

van Veelen, Matthijs; Garcia, Julian

Working Paper

In and out of Equilibrium II: Evolution in Repeated Games with Discounting and Complexity Costs

Tinbergen Institute Discussion Paper, No. 12-089/I

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: van Veelen, Matthijs; Garcia, Julian (2012) : In and out of Equilibrium II: Evolution in Repeated Games with Discounting and Complexity Costs, Tinbergen Institute Discussion Paper, No. 12-089/I, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/87540>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2012-089/I
Tinbergen Institute Discussion Paper



In and out of Equilibrium II: Evolution in Repeated Games with Discounting and Complexity Costs

Matthijs van Veelen¹
Julian Garcia²

¹ Faculty of Economics and Business, University of Amsterdam, Tinbergen Institute, and CREED;

² Max-Planck-Institute for Evolutionary Biology.

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

In and out of equilibrium II: evolution in repeated games with discounting and complexity costs.

Matthijs van Veelen and Julián García*

September 3, 2012

Abstract

We explore evolutionary dynamics for repeated games with small, but positive complexity costs. To understand the dynamics, we extend a folk theorem result by Cooper (1996) to continuation probabilities, or discount rates, smaller than 1. While this result delineates which payoffs can be supported by neutrally stable strategies, the only strategy that is evolutionarily stable, and has a uniform invasion barrier, is All D. However, with sufficiently small complexity costs, indirect invasions - but now through “almost neutral” mutants - become an important ingredient of the dynamics. These indirect invasions include stepping stone paths out of full defection.

*van Veelen: Department of Economics, Universiteit van Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, and Tinbergen Institute, the Netherlands, c.m.vanveelen@uva.nl. García: Research group for Evolutionary Theory, Max Planck Institute for Evolutionary Biology, August-Thienemann-Str 2, 24306 Plön, Germany, garcia@evolbio.mpg.de.

If more complex strategies for repeated games are also more costly than less complex ones, evolutionary dynamics in populations playing them will differ from the dynamics in the absence of complexity costs. The difference that complexity costs can make has been noticed, also in the non-evolutionary repeated games literature. Classic papers which do account for complexity costs are for instance Rubinstein (1986), Abreu and Rubinstein (1988) and Kalai and Stanford (1988), while classics without complexity costs include Friedman (1971), Aumann and Shapley (1976), Rubinstein (1979, 1980), Fudenberg and Maskin (1986), Abreu (1988) and van Damme (1989). In the evolutionary game theory literature about repeated games, most papers do not account for complexity costs; see for instance Axelrod and Hamilton (1981), Axelrod (1984), Boyd and Lorberbaum (1987), May (1987), Farrell and Ware (1989), Fudenberg and Maskin (1990), Linstér (1992, 1994), Kim (1994), Bendor and Swistak (1995, 1997, 1998), Ho (1996), Miller (1996), Lorberbaum et al. (2002). Notable exceptions are Binmore and Samuelson (1992), Cooper (1996), and Volij (2002), which all model settings where complexity costs are relatively small compared to the payoffs of the game.

Without complexity costs, we know that no pure strategy is evolutionarily stable (Selten and Hammerstein, 1984), and that there is a multitude of strategies that are neutrally stable (Bendor and Swistak, 1995). Robustness against indirect invasions (RAII, van Veelen, 2011) is a refinement of neutral stability that implies the presence of a *set* that is asymptotically stable in the replicator dynamics, and in van Veelen and García (2010) we apply this concept to repeated games. We find that no finite mixture of strategies is RAI, and that the vulnerability to indirect invasions shapes the evolutionary dynamics in finite populations; in simulations, the population finds itself in equilibrium most of the time, but regularly transitions from one equilibrium to the other through these indirect invasions.

In this paper we will again explicitly aim at understanding the dynamics

in repeated games, but now with complexity costs. We will use a Folk Theorem result by Cooper (1996), but because we also want to include repeated games with discounting, we will extend it to continuation probabilities, or discount factors, smaller than 1. With this extension, the result no longer holds true for the entire interval between the symmetric Pareto efficient payoff and the one-shot Nash equilibrium payoff, but only for parts of it. The perimeters of the payoffs that can be supported by an NSS vary with the continuation probability and can be visualized in a tree-like figure.

Another result that will help understand the dynamics is that with positive complexity costs, All D is evolutionarily stable and even has a uniform invasion barrier. Because all other equilibria from the extension of Cooper's (1996) result are not RAI, and therefore can be left through indirect invasions, this suggests that All D is more stable than all other equilibria. Yet, if we run simulations, we find that for low complexity costs, a population visits a range of strategies with different average cooperation levels. That can be understood if we realize that, as complexity costs decrease to 0, mutants that are not neutral with positive complexity costs become "almost neutral" if complexity costs decrease. Transitions that would be indirect invasions in the absence of complexity costs then become the driver of the dynamics again with vanishing complexity costs, and these include indirect invasions out of All D.

Section 2 considers continuation probabilities smaller than 1. In Section 3 we return to infinitely repeated games, and link up with results in Binmore and Samuelson (1992) and Volij (2002), which also consider infinitely repeated games. There we will see that, together with the transitions that are of primary importance to the concepts of a MESS (Binmore and Samuelson, 1992) and a LESS (Volij, 2002), indirect invasions complete the picture of relevant transitions.

1 Preliminaries

We start by repeating a few formal definitions from van Veelen and García (2010). Consider a symmetric one-shot 2-player game g characterized by a set of players $I = \{1, 2\}$, an action space A , equal for both players, and a payoff function $\pi : A \times A \rightarrow \mathbb{R}^2$. Here we will immediately restrict ourselves to prisoners dilemmas, and choose $A = \{C, D\}$. Using a discount factor δ , interpreted as a continuation probability, this one-shot game is turned into a repeated one, which will be called $\Gamma(\delta)$. A history at time t is a list of the actions played up to and including time $t - 1$, where an empty pair of brackets is used to denote the history ‘no history’. If $a_{t,i}$ is the action played by player i at time t , then these histories are:

$$\begin{aligned} h_1 &= () \\ h_t &= ((a_{1,1}, a_{1,2}), \dots, (a_{t-1,1}, a_{t-1,2})), \quad t = 2, 3, \dots \end{aligned}$$

Sometimes we will also write $(h_t, (a_{t,1}, a_{t,2}))$ for a history h_{t+1} . The set of possible histories at time t is:

$$\begin{aligned} H_1 &= \{h_1\} \\ H_t &= \prod_{i=1}^{t-1} (A \times A) \quad t = 2, 3, \dots \end{aligned}$$

and the set of all possible histories is:

$$H = \bigcup_{t=1}^{\infty} H_t.$$

It will furthermore be useful to have a way of writing down a history with the roles of the players reversed. Given a history h_t as they are defined above, its mirror image h_t^{\leftarrow} is found by simply renumbering the players:

$$\begin{aligned}
h_1^{\leftarrow} &= () \\
h_t^{\leftarrow} &= ((a_{1,2}, a_{1,1}), \dots, (a_{t-1,2}, a_{t-1,1})), \quad t = 2, 3, \dots
\end{aligned}$$

The reason why histories with roles reversed are needed, is that we assume that both players label themselves as player 1 and the other as player 2, and therefore face mirrored histories as they go along.

