Available online at www.sciencedirect.com
SCIENCE DIRECT
Journal of Theoretical Biology 239 (2006) 183–194
Journal of Theoretical Biology
www.elsevier.com/locate/yjtbi

The good, the bad and the discriminator—Errors in direct and indirect reciprocity

Hannelore Brandt, Karl Sigmund*

Fakultaet fuer Mathematik, University of Vienna and IIASA, Laxenburg, Austria

Received 11 February 2005; received in revised form 28 April 2005; accepted 23 July 2005

Abstract

This paper presents, in a series of simple diagrams, concise results about the replicator dynamics of direct and indirect reciprocity. We consider repeated interactions between donors and recipients, and analyse the relationship between three basic strategies for the donor: unconditional cooperation, all-out defection, and conditional cooperation. In other words, we investigate the competition of discriminating and indiscriminating altruists with defectors. Discriminators and defectors form a bistable community, and hence a population of discriminators cannot be invaded by defectors. But unconditional altruists can invade a discriminating population and 'soften it up' for a subsequent invasion by defectors. The resulting dynamics exhibits various forms of rock-paper-scissors cycles and depends in subtle ways on noise, in the form of errors in implementation. The probability for another round (in the case of direct reciprocity), and information about the co-player (in the case of indirect reciprocity), add further elements to the ecology of reciprocation.

Keywords: Evolutionary game theory; Replicator dynamics; Cooperation; Reputation; Prisoner's Dilemma

1. Introduction

Among the rich variety of topics treated by Maynard Smith, reciprocal altruism takes a relatively narrow place. The most explicit treatment can be found in the last chapter, and the last appendix, of his seminal book on 'Evolution and the Theory of Games' (Maynard Smith, 1982), as well as in a target paper written for *Brain and Behavioral Science* (Maynard Smith, 1984). In a commentary to that paper, Selten and Hammerstein (1984) criticized that Maynard Smith had rashly adopted the claim of Axelrod and Hamilton (1981) that tit for tat (TFT), the reciprocal strategy par excellence, is evolutionarily stable.

Indeed, Maynard Smith did take some liberty with his own definition of an evolutionarily stable strategy (or ESS). In that definition (see Maynard Smith, 1982), he had explicitly stated that for a strategy X to be an ESS, it must (a) be a best reply to itself (i.e. a Nash equilibrium) and (b) if Y were any alternative best reply, X should be a strictly better reply to Y than Y itself. Indeed, if this second condition were not satisfied, Y could invade through neutral drift.

As Maynard Smith explicitly showed in his appendix (Maynard Smith, 1982), the strategy AllC (unconditional cooperation) is an alternative best reply to TFT, and both strategies fare equally well against each other. Hence TFT is no ESS for the iterated Prisoner's Dilemma game. This is not only a mathematical pedantry. If unconditional altruists can spread, defectors can eventually invade and ultimately take over. Moreover, Selten and Hammerstein (1984), just as Axelrod and Hamilton (1981), stressed rightly that if players are only boundedly rational, an erroneous move in the iterated Prisoner's Dilemma can lead to a long, payoff-reducing vendetta between two TFT players.

In this paper, we will investigate the interplay of defectors with conditional and unconditional altruists, placing particular emphasis on the role of errors. We shall analyse this in the context of evolutionary game dynamics

*Corresponding author.
E-mail address: karl.sigmund@univie.ac.at (K. Sigmund).

0022-5193/$ - see front matter
doi:10.1016/j.jtbi.2005.08.045

(see Hofbauer and Sigmund, 1998) for both direct and indirect reciprocity, i.e. for the repeated Prisoner's Dilemma game against the same or against varying co-players.

Humans are certainly supreme reciprocators. Most the examples of reciprocity-based collaboration in other species have raised numerous objections and failed to gain universal acceptance (Dugatkin, 1997; Hammerstein, 2003). All other eusocial species achieving high levels of cooperation are based on kin-selection, to a much larger extent than we are. Since Maynard Smith made a point in professing that 'for [him], the human applications of sociobiology are peripheral' (Maynard Smith, 1988, Chapter 7), his relative neglect of reciprocity may simply have been due to the fact that he saw little evidence for it among species other than *homo reciprocans*.

In contrast, experimental economics increasingly highlights the fact that the success of our species is based on our ability to treat non-relatives, and even complete strangers, as 'honorary relatives' (to use a felicitous phrase, cf. Seabright, 2004). We seem to have a special aptitude for reciprocal interactions with our conspecifics. This tendency may well be a human universal (Gintis et al., 2003; Fehr and Fischbacher, 2003).

Reciprocal interactions are based on the principle of just return (Axelrod and Hamilton, 1981). It works if a helpful action, or a gift, is returned to the donor by the recipient. Over a period of time, such mutual support can lead to a benefit for both parties involved (Trivers, 1971). But next to this so-called direct reciprocity, one also finds, at least among humans, instances of indirect reciprocity: the return is provided, not by the recipient, but by a third party (Trivers, 1971; Alexander, 1987; Lotem et al., 1999; Wedekind and Milinski, 2000; Wedekind and Braithwaite, 2002). We note that strong reciprocity also belongs in this context: humans tend to punish wrongdoers, even if this involves a cost to themselves, and even if they are mere bystanders rather than the victims of the wrongdoer. In fact, experiments have shown that strong reciprocity and indirect reciprocity go a long way towards explaining human behaviour in public goods games (Milinski et al., 2002a, b; Fehr and Fischbacher, 2003).

