > 0$, and we show that there is $N_{A,\eta} \in \mathbb{N}$ such that

$$\mathbf{P}_{A,\tau}\left(\inf_{n \geq N_{A,\eta}} d(\bar{x}_n, F) < \frac{\delta}{4}\right) < \eta. \tag{4}$$

To this end, we show below that with probability at least $1-\eta/2$, either (A.i) there are finitely many blocks, so that the last block never terminates, and the long-run

[4] For every two sets $A, B \in \mathbb{R}^d$ and every $\lambda_1, \lambda_2 > 0$, $\lambda_1 A + \lambda_2 B := \{\lambda_1 x + \lambda_2 y : x \in A, y \in B\}$, and $d(A, B) = \inf\{d(x, y) : x \in A, y \in B\}$.

average payoff is out of F, or (A.ii) there are infinitely many blocks, and the ratio between the length of the k-th block and the total length of all preceding blocks goes to 0, as k goes to ∞.

Suppose that (A.i) holds. Let k be the index of the last (infinite) block, and let l be its type. Since at every stage of the block player 2 plays the mixed action q_l, the average payoff along any prefix of the block is in $H(q_l)$. Since the block never terminates, the average payoff is not in $B(x_l,\delta)$. Since $H(q_l)\cap \mathrm{conv}(F)\subseteq B(x_l,\frac{\delta}{2})$, the average payoff is not in $B(\mathrm{conv}(F),\frac{\delta}{2})$.

Suppose that (A.ii) holds. For $l\in\{1,\ldots,L\}$ let $\pi_{l,n}$ be the proportion of stages prior to stage n that are spent in l-blocks. We claim that $(\pi_{l,n})_{n\in I\!N}$ converges to λ_l as n goes to infinity. Indeed, letting $\mathrm{len}(k)$ denote the length of the k-th block, we obtain that $\lambda_l-|L|\times\frac{\max_k \mathrm{len}(k)}{n}\leq \pi_{l,n}\leq \lambda_l+\frac{\max_k \mathrm{len}(k)}{n}$, where the maximum is over all k such that block k does not start after stage n. By (A.ii) $\frac{\max_k \mathrm{len}(k)}{n}$ goes to 0 as n goes to infinity, and the claim follows. Since the average payoff during (a finite) block k is in $B(x_l,\delta)$, this means that the long-run average payoff is in $B(\sum_{l=1}^{L}\lambda_l x_l,\delta)=B(z,\delta)$, which is disjoint of $B(F,\delta)$.

Choose $N_{A,\eta}$ sufficiently large such that with probability $1-\eta$ (A.i) and (A.2) hold, and, in addition,

(B.1) all blocks that start after stage $\delta N_{A,\eta}$ are finite, and

(B.2) for each $n\geq N_{A,\eta}$ and every $l=1,\ldots,L$, $|\pi_{l,n}-\lambda_l|<\frac{\delta}{2}$.

As explained in the preceding two paragraphs, on the respective set $\inf_{n\geq N_{A,\eta}} d(\bar{x}_n,F)\geq\frac{\delta}{2}$, and Eq. (4) holds.

We now turn to prove that with probability at least $1-\frac{\eta}{2}$ either (A.i) or (A.ii) hold. Denote by S the set of states of the automaton. Any mixed action q of player 2 induces a Markov chain $\mathcal{M}(q)$ over S, which reflects the evolution of the automaton when player 2 plays at every stage the mixed action q, regardless of past play. We denote by $\mathbf{P}_{(A,s),q}$ the law of this Markov chain, when $s\in S$ is the initial state.

By Seneta (1981, Theorem 4.7) there are constants $c_1>0$ and $\rho\in(0,1)$ such that for every $l=1,\ldots,L$, the probability that by stage k no irreducible subset for $\mathcal{M}(q_l)$

is reached is at most $c\rho^k$:[5,6,7]

$$\mathbf{P}_{(A,s),q_l}(s_k \text{ is in some irreducible subset for } \mathcal{M}(q_l)) \leq c\rho^k.$$

Therefore, there is $k_1 \in I\!N$ such that the probability that for every $k \geq k_1$, an irreducible subset (in the respective Markov chain) is reached by stage k of block k, is at least $1 - \eta/4$. Denote by $\mathcal{E}_1$ the corresponding event. Observe that on $\mathcal{E}_1$, in each block k the process spends at least k^4 stages in an irreducible subset.

