

Evolutionary Stable Collections in Endogenously Noisy Population of Automata*

Xavier Vilà[†]
Universitat Autònoma de Barcelona

September 23, 2005

Abstract

We examine the evolutionary stability of populations of automata, in the context of the Abreu and Rubinstein's automaton selection game, in the presence of endogenously determined noise. We show that, if any population contains a small fraction of people playing a "myopic" best response to their environment, then there is a unique collection of automata that is evolutionary stable. Furthermore, these automata are efficient, boundedly complex, and, in the case of the Repeated Prisoners' Dilemma, the only stable collection turns out to be the one composed of the well known *Tit-for-Tat* and *Grim Trigger* strategies.

Keywords: Game Theory. Evolutionary Stability. Finite Automata. Complexity Costs.

*Financial support from grants BEC2002-02130 of the *Dirección General de Investigación* (Spanish Ministry of Science and Technology), SGR2001-0164 of the *Direcció General de Recerca* (Generalitat of Catalonia), and CREA (Barcelona Economics Program) are gratefully acknowledged.

[†]Departament d'Economia i d'Història Econòmica. Universitat Autònoma de Barcelona. 08193-Bellaterra (Barcelona). Spain. <Xavier.Vila@uab.es>

1 Introduction

After Aumann [3] first suggested the use of finite automata to represent the strategies used in repeated games, several papers have further explored that possibility. Rubinstein [12] focuses on the Repeated Prisoners' Dilemma and incorporates, for the first time, complexity costs into the analysis. Abreu and Rubinstein [1] extend this framework to general two-person games and show, for the case of the Repeated Prisoners' Dilemma, that the "size" of the Folk Theorem is reduced so that *only* a countable set of payoff points can be achieved in equilibrium. Binmore and Samuelson [4] consider a modification of the ESS, or MESS, as the equilibrium concept. Their results show that evolutionary stable strategies in the conventional sense (ESS) often fail to exist. They argue that this non existence is due to the fact that for any potential equilibrium automaton, a successful entrant can be constructed by replicating the first automaton and changing those states not used when playing against itself. Under their modified concept (MESS), though, there exist equilibrium automata, which achieve the highest payoff possible in the game (utilitarian outcome). Specifically, in the case of the Repeated Prisoners' Dilemma, they show that automata in a polymorphous MESS attain the cooperative payoff. Nevertheless, they express their doubts about how stable this payoff would be if mutations were allowed to overlap, assumption that seems more realistic to a social context. *If mutations were frequent*, they say, *no grounds exist for supposing that utilitarian outcomes will survive*. In a related work Cooper [5] modifies the model by Binmore and Samuelson by assuming *finite costs of complexity* (embedded in the payoff function) as opposed to the lexicographic *costs of complexity* (that operate at a different level than the payoff function) and show that, in this case, a folk theorem type of result is recovered in place of the uniqueness result of Binmore and Samuelson. Volij [10] also analyzes the repeated Prisoners's Dilemma and shows that the only automata that satisfies the standard ESS requirements (*à la* Maynard Smith [9]) is the one that plays the "always defect" strategy.

Probst [11] also builds upon Binmore and Samuelson's work and proposes, for the Repeated Prisoners' Dilemma game, a "noisy" population in which there always exists a small group of stubborn one-state automata that play "always cooperate" or "always defect". Additionally, he proposes an alternative loosening of the ESS conditions to circumvent Binmore and Samuelson's argument on the reason for the non existence of ESS in the Abreu/Rubinstein automata selection game. Probst considers collection of automata all the elements of which are indistinguishable for the evolutionary process. Then, if every element in the collection is an ESS on its own, the collection is called an *Evolutionary Stable Collection* or ESC. Probst shows that there is a unique set of five three-state automata that satisfies these conditions for ESC. Each of these five automata starts off with defection but attains the cooperative payoff (they are utilitarian in the Binmore and Samuelson's sense) when playing against each other.

Our approach is based on Probst's work but considers a different type of "noisy" population. The idea behind our hypothesis is that, in a social context, mutations are seldom random. In our view, new behaviors arise for reasons that have, at least to a certain extent, a rational explanation (or that are not completely irrational) and that respond in some way to the the particular state of the environment in which they take place. In that sense, we assume that in any population there is always a small group of people that consistently plays a short run¹ best response to the action taken by the main part of the

¹... and hence myopic since such behaviors do not take into account the full horizon of the repeated game. For instance, in the context of firms' repeated price competition, such behavior would be that of a firm that

population. For instance, in the specific context of the Repeated Prisoners' Dilemma, our approach translates into assuming that in any population there is always a small fraction of people that plays "always defect" because that is the *short run myopic* best response to any behavior prevalent in the population. By doing so, we find that there is one and only one collection (composed of two automata) that satisfies the conditions for ESC in the Repeated Prisoners' Dilemma. Moreover, the automata in this collection turn out to be the ones that appear more often in the literature: *Tit-for-Tat* and *Grim Trigger*.

We further explore this variation of Probst's idea by extending the analysis to general symmetric games with two players. We prove that, under our modified solution concept, if a game has an evolutionary stable collection, then it is unique. Furthermore, all the automata in the collection are utilitarian in the Binmore and Samuelson's sense, that is, each attains the maximum payoff possible when playing against a replica of itself. We also consider the case of polymorphous populations and show that these two results (uniqueness and efficiency) also hold. Another set of results refer to the complexity of the equilibrium automata. We show that under our conditions for ESC, the number of strategies of the stage game imposes a uniform bound on the complexity of such equilibrium automata. An even more stringent bound (although not uniform) is determined by the number of different actions that are best response to the strategy induced by the equilibrium automata. This result allows for a straightforward computation of our stability concept in any 2-players symmetric game: just consider those automata that (first) play optimally against themselves (i.e. achieve the utilitarian payoff) and (second) play optimally against the *myopic* best respondents.

In section 2, we present the formal model and introduce some notation. Section 3 presents the stability conditions. Section 4 contains the main results. In section 5 we extend the analysis by allowing polymorphous populations. In sections 6 and 7 we discuss the special implications that our approach has in the Repeated Prisoners' Dilemma and in Games of common interests respectively. Section 8 concludes.