A strategy is a function that maps histories to the action space; $S : H \rightarrow A$. For two strategies, say S and T , the course of actions is determined by recursion; all actions at all stages are determined by the initiation

$$h_1^{S,T} = ()$$

and the recursion step

$$\begin{aligned}
a_t^{S,T} &= \left(S \left(h_t^{S,T} \right), T \left(h_t^{S,T\leftarrow} \right) \right) \\
h_{t+1}^{S,T} &= \left(h_t, a_t^{S,T} \right), \quad t = 1, 2, \dots
\end{aligned}$$

For $0 \leq \delta < 1$ the discounted, normalised payoffs to (a player that uses) strategy S against strategy T is given by:

$$\Pi(S, T) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \pi_1 \left(a_t^{S,T} \right)$$

The discount factor will mostly be interpreted as a continuation probability, which makes $\Pi(S, T)$ the normalized expected payoff of strategy S against T . For $\delta = 1$ we will consider average payoffs, or the limit of means, for the infinitely repeated game:

$$\Pi(S, T) = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=1}^{\tau} \pi_1(a_t^{S, T})$$

Because we need a measure of complexity, we will restrict ourselves to finite state automata. This is also in line with the literature. If a strategy S is a finite automaton, or a Moore machine, then it can then be represented by a tuple $\{\{1, \dots, N_S\}, \lambda_S, \mu_S\}$, where N_S is the number of states, $\lambda_S : \{1, \dots, N_S\} \rightarrow \{C, D\}$ gives the output in every state and $\mu_S : \{1, \dots, N_S\} \times \{C, D\} \rightarrow \{1, \dots, N_S\}$ gives the transitions. Sometimes the formal definition of a machine also states in which state the machine starts, but because the states can always be renumbered so that the starting state is the first, we assume, without loss of generality, that the machine starts in state 1. Also without loss of generality, we assume that if it plays against itself, it will transition from state i to state $i + 1$, until for the first time it goes back to a state it has already been in, or remains in the state it is currently in.

2 Continuation probabilities smaller than 1

2.1 A “tree” full of NSS’es

Here we extend a result by Cooper (1996) to continuation probabilities smaller than 1. Theorem 2 gives perimeters of the combinations of payoffs and continuation probabilities that allow for NSS’es, provided that the complexity costs are not too large. As in Cooper (1996) we define the fitness - or profits - of the players as the payoffs from the game minus the costs:

$$U(S, T) = \Pi(S, T) - k |S|,$$

where $|S|$ is the number of states and $k > 0$ is the per state cost. The only difference with Cooper (1996) is that we also allow $\Pi(S, T)$ to be normalized,

discounted payoffs instead of the limit of means.

It will also be useful to have a value function for the states of strategy S from the point of view of its opponent. Given a continuation probability, or discount factor δ , the value to the opponent of strategy S of S being in state i is denoted by $V_S^*(i)$, $i = 1, \dots, N_S$, and is the solution to the following system of equations.

$$V_S(i) = \max_{a \in \{C, D\}} \{ \pi_1(a, \lambda_S(i)) + \delta V_S(\mu_S(i, a)) \} \quad i = 1, \dots, N_S$$

In the absence of complexity costs, strategy S is a best response to itself if $\lambda_S(i) \in \arg \max_{a \in \{C, D\}} \{ \pi_1(a, \lambda_S(i)) + \delta V_S^*(\mu_S(i, a)) \}$ for all $i = 1, \dots, N_S$. If S is indeed a best response to itself, then the value in the first state is the not normalized, discounted payoff of S playing against itself; $(1 - \delta) V_S^*(1) = \Pi(S, S)$.

If we include complexity costs, and if a strategy S is to be an NSS, or even an equilibrium, then that immediately puts some restrictions on the shape of S . As noted in the literature, from the definition of U it directly follows that an equilibrium cannot have unused states. Also, if S is to be an NSS, then it must defect in the first state.

Lemma 1. *If S is an NSS, then $\lambda_S(1) = D$.*

Proof. We begin with two simple observations. If $\lambda_S(i) = C$ and $\mu_S(i, C) = i+1$, and S is an equilibrium, then $V_S^*(i) = b + \delta V_S^*(i+1)$, which implies that $\frac{b}{1-\delta} - V_S^*(i) = \delta \left(\frac{b}{1-\delta} - V_S^*(i+1) \right)$, and therefore that $V_S^*(i+1) < V_S^*(i)$. If $\lambda_S(i) = D$ and $\mu_S(i, D) = i+1$, and S is an equilibrium, then $V_S^*(i) = c + \delta V_S^*(i+1)$, which implies that $V_S^*(i) - \frac{c}{1-\delta} = \delta \left(V_S^*(i+1) - \frac{c}{1-\delta} \right)$, and therefore that $V_S^*(i+1) > V_S^*(i)$. In other words, if S is an equilibrium, then after a state in which C is played, the value goes down, and after a state in which D is played, the value goes up.

Now assume that strategy S is an equilibrium, and that $\lambda_S(1) = C$. If $N_S = 1$ then this is All C, which obviously is not an equilibrium. Therefore $N_S > 1$. Note that, without loss of generality, we assumed that if strategy S plays against itself, it will transition from state i to state $i + 1$, until for the first time it goes to a state it has already been in, or remains in the state it is currently in. For $i = 1$ and $\lambda_S(1) = C$ this implies that either $\mu_S(1, C) = 1$ or $\mu_S(1, C) = 2$.

If $\mu_S(1, C) = 1$ then state 2 is unused when S plays against itself, which implies that S is not an equilibrium. Therefore $\mu_S(1, C) = 2$

If $\mu_S(1, C) = 2$ and S is an equilibrium, then $c + b + \delta V_S^*(\mu_S(1, D)) \leq b + \delta V_S^*(2)$, so certainly $V_S^*(\mu_S(1, D)) < V_S^*(2)$. Also $V_S^*(2) < V_S^*(1)$. Now let $\underline{V}_S = \min_{j \in \{1, \dots, N_S\}} V_S^*(j)$ and let m be a state for which this minimum is attained; $V_S^*(m) = \underline{V}_S$. From the previous observations, we can conclude that $m > 1$ and that m is preceded by a state in which C is played. But if S is an equilibrium, then $b + \delta V_S^*(m) > c + b + \delta V_S^*(\mu_S(m - 1, D))$, which is impossible because $V_S^*(m) \leq V_S^*(i)$ for all $i = 1, \dots, N_S$.

□

With all NSS'es starting with at least one defection, we will go on to look at strategies that, when playing against themselves, start with exactly $n > 0$ defections before they play cooperate for the first time. Theorem 2 gives the perimeters of equilibrium strategies for every $n > 0$. Together with the 1-state automaton All D, that covers all possible NSS'es; it is clear that if a strategy always plays D against itself, it cannot be an NSS if it has more than 1 state.