By Eq. (3) there is a constant $c_2 > 0$ such that for every $l = 1, \ldots, L$, and every pair of states s, s' in the same irreducible set E for $\mathcal{M}(q_l)$,

$$\forall n \geq k^4, \quad \mathbf{P}_{(A,s'),q_l}\left(\left|\frac{1\!\!1_{s_1=s} + \ldots + 1\!\!1_{s_n=s}}{n} - \pi_l^E(s)\right| > \frac{1}{\sqrt{k}}\right) < \frac{c_2}{k^3}, \tag{5}$$

where $\pi_l^E(s)$ is the invariant distribution of $\mathcal{M}(q_l)$ in E. This implies that

$$\mathbf{P}_{(A,s'),q_l}\left(\left|\frac{1\!\!1_{s_1=s} + \ldots + 1\!\!1_{s_n=s}}{n} - \pi_l^E(s)\right| > \frac{1}{\sqrt{k}} \qquad \forall n \geq k^4\right) < \frac{c_2}{(k-1)^2}. \tag{6}$$

Eq. (6) bounds the probability that the empirical frequency of visits to s after n stages is close to the invariant distribution at s for *every* n sufficiently large, provided the initial state is in the same irreducible set as s. If l_k is the type of block k, L_k is its length, E_k is the irreducible set the process was absorbed to, and $\nu_k(s)$ is the empirical frequency of visits to each state $s \in E_k$ along block k, then Eq. (6) implies that

$$\mathbf{P}_{A,\tau}\left(\left|\nu_k(s) - \pi_{l_k}^{E_k}(s)\right| > \frac{1}{\sqrt{k}} \mid s_k \in E_k, L_k < \infty\right) < \frac{c_2}{(k-1)^2}. \tag{7}$$

Summing up Eq. (7) over k, we deduce that there is $k_2 \in I\!N$ such that

$$\mathbf{P}_{A,\tau}\left(\forall k \geq k_2, \quad |\nu_k(s) - \pi_{l_k}^{E_k}(s)| > \frac{1}{\sqrt{k}} \mid \mathcal{E}_1, L_k < \infty \ \ \forall k\right) < \frac{\eta}{4}. \tag{8}$$

Denote by $\mathcal{E} \subseteq \mathcal{E}_1$ the respective event. We show that on $\mathcal{E}$ either (A.i) or (A.ii) hold. As $\mathbf{P}_{A,\tau}(\mathcal{E}) > 1 - \frac{\eta}{2}$, this implies the desired result.

For every irreducible set E for $\mathcal{M}(q_l)$, denote by y_l^E the long-run average payoff when the initial state of the automaton A is within E, and player 2 plays the mixed

[5] A subset of states is *irreducible* if any state in the subset can be reached from any other state in the subset, and no state outside the subset can be reached from any state in the set.

[6] Observe that the family of irreducible sets depends on q_l.

[7] Below s_k is the state of the automaton at stage k.

action q_l at every stage. It is given by $y_l^E = \sum_{s\in E}\pi_l^E(s)m_{p_s,q_l} \in H(q_l)$, where p_s is the mixed action played by the automaton at state s.

Recall that $H(q_l)\cap\text{conv}(F) \subseteq B(x_l, \frac{\delta}{2})$. Set $\zeta = \frac{1}{2}\min\{d(y_l^E, \text{conv}(F))\colon d(y_l^E, \text{conv}(F)) > \frac{\delta}{2}\}$, and $k_0 = \max\{k_1, k_2, \frac{1}{(\zeta-\delta/2)^2}\}$. Let $k \geq k_0$.

We now show that on the event $\mathcal{E}$ either (A.i) or (A.ii) hold. Consider block k of type l, and suppose that the process is absorbed to the irreducible set E. If $y_l^E \notin B(x_l, \frac{\delta}{2})$ then $y_l^E \notin B(x_l, \zeta)$. Since $k \geq 1/(\zeta - \frac{\delta}{2})^2$, the average payoff in any prefix of the block longer than $k^4 + k - 1$ remains in $B(y_l^E, \zeta - \frac{\delta}{2})$, which is disjoint of $B(x_l, \frac{\delta}{2})$, so that the block is infinite.

If $y_l^E \in B(x_l, \frac{\delta}{2})$ then average payoff in any prefix of the block longer than k^4+k-1 is in $B(y_l^E, \frac{\delta}{2}) \subseteq B(x_l, \delta)$, so that the block terminates after $k^4 + k$ stages.

