In defense of DEFECT *

Oscar Volij †

May 8, 2001

Abstract: The one-state machine that always defects is the only evolutionarily stable strategy in the machine game that is derived from the prisoners' dilemma, when preferences are lexicographic in complexity. This machine is the only stochastically stable strategy of the machine game when players are restricted to choosing machines with a uniformly bounded complexity. *Journal of Economic Literature* Classification Numbers: C70, C72.
Keywords: Automata, complexity, prisoner's dilemma.

*I thank Roberto Serrano, an anonymous referee and an associate editor for very helpful comments.
†Department of Economics, Iowa State University, Ames, Iowa 50011, and Department of Economics, Hebrew University, Jerusalem 91905, Israel. E-mail: Oscar@Volij.co.il

1 Introduction

The prisoner's dilemma is hardly a game. It consists of two independent, one person decision problems, each with one clear solution. That is why it has a unique equilibrium. When the prisoner's dilemma is repeated a finite number of times, the resulting game no longer has an equilibrium in dominant strategies but the unique equilibrium outcome still consists of mutual defection at every stage. Only when the prisoner's dilemma is repeated infinitely many times can we find several equilibrium outcomes. Eternal mutual defection, however, is still an equilibrium outcome. Abreu and Rubinstein (1988) consider a two player game –the machine game introduced in Rubinstein (1986)– where players choose a finite automaton that implements an infinitely repeated prisoner's dilemma game strategy and where a pair of strategies are evaluated lexicographically first according to the limit of the means of the stream of one shot payoffs, and second according to the complexity of the machines. Abreu and Rubinstein (1988) show that the set of equilibrium outcomes of the machine game is much smaller than the set of equilibrium outcomes in the repeated game, but still the pair of automata that defect at every stage independently of history is an equilibrium. Fudenberg and Maskin (1990) apply the concept of evolutionary stability to a noisy version of the machine game and succeed in obtaining the cooperative payoff as the unique equilibrium outcome, provided players are endowed with lexicographic beliefs about mistakes.

In an interesting and thought provoking paper, Binmore and Samuelson (1992) introduced a solution concept for the Abreu and Rubinstein (1988) machine game that prevents eternal defection from being an equilibrium outcome. Moreover, in any such equilibrium the automata cooperate at almost every stage. The novel concept, which is called *modified evolutionarily stable strategy* (MESS), has an evolutionary flavor. It intends to capture the same idea as the one behind the standard concept of *evolutionarily stable strategies* (ESS): an action is evolutionarily stable if a population composed of individuals who use that action is immune to invasion by a sufficiently small group of mutants. The difference between the two concepts lies in the importance they assign to complexity considerations for the chances of a population to be invaded. While the ESS concept takes into account the machine game's payoff function in the requirements for stability, the MESS concept first takes the stream of one shot payoffs and then, lexicographically, the complexity of the machines. This difference derives from the different preferences over lotteries that underlie the two concepts. While the formal definition of the machine game analyzed in Rubinstein (1986), Abreu and Rubinstein (1988), and Binmore and Samuelson (1992) does not require any specification of preferences over lotteries, both the application of ESS and MESS do.

This specification is necessary for the evaluation of automata facing a mixed population. Under ESS, preferences have the expected utility form, while under MESS preferences first take into account the expected mean payoff and then, lexicographically, the complexity of the machines.

Our first result is that the only machine that satisfies the standard ESS requirements in the machine game with preferences that are lexicographic in the limit of the means and in complexity, is the one-state machine that defects forever, also known as DEFECT. That is, not only is there a unique evolutionary stable outcome, but also a unique strategy that leads to that outcome.