It is clear that direct and indirect reciprocity share many common features. In particular, the so-called folk theorem on repeated games applies equally well to both cases (we will come back to this in the discussion). However, there are also many subtle differences. In this paper, we propose to compare the replicator dynamics in the two cases, restricting attention to the three most basic strategies: to cooperate, to defect, or to discriminate. Needless to say, there are many other possible strategies, and some play probably an important role. Nevertheless, we believe that the interplay of these three particular rules captures an essential aspect of the evolutionary dynamics of cooperation, and of our instinct for reciprocation. Thus we propose to investigate the logic of reciprocation by analysing the relationship of the most basic conditional strategy (do whatever the co-player did), with the two extreme unconditional strategies, those of indiscriminating altruism and all-out defection.

2. The modelling background

All interactions which we consider involve two players, one in the role of the donor, the other in the role of the recipient. The donor can confer a benefit b to the recipient, at a cost $-c$ to the donor. Thus the donor can decide whether to cooperate or to defect. We shall always assume $0<c<b$, and use the terms 'donor' and 'recipient' even if the donor refuses to donate.

We will consider repeated games. In the case of direct reciprocity, the same two players interact round after round with each other. For convenience, we shall assume that in each round, each of the two players is in both roles, and that both players have to decide simultaneously, without knowing what the other will do. In the case of indirect reciprocity, each player will be matched with a different co-player in each round. In fact, since we want to keep the parallel as close as possible, we shall again assume that the player, in each round, plays both roles (donor and recipient) and is matched against two co-players. (Alternatively, we could imagine that the player is, in each round, with the same probability in the role of the donor or the recipient. This introduces no essential change, cf. Nowak and Sigmund, 1998a.)

Let us assume, as usual, that after every round another round can occur with a constant probability $w\leqslant 1$. We number the initial round by 0 and the nth iteration by n. The probability that there will be at least n iterations is given by w^n, the probability that there are exactly n iterations by $w^n(1-w)$. In that case, the game will consist of exactly $n+1$ rounds (the first round, and then n iterations). The length of the game will be a random variable, its expectation value is $1(1-w)+2w(1-w)+\cdots+nw^{n-1}(1-w)+\cdots$ which sums up to $(1-w)^{-1}$.

If we denote by $A(n)$ the payoff in the nth round, we obtain in the case $w<1$ as expected value of the total payoff the sum

$$\sum_{n=0}^{+\infty} w^n(1-w)[A(0)+\cdots+A(n)], \tag{1}$$

which by using Abel's summation formula is $A(0)+wA(1)+\cdots$. Since all $A(n)$ are uniformly bounded, this sum always converges for $w<1$ to some value $A(w)$. The average payoff per round is given by

$$(1-w)A(w)=(1-w)^2\sum_{n=0}^{+\infty} w^n \quad [A(0)+\cdots+A(n)]. \tag{2}$$

It is often convenient to consider the limiting case $w=1$. In this case, there is always another round, the game consists of infinitely many rounds and the total payoff $\sum_n A(n)$ may diverge. It is convenient, instead, to consider the average (over time) of the payoff *per round*, i.e. the limit, for $n \to +\infty$, of

$$\frac{A(0) + \cdots + A(n)}{n+1}, \tag{3}$$

provided it exists. The theorem of Frobenius implies that in this case, the limit of the time averages is just $\lim_{w\to 1}(1-w)A(w)$.

We shall consider the interaction of three strategies only. The cooperator always decides to donate, when in the role of the donor; the defector never donates; and the discriminator donates under conditions that will be specified in the two cases of direct and indirect reciprocation considered below. Cooperators and discriminators are also called indiscriminating and discriminating altruists.

We consider a large, well-mixed population. The frequencies of the three strategies (cooperator, defector, discriminator) are given by x, y and z, respectively (with $x+y+z=1$). With P_x, P_y and P_z we denote the expected values for the total payoff obtained by these strategies, and by $\bar{P} = xP_x + yP_y + zP_z$ the average payoff in the population. We shall assume that the frequencies of the strategies change with time, such that more successful strategies increase in frequency. For instance, we may assume that from time to time, players can compare their payoff with that of another player chosen at random in the population, and imitate the strategy of that player if it is more successful. If we assume that the probability for a switch is proportional to the payoff difference, the evolution of the frequencies of the strategies in the population is given by the replicator equation

$$\dot{x} = x(P_x - \bar{P}),$$

$$\dot{y} = y(P_y - \bar{P}),$$

$$\dot{z} = z(P_z - \bar{P}) \tag{4}$$

(see e.g. Hofbauer and Sigmund, 1998). Many other dynamics show a similar behaviour. We will frequently use the fact that the replicator equation remains unchanged (in the simplex S_3) if the same function is added to each payoff term, and by abuse of notation still design them with P_x, P_y, P_z and $\bar{P}$. In particular, we can normalize the payoff matrix by adding an appropriate constant to each column. We recall that the Nash equilibria are exactly those fixed points which are saturated (i.e. if $x=0$ then $P_x \leqslant \bar{P}$ etc).

3. Direct reciprocity

The cooperator, defector and discriminator, for the case of direct reciprocation, are also known as AllC, AllD and TFT (tit for tat) player. The latter cooperates in the first round and then does whatever the co-player did in the previous round.

AllD against AllD has payoff $A(n) = 0$ in every round, so that $A(w) = 0$. A TFT player against an AllD player earns $A(0) = -c$ and, for $n \geqslant 1$, $A(n) = 0$, so that $A(w) = -c$, etc.