In particular, if (A.i) does not hold, so that all blocks are finite, the ratio between L_k and $\sum_{j<k} L_j$ is bounded by $\frac{k^4+k}{\sum_{j<k} j^4} < \frac{5k^4}{k^5} = \frac{5}{k}$, which goes to 0 as k goes to infinity, and (A.ii) holds. ■

3.2 Excludability against bounded-recall strategies

Here we prove Proposition 3, which states that a closed set that does not contain a convex approachable set is excludable against bounded-recall strategies.

When player 1 plays a k-bounded-recall strategy, we denote by m_n his memory at stage n, that is, the k-history composed of the k pairs of actions played in stages $n-k, n-k+1, \ldots, n-1$. As player 2 observes past play, he knows m_n.

The next lemma asserts that for every fixed bounded-recall strategy of player 1 there is a reply of player 2 that ensures the long-run average payoff remains far from F.

Lemma 3.1 *Let F be a closed set that does not contain any convex approachable set. Then there is $\varepsilon > 0$, and for every $k \in I\!N$ and every $\sigma : (I \times J)^k \to \Delta(I)$ there is a strategy τ of player 2 such that*

$$\forall \eta > 0, \exists N \in I\!N, \forall m \in (I \times J)^k, \qquad \mathbf{P}_{(m,\sigma),\tau}\left(\inf_{n\geq N} d(\bar{x}_n, F) < \varepsilon\right) < \eta.$$

Proof. Fix $k \in I\!N$ and $\sigma : (I \times J)^k \to \Delta(I)$.

Step 1: Irreducible sets in the space of memories. For every two memories $m, m' \in (I \times J)^k$, *m leads to m'* if and only if by playing properly, player 2 can make the game move in a single stage from memory m to memory m' with positive probability. Formally, this happens if and only if $m = (i_1, j_1, \ldots, i_k, j_k)$, $m' = (i_2, j_2, \ldots, i_k, j_k, i', j')$, and $\sigma(i' \mid m) > 0$.

Let $\mathcal{M}$ be the collection of all irreducible sets w.r.t. the "lead to" relation. That is, for every $M \in \mathcal{M}$, every memory $m \in M$ leads only to memories in M, and for every $m, m' \in M$ there is a sequence of memories $m = m_1, m_2, \ldots, m_L = m'$ such that m_l leads to m_{l+1} for each $l = 1, 2, \ldots, L-1$.

When we say that $m \in \mathcal{M}$, we mean that $m \in M$, for some $M \in \mathcal{M}$.

For every irreducible set M we fix an element $m_M \in M$. Since each $M \in \mathcal{M}$ is irreducible, there is a strategy τ_* and a positive integer N_1^*, such that

$$\forall M \in \mathcal{M}, \forall m \in M, \quad \mathbf{P}_{(m,\sigma),\tau_*}\left(m_n = m_M \text{ for some } n \leq N_1^*\right) > 1 - \delta.$$

That is, player 2 can ensure that with high probability the play moves to the memory m_M in a bounded number of stages, provided it starts in the irreducible set M.

Step 2: Irreducible sets and Approachability

Fix $m \in \mathcal{M}$. By Lehrer and Solan (2003, Proposition 3), there is no bounded-recall strategy that approaches F. Therefore, there is a strategy τ_m of player 2 and $\widehat{\delta}_m > 0$, such that on a set of play paths whose probability is at least $\widehat{\delta}_m$ the average payoff is infinitely often far from F by more than $\widehat{\delta}_m$. Since the number of possible memories is finite, $\widehat{\delta} := \frac{1}{8} \min_{m \in \cup_{M \in \mathcal{M}} M} \widehat{\delta}_m > 0$. Thus, we conclude that there is $\delta > 0$, independent of m, and an increasing sequence of integers $(N(k))_{k \in I\!N}$ such that for every $m \in \mathcal{M}$,

$$\mathbf{P}_{(m,\sigma),\tau_m}\left(\sup_{n \geq N(k)} d(\bar{x}_n, F) > 7\delta \quad \forall k \in I\!N\right) > \delta.$$