2 The Model

We consider repeated games based on 2-players symmetric stage games of the form $G = \langle S, U \rangle$ where $S = \{s_1, \dots, s_n\}$ is the space of actions for each player and $U_i : S \times S \rightarrow \mathfrak{R}$ is the (symmetric) payoff function. The repeated game $G^\infty = \langle F, \pi \rangle$ is constructed in the usual way based on the stage game G . The payoff functions π_1 and π_2 correspond to the "limit of the means":

$$\pi_i(r_1, r_2) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} U_i(r_1(h_t), r_2(h_t))$$

where h_t denotes the history of the game up to (and including) time t , and $r_i(h_t)$ is the action taken by player i at time $t + 1$

The automaton selection game of Abreu and Rubinstein [1] is $G^\# = \langle \mathcal{A}, \succ \rangle$. The strategy space is the set \mathcal{A} of finite automata. A finite automaton (with output) or "Moore machine"² is just a system that responds to discrete inputs with discrete outputs. Formally, an automaton $a \in \mathcal{A}$ is described by $a = \{Q^a, q_0^a, S, \delta^a, \lambda^a\}$. Q^a is the

behaves in a "take-the-money-and-run" fashion (i.e., set low prices, collect profits and exit the market)

²See Hopcroft and Ullman [7]

(finite) set of internal states for automaton a , q_0^a is its initial state, $\lambda^a : Q^a \rightarrow S$ maps each internal state to an action and $\delta^a : Q^a \times S \rightarrow Q^a$ is the transition function that assigns a new state to each internal state depending on the action taken by the opponent. Typically, automata are presented as directed graphs in which each node represents an internal state with the action attached to it and the directed paths represent the transitions. For instance, the two automata in Figure 1 represent the well known strategies “*Tit-for-Tat*” and “*Grim trigger*” in the Repeated Prisoners’ Dilemma. (Automata in this paper are always depicted with a curly arrow pointing to their initial state).




Figure 1: *Tit-for-Tat* and *Grim Trigger* automata

The preferences (\succ) over automata take into account not only the automata performance but also their complexity. In this sense, the complexity of an automaton $a \in \mathcal{A}$ is assumed to be its number of internal states and denoted by $|a|$.³ The preference order over automata a and a' (parametrized by the common opponent a'') satisfy:

$$a \succ_1^{a''} a' \iff \begin{cases} \{\pi_1(a, a'') > \pi_1(a', a'')\} & \text{or} \\ \{\pi_1(a, a'') = \pi_1(a', a'')\} & \text{and } |a| < |a'| \end{cases}$$

with a similar definition for $\succ_2^{a''}$.

The following fact is discussed in Abreu and Rubinstein [1] and, although it is obvious since we work with automata that have finitely many states, we want to call the attention on it because it will be of great importance for the rest of the paper.

Fact 2.1. *Let $a, a' \in \mathcal{A}$. Then, when playing against each other, they get into a finite cycle*

We will use $C^a(a') \subset Q^a(a')$ to denote such a cycle, where $Q^a(a') \subset Q^a$ is the set of internal states used by a when playing against a'

The key assumption of our model is that in any population there is always a small group of people that consistently plays a (myopic) best response to the strategy induced by the behavior of the main part of the population. In this sense, the (*stage game*) *strategy induced by a* is denoted by $\sigma(a) \in \Delta(S)$ and corresponds to the (possibly mixed) strategy for the stage game implicit in $C^a(a)$, that is:

$$\sigma(a, s_i) = \frac{\sum_{q \in C^a(a)} I_{s_i}(\lambda^a(q))}{|C^a(a)|} \quad s_i \in S$$

where

$$I_{s_i}(\lambda^a(q)) = \begin{cases} 1 & \text{if } \lambda^a(q) = s_i \\ 0 & \text{otherwise} \end{cases}$$

³See Kalai and Stanford [8] for some game theoretical properties of this measure of complexity.

and $\sigma(a) = (\sigma(a, s_1), \dots, \sigma(a, s_n))$. According to the expected utility hypothesis we define $u(s_i, \sigma(a)) = \sum_{j=1}^n \sigma(a, s_j) U(s_i, s_j)$. Now, we can define the “best response” to a population composed of replica of automaton a as:

$$\beta(a) = \{s_i \in S \mid u(s_i, \sigma(a)) \geq u(s_j, \sigma(a)) \quad \forall s_j \in S\}$$

Therefore, our main assumption is that in a population mainly formed by automata of type a , a small proportion (ϵ) of one state automata always playing some $s_i \in \beta(a)$ completes the environment. Formally, such a population will be denoted by $\mathcal{P}_\epsilon(a, \tilde{s}_i) = \epsilon \tilde{s}_i \oplus (1-\epsilon)a$, where $0 < \epsilon < \frac{1}{2}$. Thorough the paper, we will use \tilde{s}_i to represent the one state automaton that implements the strategy “always play s_i ” ($s_i \in S$). Accordingly, the expected payoff of an automaton a' when randomly matched against a population $\mathcal{P}_\epsilon(a, \tilde{s}_i)$ is

$$\pi(a', \mathcal{P}_\epsilon(a, \tilde{s}_i)) = \epsilon \pi(a', \tilde{s}_i) + (1-\epsilon) \pi(a', a).$$

3 The Stability Condition

Probst’s [11] *Payoff Indistinguishable Collection* (PIC) and *Evolutionary Stable Collection* (ESC) are defined for the specific case of the Repeated Prisoners’ Dilemma. In this section we generalize these definitions to general symmetric games with two players.

Definition 3.1. A set $A \subset \mathcal{A}$ is called Payoff Indistinguishable Collection or PIC if $\forall 0 < \epsilon < 1, \forall a, a' \in A, \forall s \in \beta(a)$ and $\forall s' \in \beta(a')$:

$$\begin{aligned} \pi(a, \mathcal{P}_\epsilon(a', \tilde{s}')) &= \pi(a, \mathcal{P}_\epsilon(a, \tilde{s})) = \pi(a', \mathcal{P}_\epsilon(a', \tilde{s}')) \\ &\text{and } |a| = |a'| \end{aligned} \tag{1}$$

An alternative and probably more intuitive characterization of a PIC is given in the following lemma.

Lemma 3.1. A set $A \subset \mathcal{A}$ is a PIC if and only if $\forall a, a' \in A$,

- (i) $\pi(a, a) = \pi(a, a') = \pi(a', a')$
- (ii) $\pi(a, \tilde{s}) = \pi(a, \tilde{s}') = \pi(a', \tilde{s}') \quad \forall s \in \beta(a), \quad \forall s' \in \beta(a')$
- (iii) $|a| = |a'|$

Proof. Necessity: Item (i) follows from a simple continuity argument letting ϵ go to zero in the definition of $\mathcal{P}_\epsilon(a, \tilde{s})$. Then, parts (ii) and (iii) follow immediately from (i) and (1)

Sufficiency: Simply multiply through by $(1-\epsilon)$ in (i) and by ϵ in (ii) and add them together to obtain (1) \square

The following definition takes Binmore and Samuelson’s [4] and Probst’s [11] modified versions of Maynard-Smith’s [9] ESS and adapts them to our framework.