Theorem 2. *If strategy S is an NSS, and if for $n \geq 1$*

1. $\lambda_S(i) = D$ if $1 \leq i \leq n$ and $\lambda_S(n + 1) = C$
2. $\mu_S(i, D) = i + 1$ if $1 \leq i \leq n$

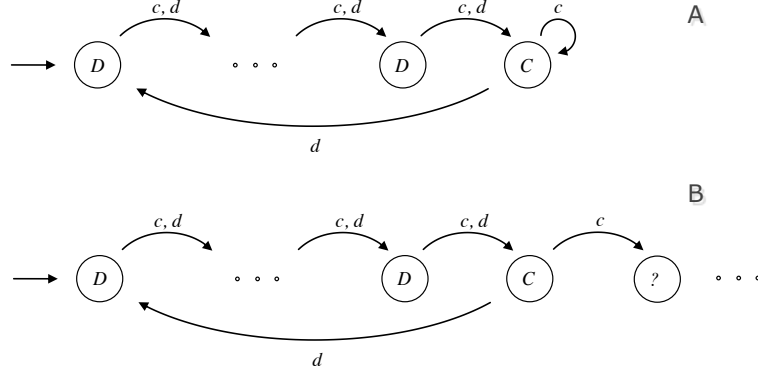


Figure 1: Two strategies that, when playing against themselves, first defect exactly n times, and then cooperate.

then

$$c + \frac{\delta^n (1 - \delta)}{1 - \delta^{n+1}} b \leq \Pi(S, S) \leq c + \delta^n (b - c)$$

Proof. If strategy S is an equilibrium, then the following two must hold.

$$(i) \quad V_S^*(1) = \sum_{i=1}^n \delta^{i-1} c + \delta^n b + \delta^{n+1} V_S^*(\mu_S(n+1, C))$$

$$(ii) \quad b + \delta V_S^*(\mu_S(n+1, C)) \geq b + c + \delta V_S^*(\mu_S(n+1, D))$$

The proof of Lemma 1 can easily be rearranged to also imply that the value in the initial state is the lowest of all values. This implies that $V_S^*(\mu_S(n+1, D)) \geq V_S^*(1)$. We can therefore replace the former with the latter in (ii).

$$\delta V_S^*(\mu_S(n+1, C)) \geq c + \delta V_S^*(1)$$

Using equation (i), we get

$$\begin{aligned}
V_S^*(1) - \frac{1 - \delta^n}{1 - \delta}c - \delta^n b &\geq \delta^n [c + \delta V_S^*(1)] \\
(1 - \delta^{n+1}) V_S^*(1) &\geq \frac{1 - \delta^n}{1 - \delta}c + \delta^n c + \delta^n b \\
(1 - \delta) V_S^*(1) &\geq c + \frac{\delta^n (1 - \delta)}{1 - \delta^{n+1}}b
\end{aligned}$$

This gives the lower bound. The upper bound follows immediately from the maximum discounted, normalized payoff that is still attainable after n mutual defections.

□

The theorem naturally implies that if $b - c < \frac{1 - \delta}{1 - \delta^{n+1}}b$, there are no equilibria that, when playing against themselves, start with exactly $n > 0$ defections before they play cooperate for the first time. As a consequence, for $\delta < \frac{c}{b}$ there are no equilibria other than All D; if $\delta < \frac{c}{b}$ then $\frac{1 - \delta}{1 - \delta^{n+1}}b > (1 - \delta)b > (1 - \frac{c}{b})b = b - c$ for all $n > 0$. On the other hand, if $b - c > \frac{1 - \delta}{1 - \delta^{n+1}}b$, then for sufficiently low costs of complexity the following strategy is an NSS, and attains the upper bound: $N_S = n + 1$; $\lambda_S(i) = D$ if $1 \leq i \leq n$ and $\lambda_S(n + 1) = C$; $\mu_S(i, j) = i + 1$ if $1 \leq i \leq n$ and $\mu_S(n + 1, D) = 1$ and $\mu_S(n + 1, C) = n + 1$ (see also Figure 1a). These strategies are NSS'es for reasons similar to those given in Cooper (1996) for $\delta = 1$. All strategies that do *not* play against themselves the way S does – that is, when playing against themselves, they do not start by playing exactly n defections before they play cooperate for the first time and ever after – earn a lower payoff against S than S does against itself. The fact that for such a strategy T we have $\Pi(T, S) < c + \delta^n(b - c) = \Pi(S, S)$ implies that for k sufficiently close to 0, we get $U(T, S) < U(S, S)$. Having more than $n + 1$ states is superfluous, so the only strategies that remain are strategies T that also have $n + 1$ states and that, when playing against themselves, do start with exactly

$n > 0$ defections before they play cooperate for the first time and ever after. For those strategies clearly $U(T, S) = U(S, S) = U(S, T) = U(T, T)$.

At the threshold, where $\frac{1-\delta}{1-\delta^{n+1}}b = b - c$, this strategy is not an equilibrium, because then $\Pi(AllD, S) = \Pi(S, S)$ while $|S| > |AllD|$.

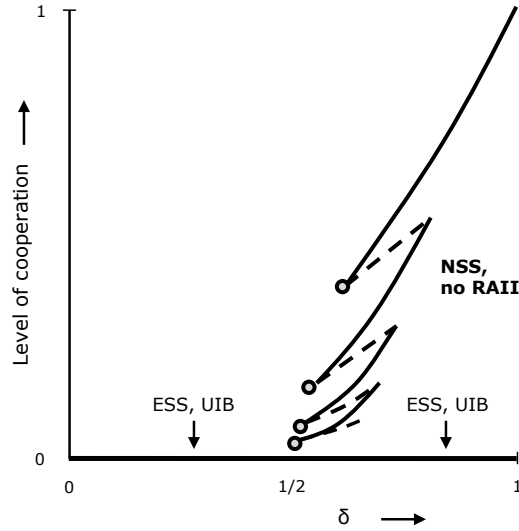


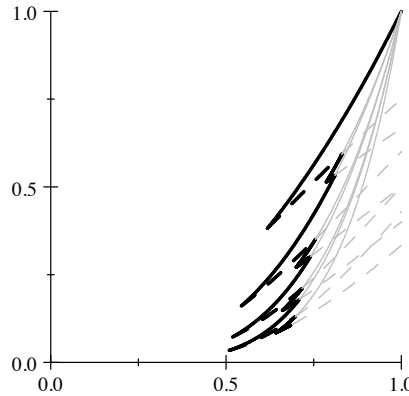
Figure 2: Equilibrium properties for $\frac{b}{c} = 2$ and $k > 0$. In the area that, roughly speaking, lies to the left of the jagged line, there are no equilibria. To the right there are many, and many are also NSS, although not every point there can be approximated arbitrarily closely with an equilibrium strategy. Also the complexity costs has to be sufficiently close to 0 to make a strategy an NSS. These NSS'es are however not RAI. All D on the other hand is an ESS for all continuation probabilities and positive complexity costs, and also has a uniform invasion barrier.

The equilibria and the perimeters are summarized for an example in Figure 2. For $b = 2$ and $c = 1$ we find that then there are no equilibria with $n = 1$ (the lower bound from Theorem 2 exceeds the upper bound for all $\delta < 1$). However, for every $n > 1$ we get a new branch on the “tree” in the picture. A few more details are worth mentioning. The first that within a

branch, not the entire area between the upper and lower bound is filled with equilibria. We can add ever more equilibria, but on the lower end of the branches, we can also carve out ever more space between the leaves. This can be done by zooming in on a branch and refining the possibilities after the first n defections and the subsequent first C . Such a procedure gives new equilibria, but also carves out extra areas without equilibria.¹ The adding of equilibria and carving out spaces without equilibria can be repeated by refining the initial sequence ever more.