The payoff matrix for the three strategies AllC, AllD and TFT is, omitting the factor $(1-w)^{-1}$, (i.e. considering the payoff *per round*)

$$M = \begin{pmatrix} b-c & -c & b-c \\ b & 0 & b(1-w) \\ b-c & -c(1-w) & b-c \end{pmatrix}. \tag{5}$$

Let us normalize the corresponding replicator equation such that P_y, the payoff for defectors, is 0. Then we obtain

$$P_x = -c + wbz, \quad P_z = P_x + wcy. \tag{6}$$

We note that $P_z - \bar{P} = yg$, with

$$g = w(b-c)z - c(1-w). \tag{7}$$

On the edge with $z=0$, AllD clearly wins. On the edge with $x=0$, i.e. in a population consisting of defectors and TFT-players, we have a bistable dynamics. The unstable equilibrium is $F_{yz} = (0, 1-\hat{z}, \hat{z})$, with

$$\hat{z} = \frac{(1-w)c}{w(b-c)}. \tag{8}$$

Since $\hat{z}$ is small if w is close to 1, this means that a small TFT-cluster is able to invade a population of defectors if w, i.e. the 'shadow of the future' is sufficiently large (Axelrod and Hamilton, 1981). The edge $y=0$ consists of fixed points only. Clearly, a population of AllC and TFT players will always cooperate, and none of the two strategies is favoured. On the edge $y=0$, those points with $z \geqslant c/wb$ are Nash equilibria, and the others are not. To see this, we have only to look at the sign of $P_y - \bar{P}$, i.e. of $P_x = -c + wbz$. The other Nash equilibria are the corner $y=1$ (defectors only) and F_{yz}. In the interior of the simplex, there is no fixed point. Indeed, we see that $P_x = P_y (=0)$ holds for the points on the line $g=0$, and that there, P_z is positive. The segment with $g=0$ consists of a single orbit parallel to the edge $z=0$, which converges to the saddle point F_{yz} and separates the simplex into two parts.

It is easy to see that the function

$$V = x^{\frac{1-w}{w}} z^{-\frac{1}{w}} g \tag{9}$$

is an invariant of motion.

In the case $c < wb$, the dynamics shows an interesting behaviour (see Fig. 1). In the absence of defectors, any mixture of TFT-players (i.e. discriminating altruists) and AllC players (indiscriminating altruists) are in equilibrium, and we have to assume that random shocks send the system up and down the defectors-free edge $y=0$. If a random shock introduces a small amount of defectors while $z > c/wb$, the defectors will forthwith be eliminated. If the defectors are introduced while $z < (1-w)c/w(b-c)$, they will take over. But if the defectors are introduced in the 'middle zone' where

$$c/wb > z > (1-w)c/w(b-c), \tag{10}$$

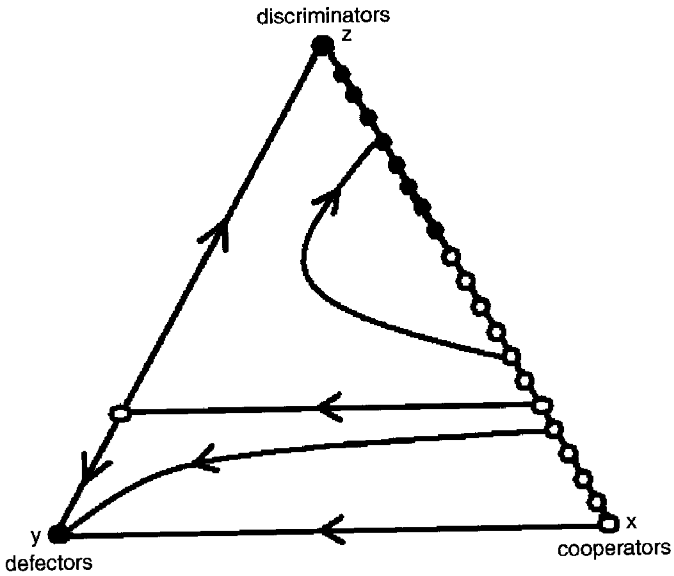

Fig. 1. The replicator dynamics of direct reciprocity in the absence of errors, assuming a constant probability $w<1$ for a further round. Here and in the other figures, full circles correspond to stable fixed points, and empty circles to unstable fixed points (stability being understood in the sense of Lyapunov: all close-by states remain close-by). We note that fixed points that are stable are Nash equilibria, but that the converse does not hold. The same dynamics shows up in the case of indirect reciprocity with a fixed number of synchronous rounds and no errors ($\varepsilon = 0$).

the amount of defectors will first increase, and then vanish. During the phase of their invasion, they will exploit and eventually deplete the AllC players. This is a kind of pyrrhic victory: the defectors end up meeting mostly TFT-players, and this will be their undoing.

Looking at it from the point of view of defectors, any invasion attempt while $z>\hat{z}$ is doomed to failure and will result in a state with $y=0$ and $z>c/wb$. The only hope for the defectors is to wait with the invasion attempt until drift, i.e. a succession of random shocks, has sent the state, along the edge $y=0$, to the region where $z<\hat{z}$. This drift needs some time. If the invasion attempts occur too often, the drift will never have the time needed to lead into the zone which favours defectors. Thus the defectors should not try too frequently to invade. In other terms, cooperators will be safe only if invasion attempts by defectors are sufficiently frequent. If they are too rare, a cooperative society might lose its immunity—random fluctuations may lead to a state with too few discriminators to repel an invasion attempt by defectors. Let us mention in this context that we assume mutations to be so rare that they do not lead to a deterministic drift term (otherwise we would not be able to keep the treatment entirely analytic).