Definition 3.2. Let A^* be a PIC. A^* is an Evolutionary Stable Collection, or ESC, if $\forall a \in A^*$ and $\forall a' \notin A^* \exists \bar{\epsilon}$ such that $\forall \epsilon \in (0, \bar{\epsilon})$:

$$\begin{aligned}
& \text{(i)} \quad \{\pi(a, \mathcal{P}_\epsilon(a, \tilde{s})) > \pi(a', \mathcal{P}_\epsilon(a, \tilde{s})) \quad \forall s \in \beta(a)\}, \\
\text{or} \quad & \text{(ii)} \quad \{\pi(a, \mathcal{P}_\epsilon(a, \tilde{s})) \geq \pi(a', \mathcal{P}_\epsilon(a, \tilde{s})) \quad \forall s \in \beta(a) \\
& \quad \text{and} \quad \pi(a, \mathcal{P}_\epsilon(a', \tilde{s}')) > \pi(a', \mathcal{P}_\epsilon(a', \tilde{s}')) \quad \forall s' \in \beta(a')\}, \\
\text{or} \quad & \text{(iii)} \quad \{\pi(a, \mathcal{P}_\epsilon(a, \tilde{s})) \geq \pi(a', \mathcal{P}_\epsilon(a, \tilde{s})) \quad \forall s \in \beta(a) \\
& \quad \text{and} \quad \pi(a, \mathcal{P}_\epsilon(a', \tilde{s}')) \geq \pi(a', \mathcal{P}_\epsilon(a', \tilde{s}')) \quad \forall s' \in \beta(a') \\
& \quad \text{and} \quad |a| < |a'|\}
\end{aligned} \tag{2}$$

Basically, this definition corresponds to the definition of ESS for automata given by Binmore and Samuelson [4] in which complexity has been taken into account. Then, it follows Probst [11] in the sense that potential invaders are restricted to come only from outside the PIC. Although the conditions given above look very similar to the ones given by Binmore and Samuelson, ESC is neither a stronger nor a weaker condition than MESS.

4 The Results

We present now our two main results. The first one states that any automaton in any ESC is utilitarian in the following sense:

Definition 4.1. An automaton $a \in \mathcal{A}$ is said to be utilitarian if

$$a \in \arg \max_{a \in \mathcal{A}} \pi(a, a)$$

That is, a is an utilitarian automaton if $\pi(a, a) = \pi^*$, where π^* is the maximum payoff achievable when an automaton plays against a replica of itself. Since the game is symmetric, an alternative characterization of π^* is $\pi^* = U(s^*, s^*)$, where $s^* \in \arg \max_{s \in S} U(s, s)$.

Our second result will establish that if an ESC exists, then it is unique.

4.1 Efficiency of ESC

For any automaton $a \in \mathcal{A}$, let $\Delta^a(q, \{s_t\}_{t=1}^n) \in Q^a$ denote the internal state of a reached following the sequence $\{s_t\}_{t=1}^n$ of actions by an opponent when a was initially at state q .

Lemma 4.1. Suppose $\pi(a, a) < \pi^*$ and select an state $q^c \in C^a(a)$. If there exists a sequence of actions $\{s_t\}_{t=1}^n$ such that

$$\begin{aligned}
& \text{(i)} \quad s_1 \neq \lambda^a(q^c) \\
& \text{(ii)} \quad \Delta^a(q^c, \{s_t\}_{t=1}^n) \in C^a(a)
\end{aligned}$$

then a cannot belong to an ESC.

Proof. Suppose that a sequence $\{s_t\}_{t=1}^n$ satisfies (i) and (ii) as above. Consider the automaton $a' \in \mathcal{A}$ as follows:

(i) In the moves before the cycle and in the first repetition of the cycle, a' behaves as follows

(i.i) If no deviation occurs, behave as a .

(i.ii) If a deviation occurs, switch to a *best reply* to that deviation and stay there.

(ii) In the second repetition of the cycle, a' plays $s_1 \neq \lambda^a(q^c)$. If matched, a' plays s^* forever. If the opponent plays $\lambda^a(q^c)$, a' plays s_2, \dots, s_n and then behaves as a .

Item (i) in the proof ensures that a' performs against any noisy player $\tilde{s} \in \beta(a)$ at least as well as a . Item (ii) defines a *secret hand-shaking* that a' uses to recognize replicas of itself and hence outperform a . If matched against a , a' behaves like a , if matched against a replica of itself achieves the utilitarian payoff π^* .

Clearly, this a' can enter a population formed by replicants of a . Hence, a cannot belong to an ESC. \square

Theorem 4.1. *Let $A^* \subset \mathcal{A}$ be an ESC. Then, $\forall a \in A^*$, $\pi(a, a) = \pi^*$ (a is utilitarian)*

Proof. Suppose the contrary, that is, suppose that $\exists a \in A^*$ such that $\pi(a, a) < \pi^*$. From here, we consider three different cases .

Case 1. $Q^a(a) = \{q\}$.

Consider in this case an automaton a' which behaves like a except for

$$\delta^{a'}(\delta^a(q, s'), \lambda^a(q)) = q \quad \text{for some } s' \neq \lambda^a(q)$$

Clearly, a' behaves like a against a and against any noisy player $\tilde{s} \in \beta(a)$. Hence, if $a \in A^*$, also $a' \in A^*$. But according to Lemma 4.1, a' cannot belong to A^* if we use the sequence $\{s_1, s_2\} = \{s', \lambda^a(q)\}$. Consequently, a cannot belong to an ESC.

Case 2. $S = \{s', s^*\}$.

Notice first that, since we are assuming that a is not utilitarian, $\lambda^a(q^c) = s'$ for some $q^c \in C^a(a)$. If $s' \in \beta(s^*)$, consider a' to be identical to a except for $\delta^{a'}(q^c, s^*) = q^c$. Clearly, a' performs as a against a and at least as well as a against any noisy player $\tilde{s} \in \beta(a)$. Hence, if $a \in A^*$, also $a' \in A^*$. A contradiction is again obtained by Lemma 4.1. If $s' \notin \beta(s^*)$, (that is, s^* is the only stage-game best response to s^*) then if $s^* \notin \beta(a)$ consider a' to be identical to a except for $\delta^{a'}(q^c, s^*) = \delta^a(q^c, \lambda^a(q^c))$. Again, a' behaves like a against a and against any noisy player. A contradiction is obtained as before. If, on the contrary, $s^* \in \beta(a)$, $a \in A^*$ implies that $\pi(a, \tilde{s}^*) = \pi^*$ and therefore $\pi(\tilde{s}^*, a) = \pi^* > \pi(a, a)$ contradicting, by definition of ESC that $a \in A^*$.