The second detail that is worth mentioning is that for continuation probabilities lower than $\frac{1}{2}$ not all normalized, discounted payoffs between b and c are feasible, even if we disregard equilibrium considerations. That can easily be seen by looking at an example; if $\delta = \frac{1}{3}$ then the only payoffs that remain

¹Consider all strategies that first defect n times and then cooperate for the first time, when playing against themselves. After that, there are two possibilities: 1) they defect at least once and 2) they cooperate at least once. By looking at the restrictions on the value functions we find, rather similar to the proof of Theorem 2, that in case 1) the normalized, discounted payoff is less than or equal to $c(1 - \delta^n + \delta^{n+1} - \delta^{n+2}) + b(\delta^n - \delta^{n+1} + \delta^{n+2})$ and in case 2) at least $\frac{c(1 - \delta^n + \delta^{n+1} - \delta^{n+2}) + b(\delta^n - \delta^{n+2})}{1 - \delta^{n+2}}$. This implies that payoffs between those two values are not attainable in equilibrium. The last inequality is found by combining $V^*(1) = \frac{1 - \delta^n}{1 - \delta}c + \delta^n b + \delta^{n+1}b + \delta^{n+2}V^*(n+3)$ and $b + \delta V^*(n+3) \geq b + c + \delta V^*(\mu(n+2, D)) \geq b + c + \delta V^*(1)$. After this second step, the tree looks as follows (where the tips of branches and sub-branches again are not equilibria. This is not indicated in the picture by lack of space).



are those that can be written as $c + x(b - c)$, where x must be an element of the Cantor set.

The NSS'es however are not RAI. That is easily seen if we realize that neutral mutants can be constructed without increasing the size of the automaton. If strategy S is an equilibrium with $\Pi(S, S) > c$, we can change the transitions in the following way: define T as equal to S , except that $\mu_T(i, D) = \mu_S(i, C)$ for all i . This is a neutral mutant – $U(T, S) = U(S, S) = U(S, T) = U(T, T)$ – and it is easily seen that $\Pi(AllD, T) > \Pi(T, T)$, and therefore a sequence of first T and then $AllD$ constitutes an indirect invasion. This observation follows the same logic as the proof of the second half of Claim 1 in Volij (2002). Also here, the fact that no strategy S with $\Pi(S, S) > c$ is RAI implies that there is no ES-set other than $\{AllD\}$ (van Veelen, 2011).

The existence of neutral mutants is also a concern in Binmore and Samuelson (1992, page 294), although there the role they play is not explored by looking explicitly at whether or not they open doors for other mutants.

2.2 All D has a uniform invasion barrier

With complexity costs $k > 0$, one can show that All D, in its simple, 1-state implementation, is an ESS for all $\delta \in [0, 1]$ and even has a uniform invasion barrier.

Theorem 3. *The strategy All D in its 1-state implementation has a uniform invasion barrier for all $\delta \in [0, 1]$ and costs $k > 0$.*

Proof. All other strategies either always play D against themselves, or they do not. In the first case the matrix of fitnesses, or profits, is:

$$\begin{array}{cc}
 & \begin{array}{c} AllD \\ S \end{array} \\
 \begin{array}{c} AllD \\ S \end{array} & \begin{array}{cc} c - k & c - k \\ c - kN_S & c - kN_S \end{array}
 \end{array}$$

If S is not also All D in its 1-state implementation, then $N_S > 1$, and therefore the invasion barrier against S is 1.

If S does not always play D against itself, then there is a state $n \geq 1$ in which it plays C for the first time, if it plays against itself or against All D. That implies that against itself, it attains at most a normalized, discounted payoff of $c + \delta^{n-1}(b - c)$. Against All D, S attains at most $c - (1 - \delta)\delta^{n-1}c$, while All D against S attains at least $c + (1 - \delta)\delta^{n-1}b$. The matrix of fitnesses, or profits, now contains lower and upper bounds.

$$\begin{array}{rcc}
 & \text{All D} & S \\
 \text{All D} & c - k & \geq c + (1 - \delta)\delta^{n-1}b - k \\
 S & \leq c - (1 - \delta)\delta^{n-1}c - kN_S & \leq c + \delta^{n-1}(b - c) - kN_S
 \end{array}$$

If $c + (1 - \delta)\delta^{n-1}b - k \geq c + \delta^{n-1}(b - c) - kN_S$, then All D dominates S and the invasion barrier is 1 again. If not, then the invasion barrier is at least

$$\begin{aligned}
 \bar{\epsilon}_S &\geq \frac{(c - k) - (c - (1 - \delta)\delta^{n-1}c - kN_S)}{(c - k) - (c + (1 - \delta)\delta^{n-1}b - k) - (c - (1 - \delta)\delta^{n-1}c - kN_S) + (c + \delta^{n-1}(b - c) - kN_S)} \\
 &= \frac{(1 - \delta)\delta^{n-1}c - (N_S - 1)k}{-(1 - \delta)\delta^{n-1}b + (1 - \delta)\delta^{n-1}c + \delta^{n-1}(b - c)} \\
 &= \frac{(1 - \delta)\delta^{n-1}c + (N_S - 1)k}{\delta^n(b - c)} \\
 &= \frac{(\frac{1}{\delta} - 1)c + \frac{(N_S - 1)k}{\delta^n}}{b - c} > 0
 \end{aligned}$$

□

2.3 What if complexity costs are small?

Theorem 3 states that with complexity costs, All D, in its simplest form, has a uniform invasion barrier. Together with the finding that with complexity costs, no other strategy is RAI – which precludes the existence of other ES-sets – this suggests that if complexity costs are positive, All D is more stable than any other equilibrium. In the absence of complexity costs, on the other hand, we know that with sufficiently high continuation probability no strategy is RAI (van Veelen and García, 2010), while NSS'es exist with levels of cooperation ranging from full defection to full cooperation (Bendor and Swistak, 1995). This suggests that without complexity costs, and for $\delta < 1$, All D is neither more nor less stable than other NSS'es.

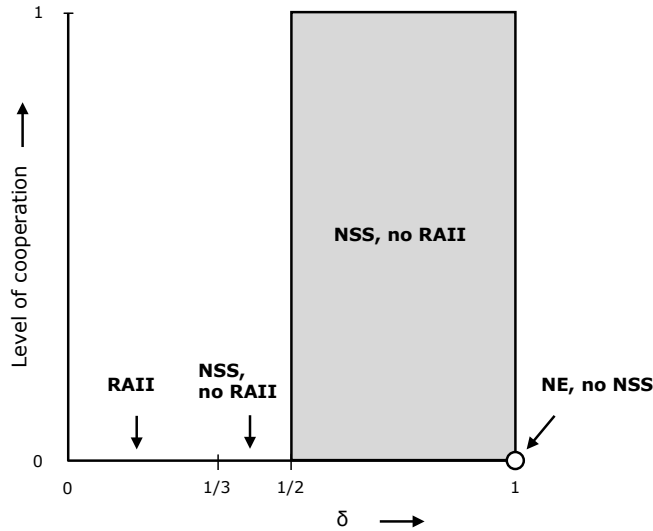


Figure 3: Equilibrium properties for $\frac{b}{c} = 2$ and $k = 0$

. See also van Veelen et al. (2012) for the equilibrium analysis left of $1/2$

Many results in the literature focus on *small* complexity costs. The results above, using static equilibrium concepts, imply that if the take the limit of complexity costs k decreasing to 0, there is a certain discontinuity; for all

$k > 0$, All D is more stable than all other NSS'es, while at $k = 0$ it is not. More precisely, the lower bound for the invasion barrier from the proof of Theorem 3 implies that if $k \downarrow 0$ the uniform invasion barrier for All D does not tend to 0, and yet at $k = 0$, there is no uniform invasion barrier, and All D is not even RAI. Moreover, there are windows of payoffs that cannot be supported by an NSS for any $k > 0$ – this is the space between “branches” in Figure 2 – while at $k = 0$ all symmetric payoffs can be approximated arbitrarily closely – see Figure 3.

