

RECURSIVE REPEATED GAMES WITH ABSORBING STATES

J. FLESCH, F. THUIJSMAN, AND O. J. VRIEZE

We show the existence of stationary limiting average ε -equilibria ($\varepsilon > 0$) for two-person recursive repeated games with absorbing states. These are stochastic games where all states but one are absorbing, and in the nonabsorbing state all payoffs are equal to zero. A state is called absorbing if the probability of a transition to any other state is zero for all available pairs of actions. For the purpose of our proof, we introduce properness for stationary strategy pairs. Our result is sharp since it extends neither to the case with more nonabsorbing states, nor to the n -person case with $n > 2$. Moreover, it is well known that the result cannot be strengthened to the existence of 0-equilibria and that repeated games with absorbing states generally do not admit stationary ε -equilibria.

1. Introduction. A recursive repeated game with absorbing states is a special kind of stochastic game with finite state and action spaces, and with $\mathbb{N} := \{1, 2, \dots\}$ as the set of stages. A two-person stochastic game can be described by a state space $S := \{1, \dots, z\}$, and a corresponding collection $\{M_1, \dots, M_z\}$ of bimatrices, where entry (i, j) of M_s consists of $r^1(s, i, j), r^2(s, i, j) \in \mathbb{R}$ and a probability vector $(p(1|s, i, j), \dots, p(z|s, i, j))$. The stochastic game is to be played in the following way. At each stage $n \in \mathbb{N}$ the play is in precisely one of the states. If the play is in state s at stage n then, simultaneously and independently, both players are to choose an action: player 1 chooses a row i of M_s , while player 2 chooses a column j of M_s . These choices induce an immediate payoff $r^1(s, i, j)$ to player 1 and $r^2(s, i, j)$ to player 2. Next, the play moves with probability $p(t|s, i, j)$ to state t , where new actions are to be chosen at stage $n + 1$.

The players are assumed to have complete information and perfect recall. A player's strategy is a specification of a probability distribution, at each stage and state, over the available actions, conditional on the history of the play up to that stage. Strategies are generally denoted by π for player 1 and σ for player 2. A strategy is called stationary if, for each state, it specifies a mixed action to be used whenever this state is being visited. Stationary strategies are denoted by x and y . A stationary strategy is called pure, if for each state, it specifies one action to be chosen. A pair of strategies (π, σ) with an initial state $s \in S$ determines a stochastic process on the payoffs. The sequences of payoffs are evaluated by the limiting average reward, given for player $k \in \{1, 2\}$ by

$$\gamma^k(s, \pi, \sigma) = E_{s\pi\sigma} \left(\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{n=1}^T R_n^k \right),$$

where R_n^k are random variables for the payoffs of player k at stage $n \in \mathbb{N}$.

Received April 14, 1995; revised October 31, 1995 and March 10, 1996.

AMS 1991 subject classification. Primary. 90D15.

OR/MS Index 1978 subject classification. Primary: Games/Stochastic.

Key words. Stochastic games, limiting average rewards, Nash-equilibria, stationary strategies.

A pair of strategies (π, σ) is a limiting average ε -equilibrium ($\varepsilon \geq 0$), if neither player can gain more than ε by unilateral deviation, i.e., if for all $s, \tilde{\pi}, \tilde{\sigma}$:

$$\gamma^1(s, \pi, \sigma) \geq \gamma^1(s, \tilde{\pi}, \sigma) - \varepsilon \quad \text{and} \quad \gamma^2(s, \pi, \sigma) \geq \gamma^2(s, \pi, \tilde{\sigma}) - \varepsilon.$$

The question of existence of ε -equilibria is the most challenging open problem in stochastic game theory these days.

Now we turn to recursive repeated games with absorbing states, a special class of stochastic games. A state is called absorbing, if the probability of ever leaving this state is zero for all pairs of actions. If all payoffs in the nonabsorbing states equal zero, then the stochastic game is called recursive. A repeated game with absorbing states is a stochastic game with only one nonabsorbing state. Thus, a recursive repeated game with absorbing states is a repeated game with absorbing states where all payoffs in the nonabsorbing state equal 0.