In order to deal with errors, it is convenient to use the results from Nowak and Sigmund (1990), where the payoffs for stochastic reactive strategies are computed. Each such strategy is given by a triplet (f,p,q), where f is the probability to cooperate in round 0 and p resp. q are the probabilities to cooperate after a cooperation resp. defection by the co-player in the previous round. In Nowak and Sigmund (1990) it is shown that if a player uses strategy (f,p,q) against a co-player using (f',p',q'), the payoff is given by

$$\frac{-c(e+wre')+b(e'+wr'e)}{(1-w)(1-uw^2)}, \tag{11}$$

where $r:=p-q$, $r':=p'-q'$, $u:=rr'$, $e:=(1-w)f+wq$ and $e':=(1-w)f'+wq'$.

AllC is given by $(1,1,1)$, AllD by $(0,0,0)$ and TFT is given by $(1,1,0)$. We will assume that an intended donation is mis-implemented with a probability ε, and an intended refusal with a probability $k\varepsilon$, for some $k\geqslant 0$. (It makes sense to distinguish between these two errors, and in particular to keep the case $k=0$ in mind.) Then the three strategies are given by $(1-\varepsilon,1-\varepsilon,1-\varepsilon)$, $(k\varepsilon,k\varepsilon,k\varepsilon)$ and $(1-\varepsilon,1-\varepsilon,k\varepsilon)$, respectively.

Applying this formula to the strategies AllC, AllD and TFT, we obtain a 3×3 payoff matrix M which, at first glance, looks somewhat daunting. But it can be simplified considerably. We will use the fact that the replicator dynamics on S_3 is unchanged if we subtract, in each column of M, the diagonal from all elements. Up to the multiplicative factor $c(1-(k+1)\varepsilon)/1-w$, the normalized matrix (which we still denote by M) is of the form

$$M=\begin{pmatrix} 0 & -1 & \delta\sigma \\ 1 & 0 & -\kappa\sigma \\ \delta & -\kappa & 0 \end{pmatrix}, \tag{12}$$

where we used

$$\delta:=w\varepsilon,\quad \kappa:=1-w+wk\varepsilon,\quad \sigma:=\frac{b\theta-c}{c-c\theta},$$
$$\theta=w(1-(k+1)\varepsilon). \tag{13}$$

We note that $\bar{P}=z(1+\sigma)P_z$. Using

$$P_z-\bar{P}=P_z[1-(1+\sigma)z], \tag{14}$$

we see that in the interior of S_3, $\dot{z}=0$ iff $g:=1-(1+\sigma)z$ vanishes. It is easy to see that $g=0$ defines an orbit connecting the fixed points $F_{yz}:=(0,1-\hat{z},\hat{z})$ and $F_{xz}:=(1-\hat{z},0,\hat{z})$, where $\hat{z}:=(1+\sigma)^{-1}$. On the edge $x=0$ there is a bistable competition between defectors and discriminators, their basins of attraction separated by F_{yz}. On the edge $y=0$ there is a stable coexistence between the discriminators and the undiscriminating altruists at the point F_{xz}. On the edge $z=0$ the defectors dominate the indiscriminating altruists.

In the interior of S_3 we obtain an invariant of motion

$$V:=x^Ay^Bz^C[1-(1+\sigma)z] \tag{15}$$

with $A=\kappa/\theta$, $B=\delta/\theta$ and $C=-1/\theta$ (note that $A+B+C+1=0$).

The interior fixed point is

$$F=(\kappa\sigma,\delta\sigma,1)\frac{1}{1+\sigma(\kappa+\delta)}. \tag{16}$$

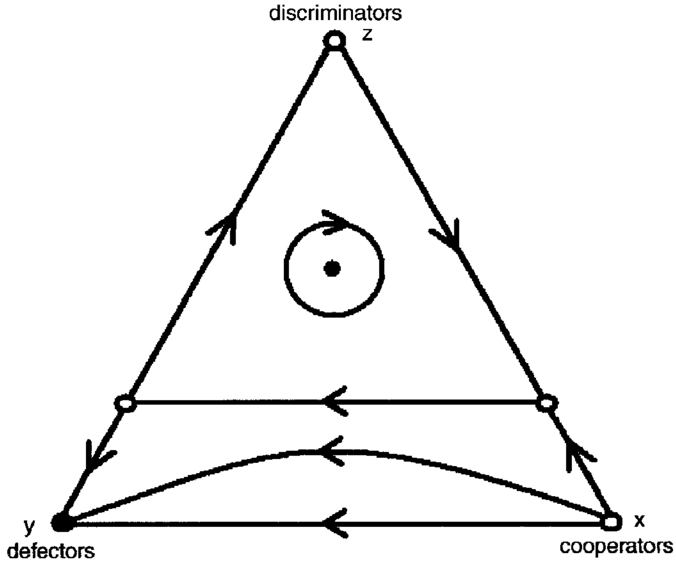

Fig. 2. The replicator dynamics of direct reciprocity with errors in implementation (i.e. $\varepsilon > 0$), for $w < 1$.

The dynamics is shown in Fig. 2. There is a horizontal orbit on the line with $z = \hat{z}$, connecting the fixed points F_{xz} and F_{yz} (the latter is a Nash equilibrium). Below this line, all orbits converge to $y = 1$, the defectors win. The part above the line is filled with periodic orbits surrounding the unique fixed point: they correspond to the constant level curves of the invariant of motion V given by (15). The time averages correspond to the values at the fixed point F. This fixed point is stable, but not asymptotically stable. We note that the amount of defectors (whose time average corresponds to $\delta\sigma$) can be made arbitrarily small if the error rate is sufficiently reduced. On the other hand, the basin of attraction of the defectors can be arbitrarily small if ε is sufficiently small and w sufficiently close to 1.