The finite population dynamics on the other hand are continuous in k ; they are ergodic Markov chains, transition probabilities of which are continuous functions of k . It is therefore to be expected that the dynamics in the limit of $k \downarrow 0$ are the same as the dynamics for $k = 0$. In order to understand this discrepancy between static concepts and dynamics, we consider the effect of a decrease in k on the stability of the different equilibria.

The first two effects do not yet explain the discrepancy, even though they are reasons why a decrease in k makes All D a little less and some cooperative equilibria a little more stable. It is relatively easy to see that the lower bound on the invasion barrier for All D in Theorem 3 can actually be attained. This makes it the largest uniform invasion barrier, which implies that the closer k is to 0, the lower this lower bound. However, for $\delta < 1$ the largest uniform invasion barrier remains bounded away from 0. Only for $\delta = 1$ does the largest uniform invasion barrier actually go to 0 if k goes to 0.

The effect on cooperative equilibria on the other hand is that a decrease in k comes with an increase in abundance of equilibria that are NSS; cooperative strategies that for instance attain the upper bound in Theorem 2 can only be equilibria for sufficiently low k , and a lowering of k adds ever more equilibria.

There are however two more effects of a decrease in k , and those do explain the discrepancy. The first is that a decrease in k also destabilizes All D in another way, even for $\delta < 1$, and in spite of the fact that the largest uniform

invasion barrier is bounded away from 0 there. The reason is that the closer k is to 0, the closer to neutrality strategies are that have more states, but that also always play D against themselves. Some of them open the door for successive mutants, and therefore these “almost neutral” strategies can induce indirect invasions. By choosing a sufficiently small k we can get the first mutant arbitrarily close to neutrality. In the limit of $k \downarrow 0$ we should therefore expect indirect invasions, through almost neutral mutants, and into cooperation, to become important in the dynamics, in spite of the fact that All D has a uniform invasion barrier.

The second effect is that a decrease in k also reduces the instability of cooperative strategies that have unused punishment states, like Tit-for-tat. Such strategies cannot be equilibria as long as $k > 0$, but they can for $k = 0$. What makes Tit-for-tat a disequilibrium strategy for $k > 0$ is the presence of All C, which has less states, and therefore does strictly better than Tit-for-tat against Tit-for-tat, as long as $k > 0$. However, what is a direct invasion of All C into Tit-for-tat for all $k > 0$ gets closer and closer to being a neutral transition if $k \downarrow 0$. All C only harms cooperation because it opens the door for All D, and the deletion of unused punishment states more generally destroys cooperation through opening the door for mutants that take advantage of that. By choosing a sufficiently small k we can, again, get the first mutant arbitrarily close to neutrality. In the limit of $k \downarrow 0$ the relevant way out of Tit-for-tat therefore becomes an indirect invasion through a neutral mutant, even though All C is not neutral yet for any $k > 0$. Therefore, even though Tit-for-tat is not an equilibrium for $k > 0$, if k decreases to 0, its stability in the dynamics gets arbitrarily close to the stability in the dynamics of Tit-for-tat at $k = 0$, where Tit-for-tat is an NSS – but not RAI, just like all other NSS'es.

Simulations with varying complexity costs confirm that in the dynamics a decrease in complexity costs increases the average level of cooperation (see

Figure 4). Moreover, when k decreases to 0, the average level of cooperation seems to approach the average level of cooperation at $k = 0$ from below, suggesting that there is indeed no discontinuity at $k = 0$ in the expected level of cooperation.

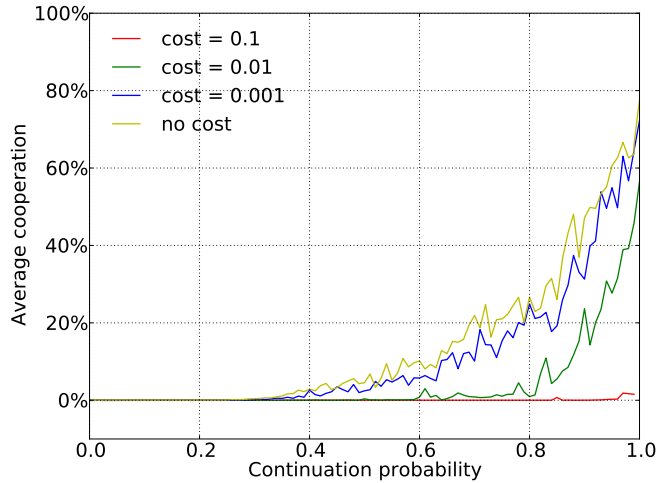


Figure 4: Levels of cooperation with different costs of complexity. The game has payoffs $P = 1$, $R = 3$, $S = 0$ and $T = 4$

The population size also plays a role. For fixed costs k , an increase in population size shifts the balance back towards what is suggested by the equilibrium analysis with positive costs (see Figure 2). This is also expected, because population size, amongst other things, magnifies the small deviations in fixation probabilities from $\frac{1}{N}$ in almost neutral transitions, and makes it relatively hard to overcome a (small) positive invasion barrier. That makes two different ways of leaving All D more difficult. Leaving All D through what would have been an indirect invasion if we disregard complexity costs is the most relevant way out for $\delta < 1$. This transition becomes relatively hard as population size increases. Direct invasions by strategies like Tit-for-tat are the most relevant ways out of All D for $\delta = 1$. For this a positive

invasion barrier has to be overcome, which also becomes increasingly hard – also in relative terms – if population size increases.

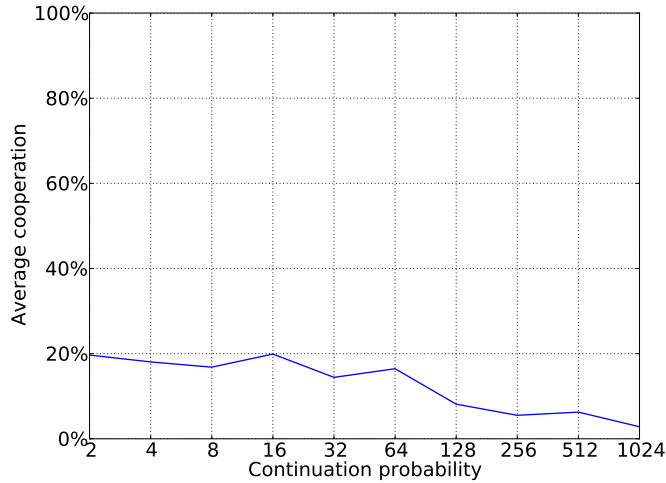


Figure 5: Levels of cooperation with increasing population size and fixed costs and continuation probability. Complexity costs k are 0.01, $\delta = 0.85$

If we compare the setting without complexity to the one with (small) complexity costs, it is not unintuitive that adding complexity costs would harm average cooperation a little. If we start with the case without complexity costs, then stepping stone paths into cooperation typically require strategies with increased complexity, while paths out of cooperation do not, and often even involve a decrease in complexity. Adding just a little cost of complexity makes those paths into cooperation become a little less likely, while the suggested paths out of cooperation become a little more likely. Simulation results in Ho (1996) also suggest that including costs that depend on the size of an automaton is not good for the level of cooperation. There however the runs are only a 100 generations long, and it seems that the runs do not contain enough transitions to draw strong conclusions.