Although repeated games with absorbing states always have limiting average ε -equilibria (cf. Vrieze and Thuijsman 1989), it is well known that they need not have stationary ones (cf. Sorin 1986). Recently, Evangelista et al. (1994) showed the existence of stationary limiting average ε -equilibria for ARAT repeated games with absorbing states. ARAT stands for additive decomposability of rewards and transitions in player controlled parts. Here we show that recursive repeated games with absorbing states also have stationary limiting average ε -equilibria. It is well known that recursive repeated games with absorbing states need not have stationary limiting average 0-equilibria (see Example 1 in §2). In §4 we provide an example showing that, contrary to the zero-sum case (cf. Everett 1957, Thuijsman and Vrieze 1992), non-zero-sum recursive games may not have a solution in stationary strategies. Although that game does have 0-equilibria, the general existence of limiting average ε -equilibria is not even known yet for recursive games.

In §2 we give the formal definitions and notations and we derive some preliminary results. For recursive repeated games with absorbing states we introduce proper strategy pairs and show their existence. In §3 of this paper we show the existence of stationary limiting average ε -equilibria. Several examples are provided to give the intuition behind the steps of the proof. Section 4 concludes with some examples showing that neither n -person recursive repeated games with absorbing states with $n > 2$, nor two-person recursive games, necessarily have stationary limiting average ε -equilibria.

2. Preliminaries. We first introduce some necessary notations. We use the notation Γ for a two-person recursive repeated game with absorbing states. Without loss of generality we suppose the absorbing states of Γ to be of size 1×1 and the nonabsorbing state to be of size $m \times n$. Thus there is only one nontrivial state with action spaces $I := \{1, \dots, m\}$ and $J := \{1, \dots, n\}$. Therefore the stationary strategy spaces X and Y have the form:

$$X = \left\{ x = (x(i))_{i \in I} \left| \sum_{i \in I} x(i) = 1, x(i) \geq 0 \forall i \in I \right. \right\},$$

$$Y = \left\{ y = (y(j))_{j \in J} \left| \sum_{j \in J} y(j) = 1, y(j) \geq 0 \forall j \in J \right. \right\}.$$

The initial state is the nonabsorbing one, and the associated bimatrix will be denoted by M . If entry (i, j) of M is chosen, then with probability p_{ij}^* a transition occurs to an

absorbing state where the payoff is a_{ij}^k to player k , and with probability $1 - p_{ij}^*$ the play stays in the initial state. In the initial state all immediate payoffs are equal to 0. For completeness we define $a_{ij}^k := 0$ if $p_{ij}^* = 0$.

DEFINITION 2.1. Let x and y be arbitrary stationary strategies. We introduce

$$p_{xy}^* := \sum_{i \in I} \sum_{j \in J} x(i) p_{ij}^* y(j), \quad \text{Car}(x) := \{i \in I | x(i) > 0\},$$

$$T^1(y) := \{i \in I | p_{iy}^* > 0\}, \quad B^1(y) := \{i \in I | \gamma^1(i, y) \geq \gamma^1(\pi, y) \forall \pi\}.$$

Sets $\text{Car}(y)$, $T^2(x)$ and $B^2(x)$ are analogously defined. The sets $\text{Car}(x)$ and $\text{Car}(y)$ are called the carriers of the strategies x and y . For the strategy pair (x, y) we have that p_{xy}^* is the one step absorption probability. If $p_{xy}^* = 0$, then we say that (x, y) is recurrent; otherwise we say (x, y) is absorbing. Now $T^1(y)$ consists of the pure strategies that are absorbing against y , $B^1(y)$ is the set of pure limiting average best replies against y , and similar interpretations apply for $T^2(x)$ and $B^2(x)$. As is well known, $B^1(y)$ and $B^2(x)$ are always nonempty.