For $w = 1$ we obtain as payoff matrix, up to the multiplicative factor $c(1 - (k+1)\varepsilon)$,

$$M = \begin{pmatrix} 0 & -1 & \beta \\ 1 & 0 & -k\beta \\ \varepsilon & -k\varepsilon & 0 \end{pmatrix}, \tag{17}$$

where

$$\beta := \frac{1}{c}\left(\frac{b-c}{1+k} - \varepsilon b\right). \tag{18}$$

(Recall that, using Frobenius, we have to multiply all values with the factor $1 - w$ in order to obtain the average payoff *per round*.) If $k > 0$ (i.e. if there is a positive probability that an intended refusal results in a donation), the dynamics is the same as in Fig. 2, the z-coordinate of the separatrix is

$$\hat{z} := \frac{c}{b-c}\left(\frac{(k+1)\varepsilon}{1-(k+1)\varepsilon}\right). \tag{19}$$

If $\varepsilon \to 0$ the separatrix merges with $z = 0$ and we obtain a system whose payoff matrix is

$$M = \begin{pmatrix} 0 & -c & (b-c)/(1+k) \\ c & 0 & -k(b-c)/(1+k) \\ 0 & 0 & 0 \end{pmatrix}. \tag{20}$$

This is a rock-paper-scissors game: AllD is outcompeted by TFT, which is outcompeted by AllC, which is outcompeted by AllD in turn. The unique fixed point in the interior of S_3 is $F = (k(b-c)/(k+1)b, (b-c)/(k+1)b, c/b)$. We conclude that for $k > 0$ (positive probability that an intended refusal turns into a donation), the replicator dynamics is as shown in Fig. 3.

If, on the other hand, we first consider the limiting case $\varepsilon = 0$ (with $w < 1$), we obtain the dynamics shown in Fig. 1. If we then consider the limit case $w = 1$, we obtain Fig. 4. We note that the passages to the limit $w = 1$ and $\varepsilon = 0$ do not commute.

Traditionally, it is assumed in most treatments of indirect reciprocity that only intended donations are mis-implemented, not intended defections (Panchanathan and Boyd, 2003; Fishman, 2003; Brandt and Sigmund, 2004). This is quite in line with everyday experience. We note that in Fishman (2003), the failure of an intended donation is not due to an error, but to a lack of resources. Such a lack of resources can occur occasionally, by pure chance, and has the same effects as an error: it results in an unintended defection. In indirect reciprocity, it turns out that if we assume that intended defections also fail, the resulting dynamics is not appreciably different. Interestingly, however, it makes a difference in direct reciprocity, for the passage to the limit $w = 1$.

To see this, let us assume that $k = 0$. In the limiting case $w = 1$, the payoff matrix is given, up to the factor

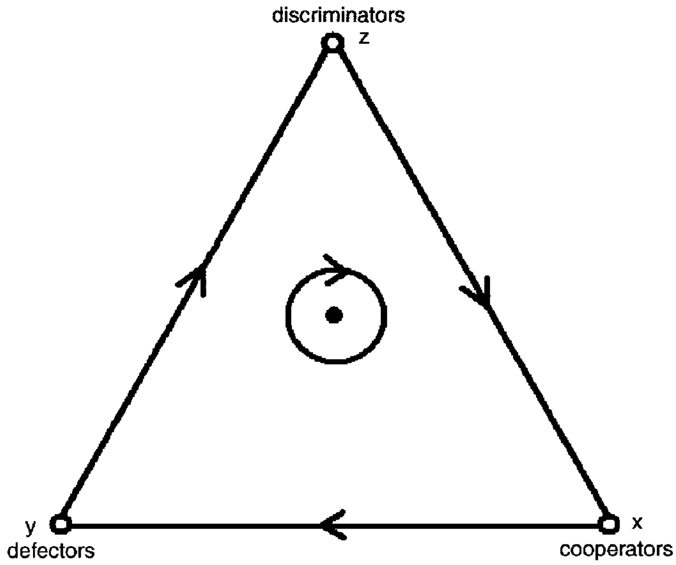

Fig. 3. The replicator dynamics of direct reciprocity, with errors in implementation, for $w = 1$ (the infinitely iterated game).

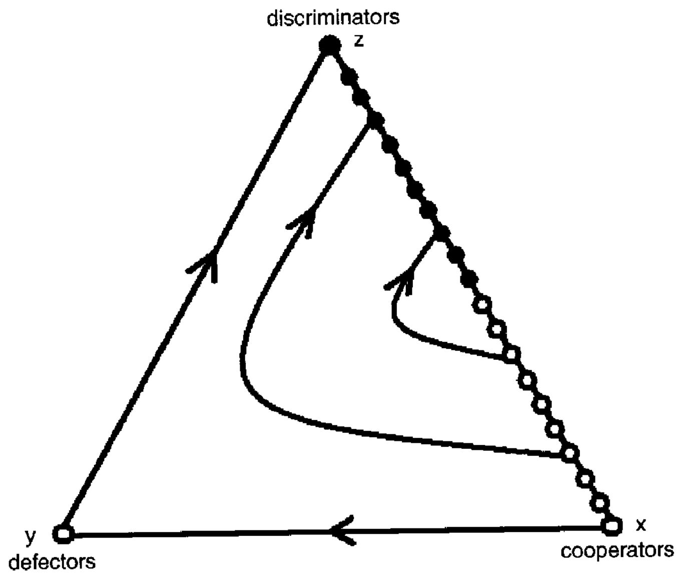

Fig. 4. The replicator dynamics of direct reciprocity, in the absence of errors, i.e. $\varepsilon = 0$, for $w = 1$.