Case 3. $S = \{s', s'', s^*\}$.

Select $q^c \in C^a(a)$, $q' \in Q^a(a)$, and $s \in S$ such that $\lambda^a(q^c) \neq s$ and $\lambda^a(q') \neq s$. By Lemma 4.1 we can assume that $\delta^a(q^c, s) = \delta^a(q', s)$. Consider the automaton a' to be identical to a except for $\delta^{a'}(q^c, s) = \delta^a(q^c, \lambda^a(q^c))$. Notice again that a' performs as a against a and against any noisy player $\tilde{s} \in \beta(a)$. Hence, if $a \in A^*$, also $a' \in A^*$. A contradiction is again obtained by Lemma 4.1.

Hence, in any case, the assumption that a belongs to an ESC but is not utilitarian leads to a contradiction. \square

The result that follows is of less importance, but indicates that an automaton in an ESC also performs efficiently against any automaton that always plays a myopic best response.

Proposition 4.1. *Let $A^* \subset \mathcal{A}$ be an ESC. Then, $\forall a \in A^*$ and $\forall s \in \beta(a)$, $\pi(a, \tilde{s}) = \max_{a' \in \mathcal{A}} \pi(a', \tilde{s}) = \max_{s_j \in S} U(s_j, s)$.*

Proof. Suppose that the proposition is not true, that is, suppose that A^* is an ESC but $\exists a \in A^*$ and $\exists s \in \beta(a)$ such that $\pi(a, \tilde{s}) < \max_{s_j \in S} U(s_j, s)$. Suppose additionally that a 's initial action is to play s , and that it always plays s thereafter if its opponent keeps playing s . It is clear then, by theorem 4.1, that $s = s^*$ for otherwise a would not belong to an ESC. Hence, in this case, s^* can not be a best response to a . Hence, either a 's initial action is different from s or a eventually takes an action different from s when playing against \tilde{s} . In either case, let q be that particular state at which a 's action is not s . Consider then the automaton a' to be an exact replica of a except that it has an additional state, q^+ , and that:

$$(i) \delta^{a'}(q, s) = q^+$$

$$(ii) \lambda^{a'}(q^+) \in \beta(s)$$

$$(iii) \delta^{a'}(q^+, s) = q^+$$

Clearly, by the way a' is constructed, we have that $\pi(a', a) = \pi(a, a) = \pi(a', a')$ and $\pi(a', \tilde{s}) > \pi(a, \tilde{s})$. Hence, a' defeats a , which contradicts the assumption that A is an ESC. \square

4.2 Uniqueness of ESC

Our uniqueness result will refer to games for which there is only one action (s^*) that yields the utilitarian payoff. Prior to it, two lemmas and one definition are needed.

The following lemma says that when an equilibrium automaton meets a replica of itself, only this particular action will be used infinitely many times.

Lemma 4.2. *Let A^* be an ESC and suppose that there is a unique $s^* \in S$ such that $U(s^*, s^*) = \pi^*$. Then, $\forall a \in A^*$, $C^a(a) = \{q^*\}$ and $\lambda^a(q^*) = s^*$.*

Proof. Note first that $\lambda^a(q) = s^* \quad \forall q \in C^a(a)$ for otherwise $\pi(a, a) \neq \pi^*$ and thus, according to theorem 4.1, a could not belong to an ESC. Next, it must be the case that $C^a(a) = \{q^*\}$ and $\lambda^a(q^*) = s^*$ for if $C^a(a)$ contained more states, we could just "chop off" all those extra states in $C^a(a)$. The result would be an automaton that would behave exactly as a but less complex and, hence, would invade A^* . \square

The lemma that follows indicates that if two ESCs have some elements in common, then the two collections have to be the same.

Lemma 4.3. *Let A and A' be two ESC such that $A \cap A' \neq \emptyset$. Then, $A = A'$.*

Proof. The proof is simple. Suppose $\exists a' \in A'$ such that $a' \notin A$ and let $a \in A \cap A'$. Then, A can not be an ESC for a' invades A as it does "as well as" a . \square

Definition 4.2. Let $S^* \subseteq S$ the smallest subset of S satisfying $\forall s' \in \beta(\tilde{s}^*) (s' \neq s^*), \exists s \in S^*$ such that $s \in \beta(\tilde{s}')$

We can now state and prove our second result. The following theorem proves that if an ESC exists, it must be unique. It also shows that the complexity of an equilibrium automaton is determined by the fact that it must attain the utilitarian payoff when playing against a replica of itself and the maximum payoff possible when playing against any of the “myopic” players (as Theorem 4.1 and Proposition 4.1 establish). That is, an equilibrium automaton must have just as many internal states as needed to optimally play against a replica of itself and against the myopic players, but no more. Its complexity is bounded,

Theorem 4.2. *Let A^* be an ESC and suppose that there is a unique $s^* \in S$ such that $U(s^*, s^*) = \pi^*$. Then,*

- (i) $\forall a \in A^*, |a| = |S^*|$
- (ii) A^* is unique.

Proof. Notice first that, by Lemma 4.2, only s^* is played in the cycle.

Consider automaton a^* as follows:

- (i) $Q^{a^*} = \{q^*\} \cup \{q_s\}_{s \in S^*}$
- (ii) $q_0^{a^*} = q^*$
- (iii) $\lambda^{a^*}(q^*) = s^*$
- (iv) $\lambda^{a^*}(q_s) = s$
- (v) $\delta^{a^*}(q^*, s^*) = q^*$
- (vi) for $s \in S^*, \delta^{a^*}(q^*, s) = q_{s'}$ with $s' \in \beta(\tilde{s})$

We will show that if A^* is an ESC, then $a^* \in A^*$.

Suppose that $a \in A^*$. Consider the automaton a' identical to a except for $\forall q, \delta^{a'}(q, s^*) = \delta^{a'}(q, \lambda^a(q))$. It is clear that if $a \in A^*$, then also $a' \in A^*$ for if $s^* \in \beta(a)$ then a' is a best reply to \tilde{s}^* ((s^*, s^*) is a Nash Equilibrium). Suppose now that s^* is played k times in the pre-cycle phase when a' plays against a replicant of itself. Consider then automaton a'' having states $\{q_1, \dots, q_k\} \cup q^* \cup \{q_s\}_{s \in S^*}$ such that:

- (i) q_1 is its initial state
- (ii) $\lambda^{a''}(q_i) = s^*, i = 1, \dots, k$
- (iii) $\delta^{a''}(q_i, s^*) = q_{i+1}, i = 1, \dots, k-1$
- (iv) $\delta^{a''}(q_k, s^*) = q^*$
- (v) $\delta^{a''}(q_i, s) = q_i$ for $s \neq s^*$
- (vi) like a^* when q^* is reached.