Basic assumptions. Each time that we write about a limit for δ to 0, we have a discrete sequence in mind, which can be assumed to converge by compactness arguments and by finiteness of the state and action spaces. Similarly, whenever we consider strategies depending on $\delta \in (0, 1)$, we implicitly assume that the carriers of these strategies are independent of δ and δ is close to 0. Also, we identify each pure stationary strategy with the corresponding action.

LEMMA 2.2. Let $(x, y) \in X \times Y$. Then $\gamma^1(x, y) = 0$ if $p_{xy}^* = 0$, and otherwise

$$\gamma^1(x, y) = \frac{\sum_{i \in I} \sum_{j \in J} x(i) p_{ij}^* y(j) a_{ij}^1}{\sum_{i \in I} \sum_{j \in J} x(i) p_{ij}^* y(j)} = \frac{\sum_{i \in I} x(i) p_{iy}^* \gamma^1(i, y)}{\sum_{i \in I} x(i) p_{iy}^*}.$$

The proof of this lemma is straightforward. For our main result we introduce proper and δ -proper strategy pairs ($\delta \in (0, 1)$) for two-person recursive repeated games with absorbing states.

DEFINITION 2.3. A pair of strategies $(x_\delta, y_\delta) \in X \times Y$ is δ -proper for $\delta > 0$, if

(1) (x_δ, y_δ) is completely mixed, i.e., $\text{Car}(x_\delta) = I$ and $\text{Car}(y_\delta) = J$,

(2) $\gamma^1(i, y_\delta) > \gamma^1(e, y_\delta)$ implies $\delta \cdot x_\delta(i) \geq x_\delta(e)$ for all $i, e \in I$,

(3) $\gamma^2(x_\delta, j) > \gamma^2(x_\delta, f)$ implies $\delta \cdot y_\delta(j) \geq y_\delta(f)$ for all $j, f \in J$.

A pair of strategies (x, y) is called proper if $(x, y) = \lim_{\delta \downarrow 0} (x_\delta, y_\delta)$ for some sequence of δ -proper strategy pairs (x_δ, y_δ) .

Proper and δ -proper strategy pairs are defined analogously to proper and δ -proper equilibria of games in normal form, but here they are not always ε -equilibria, not even for zero-sum games, as shown in the following example.

EXAMPLE 1. In this game entry (1, 1) is nonabsorbing and all other entries are absorbing with probability 1 (indicated by *) giving the corresponding absorption

0, 0	2, -2
1, -1	0, 0
	*
*	*

payoffs to players 1 and 2 respectively. Here $((1 - \delta^2, \delta^2), (1 - \delta^2, \delta^2))$ is δ -proper for small $\delta > 0$, so $((1, 0), (1, 0))$ is proper, but neither one is an ε -equilibrium.

THEOREM 2.4. *There exists a proper strategy pair in Γ .*

PROOF. By the compactness of X and Y , it suffices to show that, for δ sufficiently small, there exists a δ -proper pair in Γ . Let $\delta \in (0, 1)$ and let $X(\delta) := \{x \in X | x(i) \geq \delta^m \ \forall i \in I\}$ and $Y(\delta) := \{y \in Y | y(j) \geq \delta^n \ \forall j \in J\}$. Consider the following correspondence Ψ from $X(\delta) \times Y(\delta)$ into the set of all subsets of $X(\delta) \times Y(\delta)$: $\Psi(x, y) = (E_y, F_x)$, where $E_y := \{x \in X(\delta) | \gamma^1(i, y) > \gamma^1(e, y) \text{ implies } \delta \cdot x(i) \geq x(e) \ \forall i, e \in I\}$, and $F_x := \{y \in Y(\delta) | \gamma^2(x, j) > \gamma^2(x, f) \text{ implies } \delta \cdot y(j) \geq y(f) \ \forall j, f \in J\}$. Now Ψ has a fixed point, since all conditions of Kakutani's fixed point theorem (cf. Kakutani 1941) are satisfied. Because every fixed point is a δ -proper pair, the proof is complete. \square

3. Stationary ε -equilibria. In this section we shall prove our main result:

THEOREM 3.1. *In every recursive game with absorbing states Γ , there exists a stationary ε -equilibrium $(x_\varepsilon, y_\varepsilon)$ for all $\varepsilon > 0$.*

In order to prove this theorem we first give a sufficient condition for a strategy x_δ to be an ε -best reply against some y_δ .