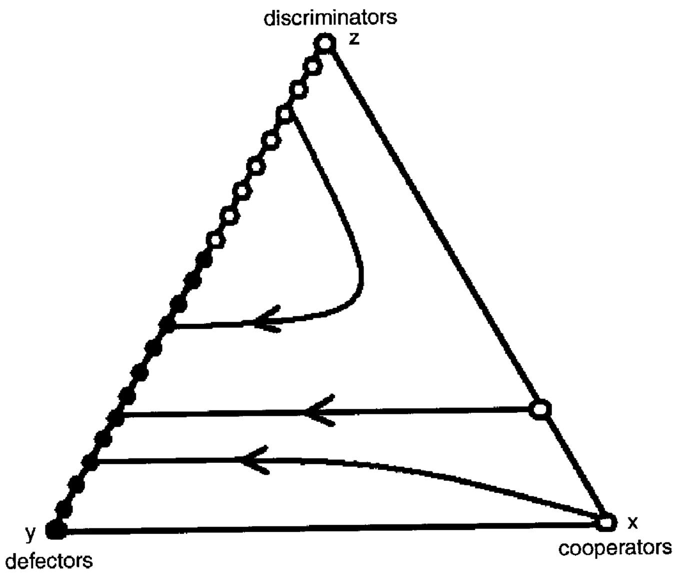

Fig. 5. The replicator dynamics of direct reciprocity, if only donations are mis-implemented, for $w < 1$.

$c(1-\varepsilon)$, by

$$M = \begin{pmatrix} 0 & -1 & \beta \\ 1 & 0 & 0 \\ \varepsilon & 0 & 0 \end{pmatrix}. \tag{21}$$

This yields a completely different picture. The edge $x = 0$ consists of fixed points. Intuitively, this is clear: errors between two TFT players will eventually lead to mutual defection, and this can never be redressed by another error. Thus their average payoff per round will be 0. The fixed points with $z \leqslant \hat{z}$ are Nash equilibria, where

$$\bar{z} = c/b(1-\varepsilon). \tag{22}$$

The dynamics looks as in Fig. 5, which is an intriguing mirror-image of Fig. 1. Finally, if we let $\varepsilon \to 0$, we obtain Fig. 6 as a mirror image of Fig. 4.

A very interesting related paper has recently been submitted (Imhof et al., 2005). It also studies, in the context of direct reciprocity, the interplay of AllC, AllD and TFT. Instead of assuming errors, it imposes a cost of complexity to the TFT strategy. The payoff matrix, therefore, is

$$M = \begin{pmatrix} b-c & -c & b-c \\ b & 0 & b(1-w) \\ b-c-v & -c(1-w)-v & b-c-v \end{pmatrix}, \tag{23}$$

where $v > 0$ is a small number corresponding to an extra cost for using a conditional strategy, rather than an unconditional one. The edge $y = 0$, now, consists of an orbit leading from $z = 1$ to $x = 1$: TFT is dominated by AllC. The dynamics on the other edges is as before. There exists a unique fixed point F in the interior of S_3:

$$F := \left(1 - \frac{v}{cw} - \frac{c}{bw}, \frac{v}{cw}, \frac{c}{bw}\right). \tag{24}$$

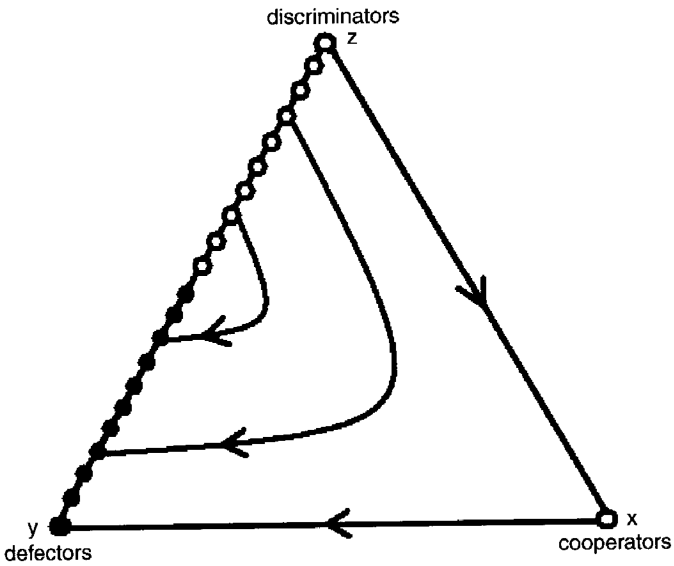

Fig. 6. The replicator dynamics of direct reciprocity if only donations are mis-implemented, for $w = 1$.