Since payoffs are bounded, for every $m \in \mathcal{M}$ there is $c_m \in I\!R^d$ satisfying $d(c_m, F) \geq 7\delta$, and two increasing sequences of integers $(N_1(k), N_2(k))_{k \in I\!N}$, such that

$$\mathbf{P}_{(m,\sigma),\tau}\left(\inf_{N_1(k) \leq n < N_2(k)} d(\bar{x}_n, c_m) < \delta\right) > \delta.$$

Denote $c_M := c_{m_M}$. We now claim that there is a strategy $\widehat{\tau}$ of player 2, and a stopping time ν, such that

$$\forall M \in \mathcal{M}, \forall m \in M, \quad \mathbf{P}_{(m,\sigma),\widehat{\tau}}(d(\bar{x}_\nu, c_M) < 2\delta) = 1. \tag{9}$$

Indeed, $\widehat{\tau}$ starts by playing randomly, until some irreducible set $M \in \mathcal{M}$ is reached (that is, until the first stage n that satisfies $m_n \in \mathcal{M}$.) τ then plays in blocks of random size. Let B_k be the first stage of block k, so that m_{B_k} is the memory of player 1 at that stage. At the beginning of block k, τ forgets past play, and follows $\tau_{m_{B_k}}$ for $\frac{B_k}{\delta}$ stages. It then continues to follow $\tau_{m_{B_k}}$, until the first stage n in which average payoff during the block is in $B(c_m, \delta)$, or until the length of the block is $N_2(l)$, where l is the minimal integer satisfying $N_1(l) \geq \frac{B_k}{\delta}$, whichever comes first.

From Eq. (9) it follows that there is $N_2^* \in I\!N$ such that

$$\forall M \in \mathcal{M}, \forall m \in M, \quad \mathbf{P}_{(m,\sigma),\widehat{\tau}}(d(\bar{x}_{\min\{\nu, N_2^*\}}, c_M) < 2\delta) > 1 - \delta.$$

Step 3: Constructing the excluding strategy. We now define a strategy τ, which guarantees that the long-run average payoff remains away from F.

τ starts by playing all actions with equal probability, until the memory of player 1 is in some irreducible set M.

From that stage on, τ plays in blocks of varying length. In each block, τ first follows the strategy τ^*, until the memory of player 1 is m_M, or for N_1^* stages, whichever comes first. If the memory of player 1 is not m_M, the block terminates. If the memory of player 1 is m_M, τ completely forgets past play; it supposes that the virtual memory of player 1 is m_M, and follows $\widehat{\tau}$ for $\min\{\nu, N_2^*\}$ stages.

Step 4: Average payoff remains out of F. We now argue that the average payoff under $((m,\sigma),\tau)$ remains far from F, for every initial memory m.

There is N_0 such that with high probability, the memory m_{N_0} is in some irreducible set:

$$\forall m \in (I \times J)^k, \qquad \mathbf{P}_{(m,\sigma),\tau}\,(m_{N_0} \text{ is in some irreducible set}) \geq 1 - \delta.$$

Denote by M the irreducible set the process was absorbed to.

The length of all blocks is bounded by $N_1^* + N_2^*$, and the probability that the average payoff along each block is in $B(c_M, 2\delta)$ is at least $1 - 2\delta$. Since payoffs are bounded by 1, conditioned that the irreducible set M is reached at stage N_0, the expectation of the distance between the average payoff along each block and c_M is at most 4δ. The lemma follows by the strong law of large numbers. ■

Remark 3.2 *Consider the proof of Lemma 2. Suppose that $\sigma' \in (I \times J)^k$ satisfies $\left|1 - \frac{\sigma(i|m)}{\sigma'(i|m)}\right| < \frac{\delta}{N_1^*+N_2^*}$. Then the strategy τ we constructed in the proof of Lemma 2 for σ is good against σ' as well. Indeed, in each block, and for every memory m, the L_1-difference between the probability distribution over the plays in the block induced by $((m,\sigma),\tau)$ and that induced by $((m,\sigma'),\tau)$ is at most δ. Hence, in at least $1-\delta$ of* *the blocks the average payoff is in $B(c_M, 5\delta)$, where M is the irreducible set that the process is absorbed to. Therefore, by the law of large numbers, the long-run average payoff is in $B(c_M, 6\delta)$, which is disjoint from F.*

We are now ready to prove Proposition 3. We will use Theorem 1 in Lehrer and Solan (2003), which states that a closed set is approachable by bounded-recall strategies if and only if it contains a convex approachable set.