LEMMA 3.2. *For all $\delta \in (0, 1)$, let x_δ and y_δ be arbitrary stationary strategies in the game Γ and let $\tilde{y} := \lim_{\delta \downarrow 0} y_\delta$. Suppose $\exists i^* \in \text{Car}(x_\delta) \cap B^1(y_\delta) \cap T^1(\tilde{y})$. If*

$$\lim_{\delta \downarrow 0} \frac{x_\delta(e)}{x_\delta(i^*)} = 0 \quad \forall e \in T^1(y_\delta) \setminus B^1(y_\delta),$$

then $\gamma^1(x_\delta, y_\delta) \geq \gamma^1(i^, y_\delta) - \varepsilon$ for every $\varepsilon > 0$ and sufficiently small δ .*

PROOF. Let $\varepsilon > 0$. Lemma 2.2 yields that for sufficiently small δ we have

$$\begin{aligned} & \gamma^1(x_\delta, y_\delta) \\ &= \frac{\sum_{e \in B^1(y_\delta)} x_\delta(e) p_{ey_\delta}^* \gamma^1(e, y_\delta) + \sum_{e \notin B^1(y_\delta)} x_\delta(e) p_{ey_\delta}^* \gamma^1(e, y_\delta)}{\sum_{e \in B^1(y_\delta)} x_\delta(e) p_{ey_\delta}^* + \sum_{e \notin B^1(y_\delta)} x_\delta(e) p_{ey_\delta}^*} \\ &= \frac{\sum_{e \in B^1(y_\delta)} \frac{x_\delta(e)}{x_\delta(i^*)} p_{ey_\delta}^* \gamma^1(e, y_\delta) + \sum_{e \notin B^1(y_\delta)} \frac{x_\delta(e)}{x_\delta(i^*)} p_{ey_\delta}^* \gamma^1(e, y_\delta)}{\sum_{e \in B^1(y_\delta)} \frac{x_\delta(e)}{x_\delta(i^*)} p_{ey_\delta}^* + \sum_{e \notin B^1(y_\delta)} \frac{x_\delta(e)}{x_\delta(i^*)} p_{ey_\delta}^*} \\ &\geq \frac{\sum_{e \in B^1(y_\delta)} \frac{x_\delta(e)}{x_\delta(i^*)} p_{ey_\delta}^* \gamma^1(e, y_\delta)}{\sum_{e \in B^1(y_\delta)} \frac{x_\delta(e)}{x_\delta(i^*)} p_{ey_\delta}^*} - \varepsilon = \gamma^1(i^*, y_\delta) - \varepsilon. \quad \square \end{aligned}$$

PROOF OF THEOREM 3.1. Let (\tilde{x}, \tilde{y}) be a proper pair, where $(\tilde{x}, \tilde{y}) = \lim_{\delta \downarrow 0} (x_\delta, y_\delta)$ for some sequence of δ -proper pairs. We show the following (illustrated in Example 2 below):

- (1) If (\tilde{x}, \tilde{y}) is absorbing, then (x_δ, y_δ) is an ε -equilibrium for small δ .
- (2) If (\tilde{x}, \tilde{y}) is recurrent, then either (\tilde{x}, \tilde{y}) or (x_δ, \tilde{y}) or (\tilde{x}, y_δ) is an ε -equilibrium for small δ .