A simple computation shows that the eigenvalues of the Jacobian of the replicator equation, at the point F, are complex conjugate and have positive real part. Hence F is unstable, and in the vicinity the orbits spiral outward, clockwise. Since F is a Nash equilibrium, and $y = 1$ is also a Nash equilibrium, it follows by the odd number theorem (see e.g. Hofbauer and Sigmund, 1998) that there must exist a third Nash equilibrium, which necessarily must be F_{yz}. This point is saturated, and hence a saddle. There must be an orbit with F as α-limit and F_{yz} as ω-limit. All other orbits in the interior of S_3 converge to $y = 1$, so that the defectors win (see Fig. 7). This follows easily from Zeeman (1980). As shown numerically in Imhof et al. (2005), the

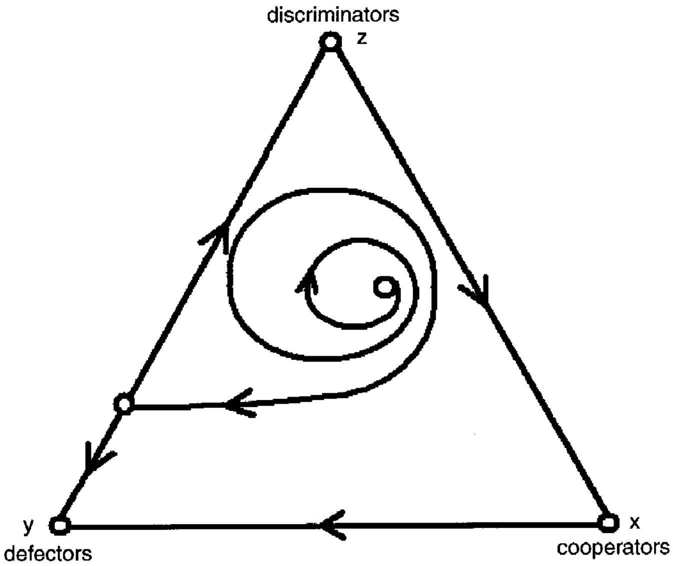

Fig. 7. The replicator dynamics of direct reciprocity if there is a cost to complexity. On the edge $y = 0$, the discriminators are dominated by the unconditional altruists.

addition of a mutation term introduces a limit cycle. The gist of this paper discusses the case of a finite population and shows that the corresponding stochastic process spends most of its time in the vicinity of the TFT corner, provided the population is sufficiently large, the number of rounds sufficiently high and the mutation rate sufficiently small. It would be of considerable interest to find out whether a corresponding result holds if players are not penalized by a cost of complexity but are liable to make errors. We stress that the bifurcation due to the cost of complexity v is quite different from the bifurcation due to the error probability ε, although in both cases the limit equilibrium is $(1 - c/bw, 0, c/bw)$.

4. Indirect reciprocity

Two of the main differences between direct and indirect reciprocation are the following.

(1) The TFT strategy discriminates according to what happened in the previous round. There are two distinct ways of translating this in the context of indirect reciprocity (see also Boyd and Richerson, 1989). Players can base their decision on what happened to them in the previous round; alternatively, they can base their decision on what their co-player did in the previous round. Roughly speaking, players can either be affected by a diffuse feeling of indebtedness ('Somebody helped me, I feel elated and therefore will help the next person'), or else, they can be moved by a feeling of appreciation ('My co-player did a noble thing, not to me but to a third party, and I will now help my co-player in turn'). In both cases, some general feeling of gratitude seems at work.

In one case, A gives to B and therefore B gives to C. In the other case, A gives to B and therefore C gives to A (see Fig. 8). In one case, the discriminator received a benefit,

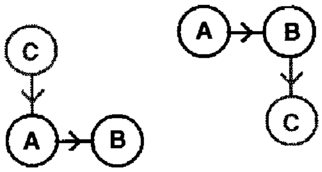

Fig. 8. Two approaches to indirect reciprocity. If A gives to B, C may either decide to reward A or expect help from B.

and thanks a person who did not help him. In the other case, the discriminator rewards a benefactor—but for an action that did not benefit him.

Interestingly, both factors seem to show up in economic experiments (cf. our remarks at the end of this section). But in the theoretical models considered so far, rewarding works fairly well and thanking not at all.

(2) The two players engaged in direct reciprocation experience in parallel the same number of rounds. By contrast, the histories of two players interacting via indirect reciprocity intersect only once, and thus each has a different numbering of his rounds: a donor in the first round may be matched with a recipient who has reached her fifth round, etc.

One more remark on the patterns of interaction between the players. In a more sophisticated direct reciprocity model, we could assume that the players alternate, either regularly or randomly, as donor and recipient, rather than acting simultaneously (Nowak and Sigmund, 1994; Frean, 1994). Similarly, in a *less* sophisticated model of indirect reciprocity, we could assume that all players start at the same time and that their rounds are synchronized (Nowak and Sigmund, 1998b; Panchanathan and Boyd, 2003; Fishman, 2003; Ohtsuki and Iwasa, 2004). This does not agree, however, with the continuous replicator dynamics, which is based on the assumption that generations blend into each other, or that learning occurs continuously. In fact, a synchronous model would better fit with a difference equation.

A high value for w, i.e. a large number of rounds, is less plausible with indirect than with direct reciprocity, since in a realistically small population, players experiencing many rounds would necessarily have to interact numerous times with the same partner, and hence be engaged in direct reciprocity. Nevertheless, the limiting case of $w \to 1$ has been considered by some authors (e.g. Ohtsuki and Iwasa, 2004). We shall see that in our model, setting $w = 1$ does not change much.

We will consider a continuous entry model, as in Brandt and Sigmund (2005). Players enter a large population one by one, interact asynchronously with different players at random times, and exit. Since we assume that the population is large, its composition will change only slowly, so that it is stationary during an individual's lifetime.

We consider the case that C gives A, i.e. that discriminators are motivated to reward players, and give if their co-player gave in the previous round. Again, we denote by ε the probability of not implementing an intended donation.