Then, by Proposition 4.1, we must have that $|a''| \leq |a'|$. Also, a'' achieves π^* against itself and against a' , and replies optimally to any $\tilde{s} \in \beta(a')$ and $\beta(a'')$. Hence, it must be the case that $a'' \in A^*$. However, if $a'' \in A^*$, then also $a^* \in A^*$. This concludes the proof of the first part of the Theorem: any ESC must contain a^* .

The second part of the Theorem, uniqueness, follows by Lemma 4.3. \square

5 Polymorphous Populations

We consider now the case in which populations can be composed not only of one type of automaton (plus some best response to it), as it was the case before, but of a mixture of the automata in a PIC. For instance, in the case of the Repeated Prisoners' Dilemma, we have seen that a population composed of automata of the type *Tit-for-Tat* or a population composed of automata of the type *Grim trigger* are both stable in the sense that each can be "invaded" only by automata of the other type and that those automata are "indistinguishable" from a evolutionary point of view. Therefore, nothing prevents any of the automata in an ESC to drift in a population composed of replica of some other automaton of the same ESC. In this situation, a natural question to ask is: what happens then if a population is composed of a mixture of the automata in the ESC ? We will see in this section that the results obtained in the case of homogeneous populations hold in this more general case as well. For that, we need some additional notation.

Since we will be working with mixtures of automata, the set of reference will be the set of probability distributions over the set of finite automata denoted by $\Delta(\mathcal{A})$ and η will denote a typical element of this set. According to the expected utility hypothesis, we define

$$\pi(\eta, \mu) = \sum_{a_i \in \text{Supp}(\eta)} \sum_{a_j \in \text{Supp}(\mu)} \eta_{a_i} \mu_{a_j} \pi(a_i, a_j) \quad \forall \eta, \mu \in \Delta(\mathcal{A})$$

Also, as we did in section 2, we define

$$\sigma(a_i, a_j, s) = \frac{\sum_{q \in C^{a_i}(a_j)} I_s(\lambda^{a_i}(q))}{|C^{a_i}(a_j)|}$$

Consequently,

$$\sigma(\eta, s) = \sum_{a_i \in \text{Supp}(\eta)} \sum_{a_j \in \text{Supp}(\eta)} \eta_{a_i} \eta_{a_j} \sigma(a_i, a_j, s)$$

and $\sigma(\eta) = \{\sigma(\eta, s_1), \dots, \sigma(\eta, s_n)\}$. Therefore, the expected payoff of using action s_i in a population characterized by η is:

$$u(s_i, \sigma(\eta)) = \sum_{s_j \in S} \sigma(\eta, s_j) U(s_i, s_j)$$

Hence, in this case, the "best response" to a population composed of the mixture η is given by:

$$\beta(\eta) = \{s_i \in S \mid u(s_i, \sigma(\eta)) \geq u(s_j, \sigma(\eta)) \quad \forall s_j \in S\}$$

The lemma that follows relates the best response to a mixture with the best response to each automaton in the support of the mixture.

Lemma 5.1. $\bigcap_{a \in \text{Supp}(\eta)} \beta(a) \subset \beta(\eta) \subset \bigcup_{a \in \text{Supp}(\eta)} \beta(a)$

Proof. Clearly, if s is a best response to all the automata in the support of η , it is also a best response to a lineal combination of those automata because of the linear property of the expected utility hypothesis. Reciprocally, if s is a best response to η , there must exist an automaton in the support of η to which s is also a best response. \square

We present next the definition of a *Polymorphous Evolutionary Stable Collection* (PESC), that is meant to be the appropriate extension of the definition of ESC to include the possibility of a population composed of a mixture of the automata in the collection.

Definition 5.1. A PIC $A^* \subset \mathcal{A}$ is a Polymorphous Evolutionary Stable Collection or PESC if $\forall \eta \in \Delta(A^*), \forall a \in \text{Supp}(\eta)$ and $\forall b \notin A^*, \exists \bar{\epsilon}$ such that $\forall \epsilon \in (0, \bar{\epsilon}),$ ⁴

- (i) $\{\pi(a, \mathcal{P}_\epsilon(\eta, \tilde{s})) > \pi(b, \mathcal{P}_\epsilon(\eta, \tilde{s})) \quad \forall s \in \beta(\eta)\},$
or (ii) $\{\pi(a, \mathcal{P}_\epsilon(\eta, \tilde{s})) \geq \pi(b, \mathcal{P}_\epsilon(\eta, \tilde{s})) \quad \forall s \in \beta(\eta)$
and $\forall \mu \in \Delta(\mathcal{A})$ such that $\text{Supp}(\mu) = \text{Supp}(\eta) \cup \{b\}$
 $\exists a' \in \text{Supp}(\eta)$ such that $\forall s' \in \beta(\mu)$
 $\pi(a', \mathcal{P}_\epsilon(\mu, \tilde{s}')) > \pi(b, \mathcal{P}_\epsilon(\mu, \tilde{s}'))\},$
or (iii) $\{\pi(a, \mathcal{P}_\epsilon(\eta, \tilde{s})) \geq \pi(b, \mathcal{P}_\epsilon(\eta, \tilde{s})) \quad \forall s \in \beta(\eta)$
and $\pi(a', \mathcal{P}_\epsilon(\mu, \tilde{s}')) \geq \pi(b, \mathcal{P}_\epsilon(\mu, \tilde{s}'))$
 $\forall a' \in \text{Supp}(\eta), \forall \mu$ as in (ii), and $\forall s' \in \beta(b)$
and $\exists a' \in \text{Supp}(\eta)$ such that $|a'| < |b|\}.$

We will show that the results obtained in section 4, that is, efficiency and uniqueness of ESC, apply to PESC as well.