Cooper (1996) shows that when complexity costs do not enter lexicographically in the payoff function, Binmore and Samuelson's result also changes dramatically. Specifically, he shows that any convex combination of the symmetric efficient payoff and the Nash equilibrium payoff can be approximately achieved by a Neutrally Stable Strategy (NSS) if a positive and small enough cost of complexity is allowed. Cooper (1996) argues that the striking difference in the number of equilibrium outcomes is due to the fact that complexity costs are positive. In contrast, we argue that it is the solution concept employed, and not the cost technology what explains the difference in the results. If one applies the standard ESS, with or without lexicographic costs of complexity, to the machine game, one gets DEFECT as the unique evolutionary stable strategy. Similarly, proceeding as in Cooper (1996), if one applies standard NSS to the machine game, with or without lexicographic costs, one obtains a folk theorem. The difference between Cooper's NSS and Binmore and Samuelson's solution concept is that the former assumes that risk preferences have the expected utility form, while the latter assumes non-expected utility preferences.

ESS is not the only solution concept that has an evolutionary idea behind it. Foster and Young (1990) developed the concept of *stochastic stability* which requires a population to be immune to persistent random mutations. This concept has been successfully applied by Kandori, Mailath, and Rob (1993) in the analysis of symmetric 2×2 games, by Young (1993a) for weakly acyclic n person games, by Young (1993b) in the analysis of bargaining, by Vega-Redondo (1997) in the analysis of competition among firms, and recently by Ben-Shoham, Serrano, and Volij (2000) in the analysis of a housing problem. The idea of stochastic stability is to add a small perturbation to the evolutionary dynamics on some game and select all those actions that are assigned positive probability by the limit of the long run distribution when the perturbation becomes arbitrarily small. It turns out that the stochastically stable actions are contained in the set of actions that are in the recurrent classes of the unperturbed process. When we apply the concept of stochastic stability to the machine games where the players are constrained to choose among automata with complexity that does not exceed some fixed number, the only stochastically stable machine is DEFECT.

Bergin and Lipman (1996) pointed out a weakness of the concept of stochastic stability: the actions selected out of the recurrent classes depend on the rate of mutation. More specifically, they show that any invariant distribution of the unperturbed process is the limit of some sequence of perturbed processes. In our model, however, the only recurrent class of the unperturbed process consists of the one-state automaton that defects forever. Consequently, DEFECT is the only stochastically stable machine *for any model of mutations*.

2 Basic Definitions

A symmetric two-player game is a pair $\langle A, u \rangle$ where A is the common set of actions and $u : A^2 \to \mathbb{R}$ is the players' common utility function that assigns to each pair of actions, one for the player and the other for his opponent, a utility level. The prisoner's dilemma is the game $G = \langle A, u \rangle$ where $A = \{C, D\}$ and u is given by the following matrix:

One player \ The opponent	C	D
C	α	γ
D	δ	β

where $\gamma < \beta < \alpha < \delta$.

A *Nash equilibrium* of a two-player game $\langle A, u \rangle$ is a pair of actions $(s, t) \in A^2$ such that for all $r \in A$, $u(s, t) \geq u(r, t)$ and $u(t, s) \geq u(r, s)$. A Nash equilibrium (s, t) is *strict* if the above inequalities are strict for all r different from s and t, respectively.

Now, we turn to the machine game first analyzed by Rubinstein (1986). A finite automaton that can play the repeated prisoners' dilemma is a quadruple $\langle Q, q^0, \lambda, \mu \rangle$ where

- Q is a finite set of states
- q^0 is the initial state of the machine
- $\lambda : Q \to A$ is the output function, that returns an action chosen as a function of the state, and
- $\mu : Q \times A \to Q$ is the transition function that returns the next state of the machine, as a function of the present state and of the action chosen by the opponent.