PART (1). Since (\tilde{x}, \tilde{y}) is absorbing, the δ -properness of (x_δ, y_δ) implies the existence of $i^* \in B^1(y_\delta) \cap T^1(\tilde{y})$. We show that x_δ is an ε -best reply against y_δ for sufficiently small δ . Suppose $T^1(y_\delta) \setminus B^1(y_\delta) \neq \emptyset$, otherwise it is obvious. By the δ -properness of (x_δ, y_δ) ,

$$\lim_{\delta \downarrow 0} \frac{x_\delta(e)}{x_\delta(i^*)} = 0 \quad \forall e \in T^1(y_\delta) \setminus B^1(y_\delta),$$

so the conditions of Lemma 3.2 are fulfilled and therefore x_δ is an ε -best reply against y_δ for sufficiently small δ .

PART (2). If (\tilde{x}, \tilde{y}) is an ε -equilibrium, then we are done. Otherwise, at least one of the players has a profitable deviation with respect to (\tilde{x}, \tilde{y}) . Without loss of generality suppose that $i^* \in B^1(\tilde{y})$ is a profitable deviation of player 1 against \tilde{y} . Then, since (\tilde{x}, \tilde{y}) is recurrent, we must have $i^* \in T^1(\tilde{y})$. We show that (x_δ, \tilde{y}) is an ε -equilibrium for sufficiently small δ .

Since $i^* \in T^1(\tilde{y})$ we also have $i^* \in T^1(y_\delta)$ for all δ . Let $e \in T^1(\tilde{y}) \setminus B^1(\tilde{y})$. Then (cf. Lemma 2.2),

$$\lim_{\delta \downarrow 0} \gamma^1(i^*, y_\delta) = \gamma^1(i^*, \tilde{y}) > \gamma^1(e, \tilde{y}) = \lim_{\delta \downarrow 0} \gamma^1(e, y_\delta)$$

yields that $\gamma^1(i^*, y_\delta) > \gamma^1(e, y_\delta)$ for sufficiently small δ . Hence, by the δ -properness of (x_δ, y_δ) we have

$$\lim_{\delta \downarrow 0} \frac{x_\delta(e)}{x_\delta(i^*)} = 0$$

and, by Lemma 3.2 with (x_δ, \tilde{y}) instead of (x_δ, y_δ) , we obtain $\gamma^1(x_\delta, \tilde{y}) \geq \gamma^1(i^*, \tilde{y}) - \varepsilon$ for small δ . So x_δ is an ε -best reply against \tilde{y} .

On the other hand \tilde{y} is a best reply against x_δ since

$$\lim_{\delta \downarrow 0} \frac{y_\delta(j^*)}{y_\delta(j)} > 0 \quad \forall j^* \in \text{Car}(\tilde{y}), \forall j \in J$$

implies, by the δ -properness of (x_δ, y_δ) , that $\gamma^2(x_\delta, j^*) \geq \gamma^2(x_\delta, j)$ for all $j^* \in \text{Car}(\tilde{y})$ and all $j \in J$, which implies $\gamma^2(x_\delta, \tilde{y}) \geq \gamma^2(x_\delta, j)$ for all $j \in J$. \square

EXAMPLE 2. As before, entry $(1, 1)$ is nonabsorbing and absorption with probability 1 is indicated by *. Here the pair $(x_\delta, y_\delta) = ((1 - \delta^2 - \delta^4, \delta^4, \delta^2), (\delta^2, 1 - \delta^2))$

0,0	4,-3
3,-2	1,-4
1,-4	3,-2

is δ -proper for small δ ; hence $(\tilde{x}, \tilde{y}) = ((1, 0, 0), (0, 1))$ is proper. Then (\tilde{x}, \tilde{y}) is absorbing, and one can easily check that (x_δ, y_δ) is an ε -equilibrium for small δ . Note that (\tilde{x}, \tilde{y}) is not an ε -equilibrium in this game.