Proposition 5.1. *Let A^* be an ESC of $G^\#$ and let G have an strictly dominant strategy. Then, A^* is a PESC of $G^\#$.*

Proof. Let $\eta \in \Delta(A^*)$. Then, $\forall a_i \in \text{Supp}(\eta)$ we have that $a_i \in A^*$. Thus, since A^* is an ESC of $G^\#$, we have that $\forall b \notin A^* \pi(a_i, a_j) = \pi(a_j, a_j) \geq \pi(b, a_j) \quad \forall a_j \in \text{Supp}(\eta)$. Therefore,

$$\sum_{a_j \in \text{Supp}(\eta)} \eta_{a_j} \pi(a_i, a_j) \geq \sum_{a_j \in \text{Supp}(\eta)} \eta_{a_j} \pi(b, a_j)$$

and thus,

$$\pi(a_i, \eta) \geq \pi(b, \eta) \tag{3}$$

Let s^+ be the strategy that is strictly dominant in G . Clearly, $\beta(a) = \{s^+\} \quad \forall a \in \mathcal{A}$. Therefore, by proposition 4.1, $\pi(a_j, \tilde{s}^+) \geq \pi(b, \tilde{s}^+) \quad \forall a_j \in A^*$. In particular,

$$\pi(a_i, \tilde{s}^+) \geq \pi(b, \tilde{s}^+) \tag{4}$$

Therefore, (3) and (4) together imply that $\forall \epsilon > 0$

$$\pi(a_i, \mathcal{P}_\epsilon(\eta, \tilde{s}^+)) \geq \pi(b, \mathcal{P}_\epsilon(\eta, \tilde{s}^+))$$

If the above inequality always holds as a strict inequality, the proposition is proved. Suppose on the contrary that $\exists \eta \in \Delta(A^*), a_i \in \text{Supp}(\eta)$ such that

$$\pi(a_i, \mathcal{P}_\epsilon(\eta, \tilde{s}^+)) = \pi(b, \mathcal{P}_\epsilon(\eta, \tilde{s}^+))$$

⁴In this definition, the entrant automata b could in fact be a mixture of several automata not in A^* , and the size of ϵ needs not to be small for the results that come next to hold. Hence, although the results obtained can in fact be more general than the ones presented here, We prefer to keep this version for the shake of simplicity in the exposition.

or, equivalently,

$$\begin{aligned} & \epsilon\pi(a_i, \tilde{s}^+) + (1 - \epsilon) \sum_{a_h \in \text{Supp}(\eta)} \mu_{a_h} \pi(a_i, a_h) = \\ & = \epsilon\pi(b, \tilde{s}^+) + (1 - \epsilon) \sum_{a_h \in \text{Supp}(\eta)} \mu_{a_h} \pi(b, a_h) \end{aligned} \quad (5)$$

Now, since $\pi(a_i, \tilde{s}^+) \geq \pi(b, \tilde{s}^+)$ (because of Proposition 4.1) and $\pi(a_i, a_h) = \pi(a_h, a_h) \geq \pi(b, a_h)$ for otherwise A^* would not be an ESC, we have that (5) implies that

Therefore, by the definition of ESC, either

- (a) $\pi(a_h, \mathcal{P}_\epsilon(b, \tilde{s}^+)) > \pi(b, \mathcal{P}_\epsilon(b, \tilde{s}^+))$ or
- (b) $\pi(a_h, \mathcal{P}_\epsilon(b, \tilde{s}^+)) = \pi(b, \mathcal{P}_\epsilon(b, \tilde{s}^+))$ and $|a_h| < |b|$.

If (a) holds we have that

$$\epsilon\pi(a_h, \tilde{s}^+) + (1 - \epsilon)\pi(a_h, b) > \epsilon\pi(b, \tilde{s}^+) + (1 - \epsilon)\pi(b, b)$$

Consider then any $\mu \in \Delta(\mathcal{A})$ satisfying $\text{Supp}(\mu) = \text{Supp}(\eta) \cup \{b\}$. If (a) holds,

$$\mu_b \epsilon\pi(a_h, \tilde{s}^+) + \mu_b(1 - \epsilon)\pi(a_h, b) > \mu_b \epsilon\pi(b, \tilde{s}^+) + \mu_b(1 - \epsilon)\pi(b, b) \quad (6)$$

Also, because of Lemma 3.1 and because A^* is an ESC, $\pi(a_h, a_j) = \pi(a_j, a_j) \geq \pi(b, a_j)$ and $\pi(a_h, \tilde{s}^+) \geq \pi(b, \tilde{s}^+)$. Therefore, $\forall a_h \in \text{Supp}(\eta)$,

$$\mu_{a_j} \epsilon\pi(a_h, \tilde{s}^+) + \mu_{a_j}(1 - \epsilon)\pi(a_h, a_j) \geq \mu_{a_j} \epsilon\pi(b, \tilde{s}^+) + \mu_{a_j}(1 - \epsilon)\pi(b, a_j) \quad (7)$$

Adding (6) and (7) we get

$$\epsilon\pi(a_h, \tilde{s}^+) + (1 - \epsilon)\pi(a_h, \mu) > \epsilon\pi(b, \tilde{s}^+) + (1 - \epsilon)\pi(b, \mu)$$

or, equivalently,

$$\pi(a_h, \mathcal{P}_\epsilon(\mu, \tilde{s}^+)) > \pi(b, \mathcal{P}_\epsilon(\mu, \tilde{s}^+))$$

Hence, if (a) holds, we have that A^* is also a PESC of $G^\#$.

If, on the contrary, (b) holds we have that, clearly

$$\pi(a_h, \mathcal{P}_\epsilon(\mu, \tilde{s}^+)) \geq \pi(b, \mathcal{P}_\epsilon(\mu, \tilde{s}^+))$$

and $|a_h| < |b|$, which also satisfy the conditions for PESC.

Hence, in either case, we have seen that if A^* is an ESC of $G^\#$ and G has a strictly dominant strategy, A^* is also a PESC of $G^\#$. \square

Proposition 5.2. *Let A^* be a PESC. Then, A^* is an ESC.*

Proof. Simply notice that the definition of PESC includes the definition of ESC as a particular case. \square

The two corollary that follow extend our two main results to polymorphous populations and are easily obtained from the previous results.

Corollary 5.1. *Let $A^* \subset \mathcal{A}$ be a PESC. Then, $\forall a \in A^*$, $\pi(a, a) = \pi^*$.*

Proof. Combine proposition 5.2 and theorem 4.1. \square

Corollary 5.2. *Suppose that there is a unique $s^* \in S$ such that $U(s^*, s^*) = \pi^*$. Then if A^* is a PES, it is unique.*

Proof. Combine proposition 5.2 and theorem 4.2. □

Therefore, if we consider polymorphous populations, the efficiency and uniqueness results obtained previously are preserved.

6 The Repeated Prisoners' Dilemma

In this section we analyze the Repeated Prisoners' Dilemma based on the game given in Figure 2 using the approach developed in the previous sections. We will find that the unique ESC in this game contains only two automata that implement the well known *Tit-for-Tat* and *Grim trigger* strategies.

	C	D
C	3,3	0,5
D	5,0	1,1

Figure 2: The Prisoners' Dilemma

Let $A^* = \{T, G\}$ be the set containing the automaton T , that implements the strategy *Tit-for-Tat*, and the automaton G , that implements the strategy *Grim trigger*, depicted in Figure 1.