Each machine represents a repeated game strategy but there are some repeated game strategies that cannot be represented by a finite automaton. One famous automaton is DEFECT, which consists of a single state at which it returns the action D. Another one is COOPERATE, which consists of a single state at which it cooperates. When two automata meet and play against each other, they determine a play in the repeated game. Formally, let $a_1 = \langle Q_1, q_1^0, \lambda_1, \mu_1 \rangle$ and $a_2 = \langle Q_2, q_2^0, \lambda_2, \mu_2 \rangle$ be two automata. The evolution of the machines' states when they play against each other is defined recursively as follows:

$$q^0 = (q_1^0, q_2^0)$$

and for $t > 0$

$$q^{t+1} = (\mu_1(q_1^t, \lambda_2(q_2^t)), \mu_2(q_2^t, \lambda_1(q_1^t)).$$

Similarly, the evolution of play when the machines play against each other is given by

$$h^0 = (\lambda_1(q_1^0), \lambda_2(q_2^0))$$

and for $t > 0$,

$$h^{t+1} = (\lambda_1(q_1^t), \lambda_2(q_2^t)).$$

We are interested in the game where players choose a finite automaton and where a pair of automata are evaluated lexicographically, first according to the limit of the means criterion and second according to the complexity of the machine. More formally, let $\mathcal{A}$ be the set of finite automata. The complexity of an automaton $a \in \mathcal{A}$ is defined to be the number of its states and denoted by $|a|$. Given two automata $a_1 = \langle Q_1, q_1^0, \lambda_1, \mu_1 \rangle$ and $a_2 = \langle Q_2, q_2^0, \lambda_2, \mu_2 \rangle$, let $P(a_1, a_2) = \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T} u(h^t)$ be the limit of the means of the stream of payoffs generated when they play each other.[1] A *machine game* is a symmetric game $\langle \mathcal{A}, U \rangle$ where U is a payoff function that satisfies

$$U(a,b) > U(c,d) \Leftrightarrow \begin{cases} P(a,b) > P(c,d) & \text{or} \\ P(a,b) = P(c,d) & \text{and } |a| < |c| \end{cases}$$

In what follows we fix an arbitrary machine game $\langle \mathcal{A}, U \rangle$ and denote it $G^{\#}$. All the results in the paper are independent of the particular choice of representation U.

The following lemma, taken from Abreu and Rubinstein (1988), is the basis of all the results that follow. Since the result is so useful and its proof is so short, I provide it here.

[1]Since we are dealing with finite automata, this limit exists, and hence $P(a_1, a_2)$ is well-defined.

Lemma 1 [Abreu and Rubinstein (1988)] For any machine $a \in \mathcal{A}$, there is a machine $b \in \mathcal{A}$ such that $|b| = |a|$ and $P(b, a) \geq P(c, a)$ for all $c \in \mathcal{A}$.

Proof : Given automaton $a = \langle Q, q^0, \lambda, \mu \rangle$, the choice of a sequence of one-shot actions that maximizes the limit of the means of the resulting sequence of one-shot payoffs is a Markov decision problem that has a stationary solution $s : Q \to A$ (see Puterman (1994), Theorem 8.4.5). Consider $b = \langle Q, q^0, s, \mu_b \rangle$, where μ_b is chosen so that $\mu_b(q, \lambda(q)) = \mu(q, s(q))$ for all $q \in Q$, while it is unrestricted otherwise. Automaton b satisfies the requirements of the lemma.

$\square$

Remark 1 It follows from the proof of Lemma 1 that if the automaton a has more than one state, then the automaton b identified in Lemma 1 is not unique. (To see this, note that if a has more than one state, there are several ways to choose a transition function μ_b that satisfies $\mu_b(q, \lambda(q)) = \mu(q, s(q))$). Therefore, if b is a best reply to a such that $|b| = |a| > 1$, then b is not the unique best reply to a. Further, one of the best replies can be chosen so that its transition function is independent of the opponent's action.

3 Evolutionary Stability

The first solution concept we want to apply is evolutionary stability.

Definition 1 [Maynard Smith and Price (1973)] Let $\langle A, u \rangle$ be a symmetric two-player game. An action $s \in A$ is an *evolutionarily stable strategy* (ESS) if for all $t \in A$, $t \neq s$ we have either

1. $u(s, s) > u(t, s)$ or
2. $u(s, s) = u(t, s)$ and $u(s, t) > u(t, t)$.