Proof of Proposition 3: Let F be a closed set that does not contain the convex hull of any approachable set. Let $\delta > 0$ such that no bounded-recall strategy approaches $B(F, \delta)$.

Remark 1 implies that there is a *countable* collection of strategies $(\tau_l)_{l \in \mathbb{N}}$, and for every $k > 0$ a countable collection of positive integers $(N_k(l))_{l \in \mathbb{N}}$, such that for each bounded-recall strategy σ (without the initial memory) there is an index $l(\sigma) \in \mathbb{N}$ satisfying

$$\mathbf{P}_{(m,\sigma),\tau_{l(\sigma)}} \left(\inf_{n \geq N_k(l(\sigma))} d(\bar{x}_n, F) < \varepsilon \right) < 1/k, \qquad \forall m \in (I \times J)^k.$$

Since for every l, τ_l excludes some bounded-recall strategies from F, and the probability by which that happens depends on the number of stages τ_l is followed, and since for every bounded-recall strategy there is l such that τ_l excludes it from F, all we have to do to exclude F against *all* bounded-recall strategies is to construct a strategy that enumerates over all strategies (τ_l) in a proper fashion.

We now construct the excluding strategy τ. Let $r = (r_1, r_2) : \mathbb{N} \to \mathbb{N}^2$ be a 1-1 and onto function. τ plays in blocks of varying (possibly infinite) size. In block n, τ follows $\tau_{r_2(n)}$. The block terminates once the following two conditions are simultaneously satisfied. (i) the length of block n is at least $N_{r_1(n)}(r_2(n))$, and (ii) the average payoff within this block is in $B(F, \varepsilon)$.

It is easy to verify that τ excludes F, provided player 1 uses bounded-recall strategies. Indeed, fix a bounded-recall strategy (m, σ) and $\eta > 0$, and let $k > 1/\eta$. Let n be the unique integer such that $r(n) = (k, l(\sigma))$. If play never reaches block n, then average payoff remains bounded away from F. If play reaches block n, there is probability at least $1 - \eta$ that block n never terminates, and payoff within this block remains bounded away from F. Therefore, there is N sufficiently large such that $\mathbf{P}_{(m,\sigma),\tau}(\inf_{n \geq N} d(\bar{x}_n, F) < \varepsilon) < \eta$, as desired.

References

[1] Aumann, R.J. and S. Sorin (1989) Cooperation and Bounded Recall. *Games Econ. Behavior* **1**, 5-39.

[2] Blackwell, D. (1956) An Analog of the Minimax Theorem for vector payoffs. *Pacific J. Math.* **6**, 1-8.

[3] Feller, W. (1968) *An Introduction to Probability Theory and its Applications.* John Wiley and Sons.

[4] Hou T.F. (1969) Weak Approachability in a Two-Person Game. *Ann. Math. Stat.* **40**, 789-813

[5] Hou T.F. (1971) Approachability in a Two-Person Game. *Ann. Math. Stat.* **42**, 735-744

[6] Kalai, E. (1990) Bounded Rationality and Strategic Complexity in Repeated Games. In *Game Theory and Applications* (Columbus, OH, 1987), 131-157, Econom. Theory Econometrics Math. Econom., Academic Press, San Diego, CA.

[7] Lehrer, E. (1988) Repeated Games with Stationary Bounded Recall Strategies. *J. Econ. Th.*, **46**, 130-144.

[8] Lehrer, E. and E. Solan (2003) No-Regret with Bounded Computational Capacity. CMS-EMS Discussion Paper #1373, Kellogg School of Management, Northwestern University.

[9] Neyman, A. (1985) Bounded Complexity Justifies Cooperation in Finitely Repeated Prisoner's Dilemma. *Econ. Let.*, **9**, 227-229.

[10] Rockafellar, R. T. (1970) *Convex analysis.* Princeton landmarks in mathematics.

[11] Rubinstein, A. (1986) Finite Automata Play the Repeated Prisoner's Dilemma. *J. Econ. Th.*, **39**, 83-96.

[12] Seneta, E. (1981) Non-Negative Matrices and Markov Chains. Springer-Verlag.

[13] Shiryayev, A.N. (1984) *Probability.* Springer-Verlag New York Inc.

[14] Spinat, X. (2002) A Necessary and Sufficient Condition for Approachability. *Math. Oper. Res.* **27**, 31-44.

[15] Vieille, N. (1992) Weak Approachability. *Math. Oper. Res.*, **17**, 781-791.