3 Infinitely repeated games

Games that are repeated indefinitely – $\delta = 1$ – have received extra attention in the literature. Also in our analysis $\delta = 1$ is an interesting case. What it shares with the case of $\delta < 1$ is that the analysis suggests that All D is *more* stable than cooperative equilibria at $k > 0$. What is different is that All D is *less* stable than cooperative equilibria at $k = 0$, because in the latter case All D can directly be invaded, and therefore it is not even an NSS. Thereby it differs from the situation at $\delta < 1$, where All D becomes an NSS that is not RAI at $k = 0$, which makes it not less and not more stable than the cooperative strategies that are also NSS but not RAI.

In order to better understand the dynamics at $\delta = 1$, and in order to link RAI with existing equilibrium concepts, we now reproduce the definitions of a MESS from Binmore and Samuelson (1992) and of a LESS from Volij (2002). We will see that the three stability concepts – MESS, LESS and RAI – focus on different transitions. As we have already seen above, the parameters k and N manipulate the relative importance of different transitions, which implies that the results of the simulations can be well understood if all relevant types of transitions are considered.

In the definition there are again two objectives; payoffs from the game, and complexity costs. Note that here the set of mutants that T can be chosen from is also the set of finite automata, not the set of all strategies.

Definition 4. MESS (Binmore and Samuelson, 1992)

A finite automaton S is a MESS if for any mutant T

1) $\Pi(S, S) \geq \Pi(T, S)$

2) *if* $\Pi(S, S) = \Pi(T, S)$ *then* $\Pi(S, T) \geq \Pi(T, T)$

3) *if* $\Pi(S, S) = \Pi(T, S)$ *and* $\Pi(S, T) = \Pi(T, T)$ *then* $|S| \geq |T|$

An automaton S is a strict MESS if one of the inequalities is strict

(Samuelson and Swinkels, 2003)². With the sets of better, equal and worse performers, defined in van Veelen (2011) and in the companion paper, one can also give an equivalent and shorter definition.

Definition 5. *An finite automaton S is a MESS if*

- 1) $S_B(S) = \emptyset$ and
- 2) If $T \in S_E(S)$ then $|T| \geq |S|$

The second definition is:

Definition 6. LESS (Volij, 2002)

A finite automaton S is a LESS if for any mutant T

- 1) $\Pi(S, S) \geq \Pi(T, S)$
- 2) if $\Pi(S, S) = \Pi(T, S)$ then $|S| \geq |T|$
- 3) if $\Pi(S, S) = \Pi(T, S)$ and $|S| = |T|$ then $\Pi(S, T) > \Pi(T, T)$

One could define a neutral LESS by allowing the last inequality to be non-strict. The difference between these concepts – and especially between a *strict* MESS and a LESS – is explained clearly in Samuelson and Swinkels (2003).

For symmetrized, infinitely repeated games, where payoffs are evaluated according to the limit of means, Binmore and Samuelson (1992) find that strategies that are a MESS must be *utilitarian* (a utilitarian strategy playing against itself maximizes the sum of payoffs in the underlying game). Applied to standard prisoners dilemmas, where the sum of payoffs is maximized if both cooperate, this implies that, after an initial phase, a MESS always cooperates, when playing against itself.

²Samuelson and Swinkels (2003) also allow for mixed strategies in their definition of a MESS and LESS. They also indicate that they follow the standard practice (cf. Rubinstein (1986) and Abreu and Rubinstein (1988)) of working with pure MESS or LESS to avoid difficulties in interpreting the complexity of mixed strategies. We do the same, also because the main results from Binmore and Samuelson (1992) and Volij (2002) concern pure strategies.

The main result in Volij (2002, Claim 1) is that only $AllD$ is a LESS, and that there is no other ES-set other than $\{AllD\}$, if the definition of an ES-set does not look just at payoffs, but instead uses lexicographic preferences where complexity matters after payoffs.³

In this paper we want to establish links with dynamics for repeated games. It might therefore be good to first start with what the dynamics would be with a simple, restricted strategy set, and see how complexity costs affect those. We therefore follow Imhof, Fudenberg and Nowak (2005) and consider a strategy space consisting of three strategies; All D, All C and Tit-for-tat. The setting in Imhof, Fudenberg and Nowak (2005) is slightly different,⁴ but what they find does carry over to our setting. With those three strategies, the population will move in circles, going from All D to Tit-for-tat, to All C, and back to All D. More importantly, we can use complexity costs and population size to manipulate the dynamics such that the population is either cooperating almost all of the time or defecting almost all of the time. The key is in the likelihood of the transition from All D to Tit-for-tat. With

³Volij (2002) actually shows that in the repeated prisoners dilemma, there is no evolutionary stable family of strategies (ESF) other than the singleton set $\{AllD\}$. An ESF as defined in Volij (2002) is a (Balkenborg & Schlag) ES-set, and therefore also a (Thomas) ES set (see Thomas, 1985; Balkenborg and Schlag, 2001; van Veelen, 2011) but the reverse is not true; in the following game, the set of all mixtures of the first two strategies is an ES-set, but not an ESF.

$$\begin{bmatrix} 1 & 2 & 0 \\ 1 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

The proof of Claim 1 in Volij (2002) however also works perfectly well for showing that there is no ES-set other than the singleton set $\{AllD\}$.

⁴They have a finite number of repetitions, over which payoffs are added, and where Tit-for-tat incurs a fixed complexity cost, independent on the number of repetitions. That implies that, compared to our setting, an increase of the number of repetitions reduces the relative importance both of the loss that Tit-for-tat incurs from its initial cooperation with All D and of the cost of complexity. Increasing the number of repetitions is therefore equivalent to bringing δ closer to 1 and k closer to 0 at the same time. See also Fudenberg and Imhof (2008) for a proof of the main result in Imhof, Fudenberg and Nowak (2005); apply Theorem 1 from the former to the appropriate strategy space.

complexity costs, the game between those two strategies is a coordination game. The size of the basin of attraction of All D depends on complexity costs; if k goes to 0, then the size of the basin of attraction goes to 0. For a fixed population size, that means that the basin of attraction can even be reduced to below $\frac{1}{N}$, which makes every single mutation push the population out of the basin of attraction of All D and into the basin of attraction of Tit-for-tat already. For a fixed k , however, if we increase the population size, then making it from All D into the basin of attraction of Tit-for-tat becomes very unlikely. The variables k and N can therefore be used to make the transition from All D to Tit-for-tat much more or much less likely than the transition from Tit-for-tat to All C. The latter transition is, for k close to 0, almost neutral; although All C strictly dominates Tit-for-tat, the identical self-play implies that the difference in fitness is only due to the (small) difference in complexity costs. If going from All D to Tit-for-tat is much more likely than from Tit-for-tat to All C – while going from All C to All D is the most likely transition – then more time is spent in Tit-for-tat than anywhere else. If going from All D to Tit-for-tat is much less likely than from Tit-for-tat to All C, then more time is spent in All D than anywhere else.

With a larger strategy set there are more transitions to consider, but the tuning of the relative likelihood of these three will still play an important role.