Implicit in the above definition of evolutionary stability is the assumption that u is a von-Neumann-Morgenstern utility function. In particular, when opposed to a population mixture $(1-\epsilon)x+\epsilon y$, action s is preferred to action t if and only if $(1-\epsilon)u(s, x)+\epsilon u(s, y) > (1-\epsilon)u(t, x)+\epsilon u(t, y)$. Therefore, although probability mixtures are not present in the above definition, one should bear in mind that preferences over lotteries have the expected utility form.

There is a problem with the concept of an evolutionarily stable strategy: an action may not be ESS just because it can be invaded by an essentially equivalent action. To see this, consider the following game:

One player \ The opponent	s_1	s_2	s_3
s_1	3	3	0
s_2	3	3	1
s_3	1	0	2

This game has a unique ESS, s_3, which is its only strict Nash equilibrium. The reason s_1 is not an ESS is that it can be invaded by s_2. Similarly, s_2 is not an ESS because it can be invaded by s_1. But s_1 and s_2 are indistinguishable in a population composed solely by elements of $\{s_1, s_2\}$. There is a sense in which the set $\{s_1, s_2\}$ is evolutionarily stable. If the population is composed only of elements of $\{s_1, s_2\}$, in whatever proportion, strategy s_3 cannot invade. There will be drift, but the population will remain composed of s_1 and s_2 alone. In order to develop this idea, we need to extend the definition of an evolutionarily stable strategy so as to allow for evolutionarily stable sets of strategies: we say that $B \subseteq A$ is an *evolutionarily stable family of strategies* (ESF) if for all $s, t \in B$ and for all $r \in A \setminus B$ we have

- $u(s,s) = u(s,t) = u(t,s) = u(t,t)$ and
- 1. $u(s,s) > u(r,s)$ or
  2. $u(s,s) = u(r,s)$ and $u(s,r) > u(r,r)$.

Note that if s is an evolutionarily stable strategy, then $\{s\}$ is an ESF. Also, if s belongs to an ESF then (s,s) is a Nash equilibrium. This definition captures the idea that any population composed by elements of an ESF, B, will be able to reject any attempt of invasion by any mutant outside B, and only by a mutant outside B.

We are now ready to state the following:

Claim 1 The only evolutionarily stable strategy of $G^{\#}$ is DEFECT. Further, {DEFECT} is the only evolutionarily stable family of $G^{\#}$.

Proof : Since (D,D) is a Nash equilibrium of the one-shot prisoners' dilemma G, it follows that $P(\text{DEFECT}, \text{DEFECT}) \geq P(b, \text{DEFECT})$ for all automata $b \in \mathcal{A}$. The fact that DEFECT is an ESS follows from the observation that the only one-state machine that performs against DEFECT as well as DEFECT is DEFECT itself. To see that there is no other ESF than {DEFECT}, note first that COOPERATE cannot belong to an ESF because (COOPERATE, COOPERATE) is not a Nash equilibrium. Let $a = \langle Q, q^0, \lambda, \mu \rangle$ be an automaton with at least two states and assume by contradiction that it belongs to an ESF, denoted by B. Consider the automaton $b = \langle Q, q^0, \lambda, \mu_b \rangle$ where $\mu_b(q, C) = \mu_b(q, D) = \mu(q, \lambda(q))$ for all $q \in Q$. By construction, $U(a, a) = U(a, b) = U(b, a) = U(b, b)$. Therefore, since a belongs to B, so does b. Since b's transition function is independent of the opponent's action, DEFECT is a best response to b. That is, $U(\text{DEFECT}, b) \geq U(b, b)$. But since $|DEFECT| < |b|$, we have $U(\text{DEFECT}, b) > U(b, b)$ which contradicts the fact that $b \in B$. □

No Nash equilibrium of the machine game, except for DEFECT, is evolutionarily stable.[2] The reason is that if a is an automaton with at least two states that is part of a symmetric Nash equilibrium of the machine game, there is another automaton, b, with the same number of states, which, when matched with a, behaves identically to a. Automata a and b differ only in their "out of equilibrium" behavior. Claim 1 shows, however, that a and b cannot co-exist even in an evolutionarily stable polymorphic population.