Lemma 6.1. A^* is a PIC

Proof. This result is straightforward, just notice that:

$$\begin{aligned} \pi(G, \mathcal{P}_\epsilon(T, \tilde{D})) &= \pi(T, \mathcal{P}_\epsilon(T, \tilde{D})) = \\ &= \pi(G, \mathcal{P}_\epsilon(G, \tilde{D})) = \pi(T, \mathcal{P}_\epsilon(G, \tilde{D})) = \epsilon + 3(1 - \epsilon) \end{aligned}$$

□

Proposition 6.1. A^* is an ESC

Proof. First, we will prove that no automaton from outside A^* can successfully invade a population consisting mainly of automata of type G . For that, note that

$$\pi(G, \mathcal{P}_\epsilon(G, \tilde{D})) = \epsilon + 3(1 - \epsilon)$$

is the highest payoff achievable against $\mathcal{P}_\epsilon(G, \tilde{D})$. Thus, no automaton can do strictly better against $\mathcal{P}_\epsilon(G, \tilde{D})$ than G itself. Therefore, if $a \in \mathcal{A} \setminus A^*$ is a potential invader of $\mathcal{P}_\epsilon(G, \tilde{D})$, it must be the case that:

$$\pi(a, \mathcal{P}_\epsilon(G, \tilde{D})) = \epsilon + 3(1 - \epsilon)$$

In order to achieve a payoff of 3 against G , a must necessarily be such that

$$\lambda^a(q_0^a) = C \quad \text{and} \quad \lambda^a(\delta^a(q_0^a, C)^n) = C \quad \forall n \geq 1 \quad (8)$$

where $\delta^a(q, C)^n = \delta^a(\delta^a(q, C)^{n-1}, C)$ for $n \geq 2$ (similarly for $\delta^a(q, D)^n$). That is, $\delta^a(q, C)^n$ is the active state for automaton a when, being at q , its opponent plays C n consecutive times.

Also, a would need to attain the maximum payoff against \tilde{D} , so that it must be the case that

$$\exists \hat{q} \in Q^a \quad \text{such that} \quad \lambda^a(\delta^a(\hat{q}, D)^n) = D \quad \forall n \geq 1 \quad (9)$$

Now, (8) and (9) together imply that $\pi(a, \mathcal{P}_\epsilon(a, \tilde{D})) = \pi(G, \mathcal{P}_\epsilon(a, \tilde{D}))$ and also that $|a| \geq 2$. Thus, according to (iii) in the definition of ESC (2), if a can invade $\mathcal{P}_\epsilon(G, \tilde{D})$, it is necessary that

$$2 \leq |a| \leq |G| = 2 \quad (10)$$

It is easy to verify that only two automata satisfy (8), (9) and (10), the one implementing the *Grim Trigger* strategy itself and the one implementing the *Tit-for-Tat* strategy. Therefore, we must conclude that no automaton from outside A^* can enter the population $\mathcal{P}_\epsilon(G, \tilde{D})$.

It remains to prove the analogous for a population consisting mainly on automata of type T , that is, no automaton from outside A^* can successfully invade a population $\mathcal{P}_\epsilon(T, \tilde{D})$. As before, no automaton can do better against $\mathcal{P}_\epsilon(T, \tilde{D})$ than T itself. Therefore, we must look for automata that satisfy $\pi(a, \mathcal{P}_\epsilon(T, \tilde{D})) = \epsilon + 3(1 - \epsilon)$. Hence, such automata must satisfy $\pi(a, \tilde{D}) = 1$ and $\pi(a, T) = 3$, which implies

$$\exists \hat{q} \in Q^a \quad \text{such that} \quad \lambda^a(\delta^a(\hat{q}, D)^n) = D \quad \forall n \geq 1 \quad (11)$$

$$\exists q^* \in Q^a \quad \text{such that} \quad \lambda^a(\delta^a(q^*, C)^n) = C \quad \forall n \geq 1 \quad (12)$$

Clearly, (11) and (12) imply that $|a| \geq 2$. Additionally, we have that $\pi(T, \mathcal{P}_\epsilon(a, \tilde{D})) = \epsilon + 3(1 - \epsilon)$, so that if a can invade $\mathcal{P}_\epsilon(T, \tilde{D})$, it is necessary that

$$2 \leq |a| \leq |T| = 2 \quad (13)$$

As before, only two automata satisfy (11), (12) and (13), those in A^* . Therefore, we must conclude that no automaton from outside A^* can enter the population $\mathcal{P}_\epsilon(T, \tilde{D})$. \square

Clearly, since the Prisoners' Dilemma has a strategy that is strictly dominant, the strategies *Tit-for-Tat* and *Grim trigger* also form a PESC. The next proposition formally states this fact.

Proposition 6.2. $A^* = \{T, G\}$ is the unique PESC of the Repeated Prisoners' Dilemma.

Proof. Clearly, D is a strictly dominant strategy in the Prisoners' Dilemma. Therefore, according to propositions 6.1 and 5.1 A^* is a PESC. To see that it is unique, apply corollary 5.2. \square

7 Games of Common Interests

The games of *common interests* where first studied by Aumann and Sorin [2] and are related to a problem considered one year earlier by Harsanyi and Selten [6]. The problem refers to the question of what equilibrium should be used as the prediction for the outcome of a game if there exist more than one equilibrium. Consider, for example, the game in Figure 3 (from Harsanyi and Selten [6]):

	s_1	s_2
s_1	9,9	0,8
s_2	8,0	7,7

Figure 3: A game of common interests

Both (s_1, s_1) and (s_2, s_2) are Nash equilibria of the game, but the pair of payoffs (9,9) associated to (s_1, s_1) strictly Pareto-dominates any other payoff vector. For that reason, one might think of (s_1, s_1) as the “natural” outcome of the game. Nevertheless, (s_2, s_2) also constitutes an equilibrium of the game, the so called *risk-dominant* equilibrium according to Harsanyi and Selten’s terminology. The problem in this situation is the following: what assumptions can we make on the way the game is played so that the *Pareto-dominant* equilibrium is selected ?

Aumann and Sorin define a game of *common interests* as a (2-persons) game that has a simple payoff pair that strongly Pareto-dominates all other payoff pairs.⁵ They consider repetitions of the game with each player attaches a small but positive probability to the other playing some fixed strategy with bounded recall. They show that this game has an equilibrium in pure strategies with payoffs “close” to the Pareto-dominant vector.