If we look at the difference between a MESS and a LESS, then we see that, provided that a strategy is a Nash equilibrium, the performance in self-play matters first for being a MESS, and complexity matters first for being a LESS. That makes MESS align with relatively low k and low N and LESS with relatively high k and high N (or with relatively small payoff differences). Also more in general, transitions with increasing cooperation typically require increasing complexity, as the more cooperative strategies need an exploration phase in order to determine which strategy they are

playing against, a part for self-play and a part for playing against the other (see Binmore and Samuelson 1992, page 292). The increase in cooperation in self-play is more important than the increase in complexity in the concept of a MESS and less important in the concept of a LESS. Again, such a transition is relatively likely for low k and low N and relatively unlikely for high k and high N . These two concepts therefore explain why cooperation is low for high k and high N and high for low k and low N , as it is in Imhof, Fudenberg and Nowak (2005).

With this in mind, one might expect that, for a fixed population size N , the average level of cooperation can be pushed down to 0 by increasing complexity costs, because All D is the only LESS. Moreover, one might also expect that in the limit of $k \downarrow 0$ the average level of cooperation will be 1, because there all MESS'es are efficient (and the set of MESS'es is not empty). The first is true, but the latter is not. If we look at the level of cooperation in simulations for fixed N , $\delta = 1$ and varying k , we find that the level of cooperation is close to 0 already for $k = 0.1$. For $k \downarrow 0$, on the other hand, we do *not* find that the level of cooperation goes to 1 (see Figure 4); the simulations suggest that at $k = 0$ the average level of cooperation is a value smaller than 1, which is approached from below if k decreases to 0. Why that is, can be understood by looking at the dynamics at $k = 0$, provided that we consider a part of the web of transitions that includes strategies with unused states. Strategies with unused states obviously cannot be a MESS or a LESS, nor an NSS if we have explicit complexity costs $k > 0$. Such strategies are therefore ruled out at the onset by the definitions of a MESS, LESS and an NSS with explicit complexity costs. This has an important implication in the case of infinitely repeated games, where transient states do not matter for the payoff. In infinitely repeated games, and in the absence of unused states, any strategy can be figured out, and optimally responded to, while a strategy can also cooperate fully when it figures out that it is

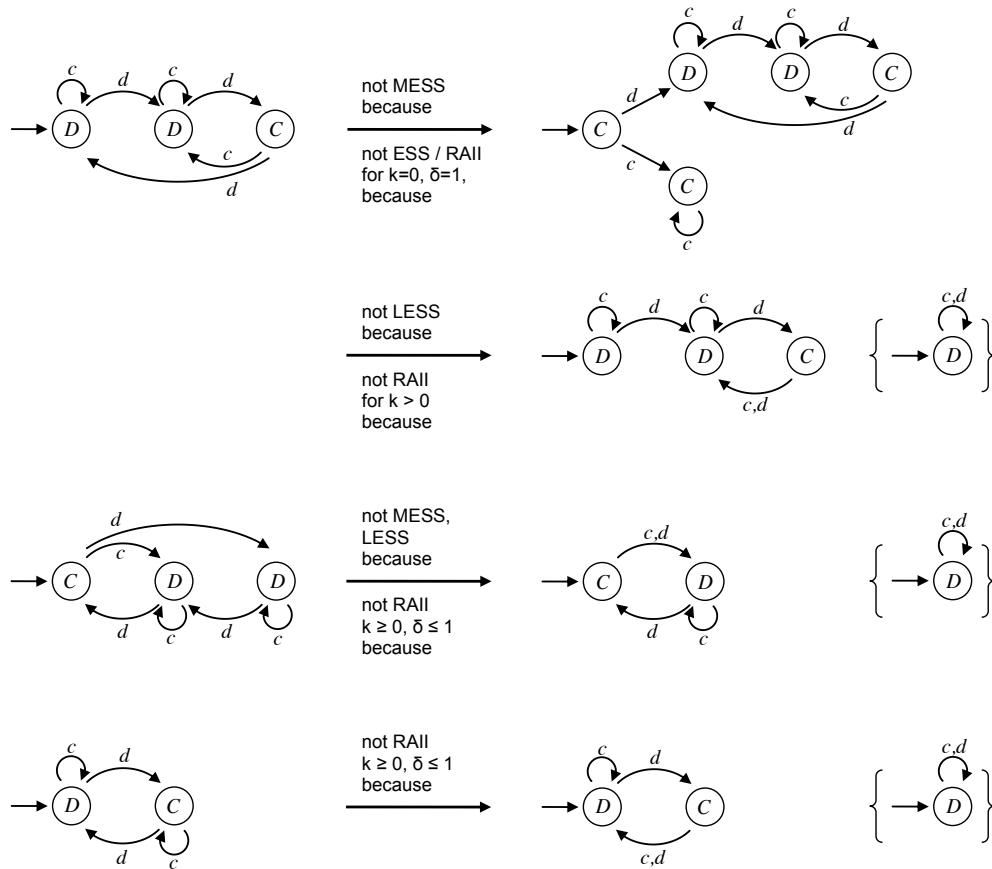


Figure 6: Three strategies with mutants that make them not satisfy definitions of different static stability concepts. The first two are not fully cooperative at $\delta = 1$, the third is. With complexity costs k decreasing to 0, the first transition away from the first strategy becomes ever more likely in finite population dynamics. The transition away from the second strategy becomes less likely with decreasing k . The fixation probability of the mutant that leads away from it becomes arbitrarily close to the fixation probability of the mutant that leads away from the third, fully cooperative, strategy. In the limit of $k \downarrow 0$ these two strategies become equally (in)stable in the finite population dynamics.

playing against a copy of itself (see Binmore & Samuelson, 1992). Invasions that exploit that possibility become more likely as complexity costs go to 0. The concept of a MESS captures this transition out of strategies that are not fully cooperative, and that do not have unused punishment states, while strategies with unused states are ruled out at the onset.

However, as complexity costs go to 0, the disadvantage of having unused states also goes to 0. In the limit of k going to 0, strategies that do not cooperate fully, when playing against themselves, and that do have punishment states that are *not* used in self-play, thereby become as stable as they are for $k = 0$. Some of these strategies cannot be figured out without triggering everlasting punishment. That explains why we do not achieve full cooperation if we let k go to 0.

For $k \downarrow 0$ the transition probabilities in the entire web of transitions go to the transition probabilities at $k = 0$. On the one hand some not fully cooperative strategies – ones that do not have unused punishment states – have become disequilibrium states, because at $k = 0$ it is free to build a machine that can figure them out. On the other hand, some other not fully cooperative strategies – ones that do have unused punishment states; see Figure 6 – move in the opposite direction; at $k = 0$ they have become equilibria, more or less getting in through the back door. Furthermore, at $k = 0$ we also know that no strategy is RAI, so the remaining NSS'es are all susceptible to indirect invasions, the fully cooperative ones as well as the not fully cooperative ones (while the fully defecting ones are not even NSS'es). Since the dynamics for $k \downarrow 0$ go to the dynamics at $k = 0$, this explains why we do not see average behaviour converge to full cooperation if $k \downarrow 0$. With small complexity costs, transitions that would be indirect invasions without complexity costs therefore turn out to complete the picture of the dynamics at $\delta = 1$.