4 Comparison with MESS

In this section we compare the concept of ESS with an alternative solution concept, proposed by Binmore and Samuelson (1992). To facilitate the comparison, we apply the definition of ESS to the machine game and get the following observation, the proof of which is left to the reader.

Observation 1 Automaton $a \in \mathcal{A}$ is an *evolutionarily stable strategy* (ESS) of the machine game $G^{\#}$ if, and only if, for all $b \in \mathcal{A}$, $b \neq a$ we have either

1. $P(a, a) > P(b, a)$ or
2. $P(a, a) = P(b, a)$ and $|a| < |b|$ or
3. $P(a, a) = P(b, a)$ and $|a| = |b|$ and $P(a, b) > P(b, b)$.

[2] Claim 1 seems to contradict some statements in the literature; see Binmore and Samuelson (1992) pages 282 and 287, and Fudenberg (1992), endnote 39. We comment on the Binmore-Samuelson approach in Section 4.

The alternative solution concept is defined as follows: [3]

Definition 2 Automaton $a \in \mathcal{A}$ is an ESS* of the machine game $G^\#$ if, and only if, for all $b \in \mathcal{A}$, $b \neq a$ we have either

1. $P(a,a) > P(b,a)$ or

2*. $P(a,a) = P(b,a)$ and $P(a,b) > P(b,b)$ or

3*. $P(a,a) = P(b,a)$ and $P(a,b) = P(b,b)$ and $|a| < |b|$.

Note that the requirements in Observation 1 differ from Binmore and Samuelson's ESS*. Loosely speaking, ESS* is ESS applied to P rather than U with an extra line, (3*), that applies when $P(a,a) = P(b,a) = P(a,b) = P(b,b)$ and in which complexity matters. The reader may find the following alternative way to state the above definitions suggestive: Automaton a is ESS if there is an $\bar{\epsilon} \in (0,1)$ such that for all $\epsilon \in (0, \bar{\epsilon})$ and for all $b \neq a$

$$(1-\epsilon)U(a,a) + \epsilon U(a,b) > (1-\epsilon)U(b,a) + \epsilon U(b,b),$$

while automaton a is ESS* if there is $\bar{\epsilon} \in (0,1)$ such that for all $\epsilon \in (0, \bar{\epsilon})$ and for all $b \neq a$

1. $(1-\epsilon)P(a,a) + \epsilon P(a,b) > (1-\epsilon)P(b,a) + \epsilon P(b,b)$ or
2. $(1-\epsilon)P(a,a) + \epsilon P(a,b) = (1-\epsilon)P(b,a) + \epsilon P(b,b)$ and $|a| < |b|$.

One can see that ESS takes stock of the *expected utility* obtained by each automaton facing a mixed population. That is, the expectation is taken over the utilities derived from each possible match. The concept of ESS*, on the other hand, lexicographically considers first the *expected mean stage-payoff* and then complexity. In other words, the implicit assumption behind the definition of ESS* is that when opposed to a population mixture $(1-\epsilon)x + \epsilon y$, automaton a is preferred to automaton b if and only if

1. $(1-\epsilon)P(a,x) + \epsilon P(a,y) > (1-\epsilon)P(b,x) + \epsilon P(b,y)$ or
2. $(1-\epsilon)P(a,x) + \epsilon P(a,y) = (1-\epsilon)P(b,x) + \epsilon P(b,y)$ and $|a| < |b|$.

These preferences do not have the expected utility form. Consider for instance the automata a and b in figure 1.