Analyzing this type of games using the approach developed in the previous sections, it turns out that a repeated game based on a game of common interests always has an ESC. Moreover, this ESC consists of an unique automaton that always plays the action associated to the Pareto-dominant equilibrium. The simplicity of this result might be surprising, but it is a rather natural consequence of the the introduction of complexity costs in the evolutionary framework. Indeed, any potential equilibrium automaton should, in the first place, attain the Pareto-dominant payoff for otherwise a successful invader could be constructed (very much like in the proof of Theorem 4.1) in such a way that attains this Pareto-dominant payoff when playing against itself but imitates the original automaton if faced against it. Therefore, the “noisy” players that respond “myopically” to the behavior of the majority of the population will play this Pareto-dominant strategy all the time. In this situation, it is unnecessarily costly to carry extra states that will not be used in any case.⁶

The following definition and proposition formalize this discussion.

Definition 7.1. [Aumann and Sorin] A game G is of common interests if there is a payoff pair $(U_1(s_1^*, s_2^*), U_2(s_1^*, s_2^*))$ that strongly Pareto-dominates all other outcomes, i.e., such that

$$U_i(s_1^*, s_2^*) > U_i(s_1, s_2) \quad \forall s_1, s_2 \in S \quad \forall i = 1, 2$$

Proposition 7.1. *Let G be a game of common interest and let s^* denote the action associated to the Pareto-dominant equilibrium. Then $G^\#$ has a unique ESC consisting of the one-state automaton that always plays the action s^* .*

Proof. Since (s^*, s^*) strictly Pareto-dominates any other pair of payoffs, we have that $\beta(\tilde{s}^*) = \{s^*\}$. Hence, $\pi(s^*, \mathcal{P}(s^*, s^*)) = U(s^*, s^*)$. Let $a \in \mathcal{A}(a \neq \tilde{s}^*)$ be a

⁵There might or might not exist other equilibria of the game.

⁶For a similar reason, the one-state automaton that always plays the action associated to the Pareto-dominant equilibrium is also a MESS (Binmore and Samuelson (1992)). Nevertheless, it is not an ESS (Maynard-Smith (1982))

potential invader. Then a must satisfy $\pi(a, \mathcal{P}(s^*, s^*)) = U(s^*, s^*)$, which implies that $\beta(a) = \{s^*\}$ and hence $\pi(s^*, \mathcal{P}(a, s^*)) = \pi(a, \mathcal{P}(a, s^*))$ and $|\tilde{s}^*| < |a|$. That is, \tilde{s}^* cannot be invade by any automaton thanks to its optimal performance and minimal complexity. Therefore, $\{\tilde{s}^*\}$ is an ESC. Moreover, it is unique in virtue of Theorem 4.2. \square

8 Conclusions

In this paper we analyze the evolutionary stability of repeated symmetric games in the context of the Abreu and Rubinstein's automaton selection games. We focus on a modification of Probst's [11] solution concept (ESC) which is a modification of Binmore and Samuelson's [4] MESS which, in turn, is a modification of the conventional ESS to adapt it to the automata context with complexity costs. We show that, if any population contains always a small fraction of people playing a short run best response to their environment, then there is only one set of automata that is evolutionary stable in the sense that no automaton from outside the set can successfully enter a population composed mainly of such automata. Furthermore, those automata are efficient in the sense that they maximize the payoff that an automaton obtains when playing against a replica of itself, and their complexity is bounded.

Some interesting results are obtained when analyzing typical symmetric games such as the Repeated Prisoners' Dilemma. It turns out that the unique population that is evolutionary stable according to the conditions imposed in this paper is the one composed of the most renowned strategies in the broad literature devoted to this game: *Tit-for-Tat* and *Grim trigger*. Also, for *games of common interests*, there always exists an ESC consisting of the automata that always Plays the Pareto-dominant strategy

We believe that the main attractive of this approach is that simple hypothesis, such as assuming that in any population there is always some people that respond myopically to the behavior of the majority, lead to desirable results such as uniqueness and efficiency. Additionally, although the definition of the ESC solution concept might look involved, it turns out that the actual computation of Evolutionary Stable Collections in any specific 2-players game is rather straightforward. For that, we just need to determine first what strategy of the stage game allows to achieve the utilitarian payoff. Such strategy will configure the initial state of the equilibrium automata and the corresponding action. Second (and last) we must determine what is the "one shot" (myopic) best reply to that action and find its corresponding best reply to configure the remaining states and associated actions (and transitions) of the equilibrium automata.

The picture, though, is far from been complete. We think that further research in the direction of relating this approach to some "traditional" or "rationality-based" solution concepts like the ones considered in game theory would be of interest. The goal of such an exercise would be to determine whether the type of evolutionary stability proposed in this paper induces some sort of "rationality" and, if so, of what kind. Furthermore, the solution concept proposed here is a purely static one. The study of dynamical processes that encompass the ideas discussed in this paper would be, in our opinion, of great interest.

References

- [1] Dilip Abreu and Ariel Rubinstein, *The structure of nash equilibrium in repeated games with finite automata*, *Econometrica* **56** (1988), 1259–82.
- [2] Robert Aumann and Sylvain Sorin, *Cooperation and bounded recall*, *Games and Economic Behavior* **1** (1989), 5–39.
- [3] Robert J. Aumann, *Survey of repeated games*, *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern* (R.J. Aumann et al., ed.), Bibliographisches Institut, Zurich, 1981, pp. 11–42.
- [4] Ken Binmore and Larry Samuelson, *Evolutionary stability in repeated games played by finite automata*, *Journal of Economic Theory* **57** (1992), 278–305.
- [5] David J. Cooper, *Supergames Played by Finite Automata with Finite Costs of Complexity in an Evolutionary Setting*, *Journal of Economic Theory* **68** (1996), 266–275.
- [6] John C. Harsanyi and Reinhard Selten, *A general theory of equilibrium selection in games*, MIT Press, Cambridge, Massachusetts, 1988.
- [7] J.E. Hopcroft and J.D. Ullman, *Introduction to automata theory, languages, and computation*, Addison-Wesley, Reading, MA, 1979.
- [8] Ehud Kalai and W. Stanford, *Finite rationality and interpersonal complexity in repeated games*, *Econometrica* **56** (1988), 397–410.
- [9] John Maynard Smith, *Evolution and the theory of games*, Cambridge University Press, Cambridge, 1982.
- [10] Oscar Volij, *In Defense of DEFECT*, *Games and Economic Behavior* **39** (2002), 309–321.
- [11] Daniel A. Probst, *Evolution, automata and the repeated prisoner's dilemma*, no. 99-55 Sonderforschungsbereich 504, Universität Mannheim. (1992).
- [12] Ariel Rubinstein, *Finite automata play the repeated prisoners' dilemma*, *Journal of Economic Theory* **39** (1986), 83–96.