4 Conclusion and discussion

If we include complexity costs, we find that there are strategies with positive levels of cooperation that are NSS'es for sufficiently low complexity costs and sufficiently high continuation probability. This extends a folk theorem result in Cooper (1996) for infinitely repeated prisoners dilemmas without discounting to repeated prisoners dilemmas with discounting, or, in our interpretation, with continuation probabilities smaller than 1. The difference between the result in Cooper (1996) and the extension is that in the former all levels of cooperation can be approximated arbitrarily closely by a strategy that is an NSS for sufficiently low costs of complexity, while this is not the case anymore for continuation probabilities (or discount factors) smaller than 1. The NSS'es with positive levels of cooperation are however not robust against indirect invasions – which was also the case without complexity costs – while All D, in its 1-state implementation, is an ESS, and even has a uniform invasion barrier, which was not the case without complexity costs. The simulations also suggest that the dynamical behaviour in the limit of complexity costs going to 0 equals the dynamical behaviour at complexity costs 0, where indirect invasions drive the dynamics for large enough population size. Furthermore, high continuation probability and low costs of complexity are good for cooperation, while large population sizes are bad for cooperation. The simulation results for the case of infinitely many repetitions and no discounting, or for continuation probabilities close to 1, can be understood by looking at typical invasions considered by the concepts of a MESS (Binmore and Samuelson, 1992) and a LESS (Volij, 2002) as well as RAI. With small complexity costs, neutral and almost neutral mutants that open doors for mutants with a selective advantage complete the picture of the relevant transitions.

References

- Abreu, D.** 1988. "On the theory of infinitely repeated games with discounting." *Econometrica*, 56(2): 383–396.
- Abreu, D., and A. Rubinstein.** 1988. "The structure of Nash equilibrium in repeated games with finite automata." *Econometrica: Journal of the Econometric Society*, 1259–1281.
- Aumann, R., and L. Shapley.** 1976. "Long term competition: a game theoretic analysis." *Mimeo Hebrew University*.
- Axelrod, R.** 1984. *The Evolution of Cooperation*. New York:Basic Books.
- Axelrod, R., and W. D. Hamilton.** 1981. "The Evolution of Cooperation." *Science*, 211: 1390–1396.
- Balkenborg, D., and K.H. Schlag.** 2001. "Evolutionarily stable sets." *International Journal of Game Theory*, 29(4): 571–595.
- Bendor, J., and P. Swistak.** 1995. "Types of evolutionary stability and the problem of cooperation." *Proc. Natl. Acad. Sci. USA*, 92: 3596–3600.
- Bendor, J., and P Swistak.** 1997. "The evolutionary stability of cooperation." *American Political Science Review*, 91(2): 290–307.
- Bendor, J., and P Swistak.** 1998. "Evolutionary equilibria: Characterization theorems and their implications." *Theory and decision*, 45(2): 99–159.
- Binmore, K., and L Samuelson.** 1992. "Evolutionary stability in repeated games played by finite automata." *Journal of economic theory*, 57(2): 278–305.

- Boyd, R., and J.M.D. Lorberbaum.** 1987. “No Pure Strategy is Evolutionary Stable in the Iterated Prisoner’s Dilemma Game.” *Nature*, 327: 58–59.
- Cooper, D.J.** 1996. “Supergames played by finite automata with finite costs of complexity in an evolutionary setting.” *Journal of Economic Theory*, 68(1): 266–275.
- Farrell, J, and R Ware.** 1989. “Evolutionary stability in the repeated Prisoner’s Dilemma.” *Theoretical population biology*, 36(2): 161–166.
- Friedman, J W.** 1971. “A non-cooperative equilibrium for supergames.” *The Review of Economic Studies*, 38(1): 1–12.
- Fudenberg, D., and E. Maskin.** 1986. “The folk theorem in repeated games with discounting or with incomplete information.” *Econometrica*, 54(3): 533–554.
- Fudenberg, D., and E. Maskin.** 1990. “Evolution and cooperation in noisy repeated games.” *American Economic Review*, 80: 274–279.
- Fudenberg, D., and L. A. Imhof.** 2008. “Monotone imitation dynamics in large populations.” *Journal of Economic Theory*, 140: 229–245.
- Ho, T.H.** 1996. “Finite automata play repeated prisoner’s dilemma with information processing costs.” *Journal of economic dynamics and control*, 20(1-3): 173–207.
- Imhof, L. A., D. Fudenberg, and M. A. Nowak.** 2005. “Evolutionary cycles of cooperation and defection.” *Proc. Natl. Acad. Sci. USA*, 102: 10797–10800.
- Kalai, E., and W. Stanford.** 1988. “Finite rationality and interpersonal complexity in repeated games.” *Econometrica*, 397–410.

- Kim, Y.G.** 1994. "Evolutionarily stable strategies in the repeated prisoner's dilemma." *Mathematical Social Sciences*, 28(3): 167–197.
- Linster, B.G.** 1992. "Evolutionary stability in the infinitely repeated Prisoners' Dilemma played by two-state Moore machines." *Southern Economic Journal*, 880–903.
- Linster, B.G.** 1994. "Stochastic Evolutionary Dynamics in the Repeated Prisoners' dilemma." *Economic Inquiry*, 32(2): 342–357.
- Lorberbaum, J P, D E Bohning, A Shastri, and L E Sine.** 2002. "Are there really no evolutionarily stable strategies in the iterated prisoner's dilemma?" *Journal of Theoretical Biology*, 214(2): 155–169.
- May, Robert M.** 1987. "More Evolution of Cooperation." *Nature*, 327: 15–17.
- Miller, J.** 1996. "The coevolution of automata in the repeated Prisoner's Dilemma." *Journal of Economic Behavior & Organization*, 29(1): 87–112.
- Rubinstein, A.** 1979. "Equilibrium in supergames with the overtaking criterion." *Journal of Economic Theory*, 21(1): 1–9.
- Rubinstein, A.** 1980. "Strong perfect equilibrium in supergames." *International Journal of Game Theory*, 9(1): 1–12.
- Rubinstein, A.** 1986. "Finite automata play the repeated prisoner's dilemma." *Journal of economic theory*, 39(1): 83–96.
- Samuelson, L, and J M Swinkels.** 2003. "Evolutionary stability and lexicographic preferences." *Games and Economic Behavior*, 44(2): 332–342.

- Selten, R, and P Hammerstein.** 1984. “Gaps in Harley’s argument on evolutionarily stable learning rules and in the logic of “tit for tat”.” *Behavioral and Brain Sciences*, 7(01): 115–116.
- Thomas, B.** 1985. “On evolutionarily stable sets.” *Journal of mathematical Biology*, 22(1): 105–115.
- van Damme, E.** 1989. “Renegotiation-proof equilibria in repeated prisoners’ dilemma.” *Journal of Economic Theory*, 47(1): 206–217.
- van Veelen, M.** 2011. “Robustness against indirect invasions.” *Games and Economic Behavior*, 74(1): 382–393.
- van Veelen, M., and J. García.** 2010. “In and Out of Equilibrium: Evolution of Strategies in Repeated Games with Discounting.” *TI Working Paper*, 10-037/1.
- van Veelen, M., J. García, D.G Rand, and M. A. Nowak.** 2012. “Direct reciprocity in structured populations.” *Proc. Natl. Acad. Sci. USA*, 109: 9929–9934.
- Volij, O.** 2002. “In defense of DEFECT.” *Games and Economic Behavior*, 39(2): 309–321.

A Acknowledgements

We would like to thank Gerard van der Laan, Maus Sabelis, Martin Nowak, Corina Tarnita, Dave Rand, Arthur Schram, Peter Spreij, Aljaz Ule, Drew Fudenberg, Jörgen Weibull and other participants of the Tinbergen Institute 2010 conference. Both authors gratefully acknowledge financial support by the Netherlands’ Organization for Scientific Research (NWO).