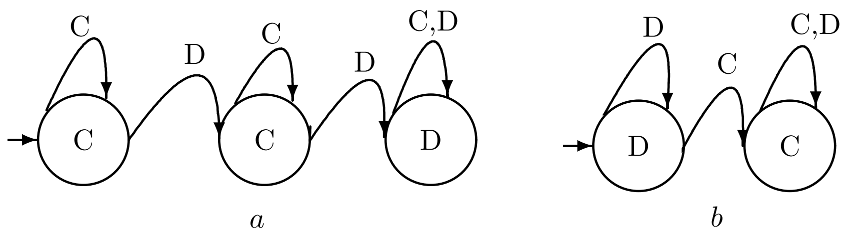

Figure 1: Two automata.

Automaton a is a three-state machine that cooperates until the opponent defects for the second time. Once this happens, it switches to defection forever. Automaton b, on the other hand, is a two-state automaton that starts defecting and continues to do so as long as its opponent defects. As soon as the opponent cooperates, however, it switches to cooperation forever. One can check that $P(a,a) = P(a,b) = P(b,a) > P(b,b)$. Therefore, against a population $(1-\epsilon)a+\epsilon b$ automaton a is preferred to automaton b, forevery $\epsilon \in (0,1)$. Against a monomorphic population of a (that is, for $\epsilon = 0$), however, b is preferred to a since $|b| < |a|$. That is, the continuity axiom is violated.

This discontinuity of the ESS* preferences has substantial consequences. It means that the ESS* solution concept is not a refinement of Nash equilibrium. To see this, consider a 2×2 machine selection game $\langle \mathcal{A}', U \rangle$ where the set $\mathcal{A}'$ consists of the two automata of figure 1. Automaton a is an ESS* because for all $\epsilon \in (0,1)$ we have

$$(1-\epsilon)P(a,a) + \epsilon P(a,b) > (1-\epsilon)P(b,a) + \epsilon P(b,b)$$

(automaton a enjoys the fruits of cooperation almost every stage while the mutant b gets mutual defection forever with probability ϵ). On the other hand, (a,a) is not a Nash equilibrium *of the machine game* because b is a profitable deviation for any of the players. This is so because, although $P(a,a) = P(b,a)$, automaton b has a lower complexity than that of automaton a: $|b| = 2 < 3 = |a|$. This example shows that an ESS* need not be a Nash equilibrium. Needless to say, it has long been proven that a standard ESS is a Nash equilibrium.

[3] Binmore and Samuelson (1992) note that no automaton satisfies the ESS* requirements when aplied to $G^{\#}$ and thus they define MESS to be a weaker version of ESS* where the last inequality of 3* is replaced by a weak inequality. In the sequel, all the statements concerning ESS* also apply to MESS.

5 Stochastic Stability

We now turn to the concept of stochastic stability which is the second well-established evolutionary solution concept. As opposed to the ESS concept, it allows for distinct mutants to co-exist. For each $n \in \mathbb{N}$, let $\mathcal{A}(n)$ be the set of automata with at most n states and let G^n be the game $G^n = \langle \mathcal{A}(n), U \rangle$. There is only one difference between the game G^n and $G^{\#}$. In the former, players can choose only finite automata whose complexity does not exceed n while in the latter there is no bound in the complexity of the allowed automata. Consequently, G^n is a finite normal form game. Consider the following learning model. The game G^n is played by a single population of $m > 2$ individuals, where for simplicity m is assumed to be even. At each period, the individuals are randomly matched in pairs and play G^n. For each $a \in \mathcal{A}(n)$ let k_a^t denote the number of individuals that choose automaton a in period t. Consider a Markov process where the state space is the set of all vectors $k = (k_a)_{a \in A(n)}$ such that k_a is a nonnegative integer for all $a \in \mathcal{A}(n)$, and $\sum_{a \in A(n)} k_a = m$. In other words, a state is a list of the numbers of agents that choose any given automaton. As for the transition function, assume that the current state is k. At the beginning of next period each agent draws a sample of size $s < m/2$, without sampling himself and without replacement, from the previous period's actions, and plays a best reply to the resulting sample proportions. If there is more than one best reply, each is played with equal probability.