REMARK 1. Another way to establish equilibria having similar properties as δ -proper pairs is by defining for $\delta \in (0, 1)$ the following restricted strategy spaces:

$$\bar{X}(\delta) := \left\{ x \in X \left| \sum_{i \in U} x(i) \geq \delta^{m-|U|} \quad \forall \emptyset \neq U \subset I \right. \right\},$$

$$\bar{Y}(\delta) := \left\{ y \in Y \left| \sum_{j \in V} y(j) \geq \delta^{n-|V|} \quad \forall \emptyset \neq V \subset J \right. \right\},$$

and by defining linearized rewards

$$\bar{\gamma}^1(x, y) := \sum_{i \in I} x(i) \gamma^1(i, y) \quad \text{and} \quad \bar{\gamma}^2(x, y) := \sum_{j \in J} y(j) \gamma^2(x, j).$$

Using Kakutani's fixed point theorem, one can show the existence of stationary equilibria $(\bar{x}_\delta, \bar{y}_\delta)$ in $\bar{X}(\delta) \times \bar{Y}(\delta)$ with respect to the rewards $(\bar{\gamma}^1, \bar{\gamma}^2)$. Such equilibria have similar properties as δ -proper pairs, and the existence of limiting average ε -equilibria can be established analogously.

Notice that a stationary equilibrium (z_δ, w_δ) would also exist in $X(\delta) \times Y(\delta)$ with respect to the original rewards (γ^1, γ^2) , but for such an equilibrium $\gamma^1(i, w_\delta) > \gamma^1(e, w_\delta)$ would not necessarily imply $\delta \cdot z_\delta(i) \geq z_\delta(e)$. This causes a discontinuity in the best reply structures when approaching $X \times Y$ by $X(\delta) \times Y(\delta)$.

REMARK 2. Limiting average rewards are often defined as $\liminf_{T \rightarrow \infty} ((1/T) \sum_{n=1}^T E_{s\pi\sigma}(R_n^k))$. Since for stationary strategies the two definitions coincide, our results apply to this alternative definition as well.

REMARK 3. Observe that the ε -equilibria constructed are also ε -equilibria in the finitely repeated game, if the number of repetitions is large enough.

Acknowledgement. We are indebted to anonymous referees. Their remarks contributed to a better presentation of the results.

References

- Everett, H. (1957). Recursive games. In *Contributions to the Theory of Games III*, M. Dresher, A. W. Tucker and P. Wolfe, eds., *Ann. Math. Studies* **39**, Princeton University Press, Princeton, 47–78.
- Evangelista, E. S., T. E. S. Raghavan, O. J. Vrieze (1994). *ARAT repeated games with absorbing states*. Mimeo, University of Illinois at Chicago.
- Flesch, J., F. Thuijsman, O. J. Vrieze (1995). *Cyclic Markov equilibria in a cubic game*. Report M95-06, University of Limburg, Maastricht, The Netherlands.
- Kakutani, S. (1941). A generalization of Brouwer's fixed point theorem. *Duke Math. J.* **8** 161–167.
- Sorin, S. (1986). Asymptotic properties of a non-zero-sum stochastic game. *Internat. J. Game Theory* **15** 101–107.
- Thuijsman, F., O. J. Vrieze (1992). Note on recursive games. In *Game Theory and Economic Applications*, B. Dutta et al., eds., *Lecture Notes in Econ. and Math. Systems* **389**, Springer Verlag, Berlin, 133–145.
- Vrieze, O. J., F. Thuijsman (1989). On equilibria in repeated games with absorbing states. *Internat. J. Game Theory* **18** 293–310.

J. Flesch: Department of Mathematics, University of Limburg, P.O. Box 616, 6200 MD Maastricht, The Netherlands; e-mail: flesch@math.rulimburg.nl

F. Thuijsman: Department of Mathematics, University of Limburg, P.O. Box 616, 6200 MD Maastricht, The Netherlands; e-mail: frank@math.rulimburg.nl

O. J. Vrieze: Department of Mathematics, University of Limburg, P.O. Box 616, 6200 MD Maastricht, The Netherlands; e-mail: vrieze@math.rulimburg.nl