The dynamic process just described is chosen for its simplicity. Other dynamic processes could have been chosen without affecting the following result.[4]

Claim 2 The process defined above on G^n converges with probability one to the state where all the agents choose DEFECT.

Proof : It is enough to show that the only absorbing state is the one where all agents chose DEFECT, and that there are no recurrent classes other than the singleton containing the unique absorbing state. This is shown in the following lemmas.

Lemma 2 The unique strict Nash equilibrium of G^n is (DEFECT, DEFECT).

Proof : It is clear that the only best reply to DEFECT is DEFECT so (DEFECT, DEFECT) is a strict Nash equilibrium. Let now (a, b) be a Nash equilibrium and assume that one of the

[4] See Cooper (1993) for a similar result.

automata, say b, is not DEFECT. Automaton b cannot be COOPERATE because the only best reply to COOPERATE is DEFECT and (DEFECT,COOPERATE) is not a Nash equilibrium. Therefore $|b| > 1$. By Theorem 1 (a) in Abreu and Rubinstein (1988), $|a| = |b| > 1$. Then, Remark 1 tells us that there are more than one best reply to a, which means that (a, b) is not strict. □

Lemma 3 The only absorbing state of this process is the state where all individuals play DEFECT.

Proof : Since (DEFECT, DEFECT) is a strict Nash equilibrium, the state where all the individuals play DEFECT is an absorbing state. Now let $k = (k_a)_{a \in \mathcal{A}(n)}$ be an absorbing state and let a be an automaton such that $k_a \geq 1$. This means that at every period there is an agent that chooses a. It must be the case that a is this agent's only best reply to every sample he can possibly pick because, otherwise, there would be a positive probability that the process passes from k to some other state k' with $k'_a < k_a$. Assume now there are two automata, a_1 and a_2 with $k_{a_1} \geq 1$ and $k_{a_2} \geq 1$. Then, by the previous argument, there are two agents such that a_1 is the only best reply to every possible sample of the first agent, and a_2 is the only best reply to every possible sample of the second agent. But since the sample size is less than $m/2$, there must be a sample that both agents can simultaneously pick, which implies that $a_1 = a_2$. Therefore, at k, every agent chooses the same automaton a. Consequently, (a, a) must be a strict Nash equilibrium for otherwise there would be a positive probability that the system moves from k to some state other than k. But by Lemma 2, a=DEFECT. □

Lemma 4 There is a positive probability that the system reaches the only absorbing state in a finite number of stages, independently of the initial state.

Proof : Let k^0 be the current state. Since $s < m/2$, there is a positive probability that half the population picks a common sample and the other half picks another common sample. Consequently, there is a positive probability that next period all agents sample the same distribution of automata, and since all best replies have positive probability, there is a positive probability that all m agents choose the same automaton. If this automaton is DEFECT, we are done. If this automaton is COOPERATE, then next period everybody will choose DEFECT and we are done again. So assume that the state is one in which every agent chooses the same automaton a, with $|a| > 1$. Let $(a_n)_{n=1}^{\infty}$ be a sequence of automata such that $a_1 = a$ and for all $n > 1$, a_n is a best reply to a_{n-1}. It follows from Lemma 1 that the sequence $(|a_n|)_{n=1}^{\infty}$ of the corresponding complexities is a non increasing sequence of positive integers. Consequently, there is an $N \in \mathbb{N}$ such that the sequence in $(a_n)_{n=1}^{\infty}$ consists of automata that have the same complexity from N on. This constant complexity is $|a_N|$. By the definition of the process there is a positive probability that the process follows the path $(a_1, \ldots, a_N)$ where t periods after the current state, all agents choose a_t. If $|a_N| = 1$ then we are done because from the next period on, all the agents will choose DEFECT. If $|a_N| > 1$ then $|a_N| = |a_{N+1}| > 1$ and by Remark 1 there is an automaton a^* with $|a^*| = |a_N|$ that is a best reply to a_N and whose transition function is independent of the action of its opponent (it depends only on its own present state). Since the number of best replies to a given automaton is finite, there is a positive probability that the system moves to a state where all the agents choose automaton a^*. But then, since the only best reply to an automaton with a transition function that is insensitive to the opponent's actions is DEFECT, from the next period on, all the agents will play DEFECT. This shows that from any state, there is a positive probability to reach the only absorbing state of the process which implies that this state constitutes its only recurrent class. □

□

The uniform bound on complexity is necessary for the above proof to work. To see this, note that a best reply to a given sample could be an automaton of complexity higher than the complexity of any automaton in the sample. Therefore, if all finite automata are available we cannot preclude the case where the system visits an infinite amount of different states. Consequently, without the bound on complexity Lemma 4 does not imply that the process converges to the unique absorbing state with probability one. Lemma 1 assures, however, that all best replies to any sample of size 1 have a complexity which is no higher than the complexity of the sampled automaton. As a result, if we restrict the agents to take samples of size 1, the corresponding process will visit a finite number of states. This observation, together with the proof of Claim 2, show the following.

Claim 3 When the sample size is restricted to be 1, the process defined above on $G^{\#}$ converges with probability one to the state where all the agents choose DEFECT.

6 Conclusion

We have shown that two well-established solution concepts of evolutionary game theory — standard evolutionarily stable strategies and stochastic stability— have a very strong selection power when applied to the repeated prisoner's dilemma machine game. Not only do they point to a unique equilibrium outcome, but they also select a unique strategy: the never fully appreciated one-state machine that defects forever.

References

Abreu, D. and A. Rubinstein (1988). The structure of Nash equilibrium in repeated games with finite automata. *Econometrica 56*, 1259–1282.

Ben-Shoham, A., R. Serrano, and O. Volij (2000). The evolution of exchange. Working paper, Center for Rationality, Hebrew University of Jerusalem.

Bergin, J. and B. L. Lipman (1996). Evolution with state-dependent mutations. *Econometrica 64*, 943–956.

Binmore, K. G. and L. Samuelson (1992). Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory 57*, 278–305.

Cooper, D. J. (1993). Supergames played by finite automata with finite costs of complexity in an evolutionary setting. Working paper, University of Pittsburgh.

Cooper, D. J. (1996). Supergames played by finite automata with finite costs of complexity in an evolutionary setting. *Journal of Economic Theory 68*, 266–275.

Foster, D. P. and H. P. Young (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology 38*, 219–232.

Fudenberg, D. (1992). Explaining cooperation and commitment in repeated games. In J.-J. Laffont (Ed.), *Advances in Economic Theory*, Volume 1, pp. 89–131. Cambridge: Cambridge University Press.

Fudenberg, D. and E. Maskin (1990). Evolution and cooperation in noisy repeated games. *AER 80*, 274–279.

Kandori, M., G. Mailath, and R. Rob (1993). Learning, mutations and long run equilibria in games. *Econometrica 61*, 29–56.

Maynard Smith, J. and G. R. Price (1973). The logic of animal conflict. *Nature 246*, 15–18.

Puterman, M. L. (1994). *Markov Decision Processes*. John Wiley & Sons, Inc.

Rubinstein, A. (1986). Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory 39*, 83–96.

Vega-Redondo, F. (1997). The evolution of Walrasian behavior. *Econometrica 65*, 375–384.

Young, H. P. (1993a). The evolution of conventions. *Econometrica 61*, 57–84.

Young, H. P. (1993b). An evolutionary model of bargaining. *Journal of Economic Theory 59*, 145–168.