Let q be the probability that a player knows (either through direct observation or via gossip) what a randomly chosen co-player did in the previous round. Furthermore, let us posit that discriminators are trustful in the sense that if they have no information, they assume that their recipient gave help in the previous round. With h we denote the frequency of players with a good reputation (i.e. having given in their previous round). It is easy to see that $h = \bar{\varepsilon}(x + z(1 - q + qh))$, so that we obtain

$$h = \frac{\bar{\varepsilon}(x + (1-q)z)}{1 - \bar{\varepsilon}qz}. \tag{25}$$

The payoff in round n (with $n \geq 1$) for an indiscriminate altruist is

$$P_x(n) = -c\bar{\varepsilon} + b\bar{\varepsilon}[x + z[(1-q) + \bar{\varepsilon}q]]. \tag{26}$$

Indeed, such a player always tries to donate, at a cost $-c$ (this succeeds with probability $\bar{\varepsilon}$). On the other hand, a player is the object of an intended donation if the co-player who donates is either an unconditional cooperator (probability x) or a discriminator (probability y) who either does not know the player's reputation (probability $1 - q$) or else knows the reputation (probability q), and that reputation is good (probability $\bar{\varepsilon}$, because it can only be bad if the player, an unconditional altruist, made a mistake in the previous round). The benefit resulting from an intended donation is $b\bar{\varepsilon}$, because the donation can fail with probability ε. Similarly, the payoff for a defector is

$$P_y(n) = b\bar{\varepsilon}[x + (1-q)z], \tag{27}$$

and for a discriminator which we call A, it is

$$P_z(n) = -c\bar{\varepsilon}(1 - q + qh) + b\bar{\varepsilon}[x + z[(1-q) + \bar{\varepsilon}q(1 - q + qh)]]. \tag{28}$$

The second term in the sum is (up to the expected benefit $b\bar{\varepsilon}$) just the probability that the co-player intends to make a donation to player A. This happens either if the player is an unconditional altruist (probability x), or if he is a discriminator (probability z) who either does not know the reputation of A (probability $1 - q$) or else knows the reputation (probability q), and this reputation is good. The reputation of A is good if in the previous round, A intended to donate (either because A did not know the co-player's reputation or else because that reputation was good, an event whose probability is h), and if, moreover, A succeeded in the intended donation (probability $\bar{\varepsilon}$).

A straightforward computation shows that

$$P_z(n) - P_y(n) = [P_x(n) - P_y(n)](1 - q + qh). \tag{29}$$

The same relation holds for the first round, although the payoffs for the first round are slightly different: $P_x(0) = -c\bar{\varepsilon} + b\bar{\varepsilon}(x + z)$, $P_y(0) = b\bar{\varepsilon}(x + z)$ and $P_z(0) = -c\bar{\varepsilon}[(1-q) + qh] + b\bar{\varepsilon}(x + z)$. Hence the total payoff values P_x, P_y and P_z also satisfy

$$P_z - P_y = [P_x - P_y](1 - q + qh). \tag{30}$$

Clearly $P_x(n) - P_y(n) = \bar{\varepsilon}(-c + b\bar{\varepsilon}qz)$ (for $n \geq 1$) and $P_x(0) - P_y(0) = -c\bar{\varepsilon}$. Thus if $w = 1$ the payoff values per round satisfy

$$P_x - P_y = \bar{\varepsilon}(-c + b\bar{\varepsilon}qz) \tag{31}$$

and for $w < 1$,

$$P_x - P_y = \bar{\varepsilon}(-c + wb\bar{\varepsilon}qz). \tag{32}$$

If we normalize by setting $P_y = 0$ then, up to the factor $\bar{\varepsilon}$, we obtain

$$P_x = f, \quad P_z = f(1 - q + qh), \tag{33}$$

where $f = -c + wb\bar{\varepsilon}qz$.

Let us first consider the corresponding replicator equation without the common factor f. Since $h = \bar{\varepsilon}(x + (1-q)z)/(1 - \bar{\varepsilon}qz)$, this equation has the same orbits as the equation with

$$P_x = 1 - \bar{\varepsilon}qz, \quad P_z = 1 - q + \bar{\varepsilon}qx. \tag{34}$$

If $q < 1$ and $\varepsilon > 0$, we have $0 = P_y < P_z < P_x$ and hence all orbits in S_3 converge to $x = 1$, with the exception of the edge $x = 0$. An invariant of motion is given by $V = zx^{q-1}y^{-\varepsilon q}$.

If $\varepsilon = 0$ (no errors), the edge $y = 0$ consists of fixed points and the invariant of motion is $V = zxq^{-1}$. If $q = 1$ (full information about the co-players) the edge $x = 0$ consists of fixed points and the invariant of motion is $V = zy^{-\varepsilon}$.

Let us now consider the replicator dynamics for (33).

If $q < c/wb\bar{\varepsilon}$ then f is negative for all values of z between 0 and 1, and hence on all of S_3. Multiplication with f corresponds thus to a time-reversal. This means that the indiscriminating altruists are dominated by both the discriminators and the defectors, while the discriminators are dominated by the defectors. All orbits in the interior of the simplex lead from $x = 1$ (indiscriminating altruists only) to $y = 1$ (defectors only). This means that if the probability q to know the co-players past is too small (i.e. if there is not much scope for reputation), cooperation cannot evolve, a well-known result from Nowak and Sigmund (1998a) (see Fig. 9).

If $q > c/wb\bar{\varepsilon}$, then the line $z = c/wbq\bar{\varepsilon}$ intersects the interior of the simplex S_3 and defines a segment of fixed points. Indeed, on that line, $0 = P_y = P_x = P_z$. These fixed points are all Nash equilibria. In the simplex S_3, all orbits lie on the same curves as with (34), but the orientation has not changed in the region with $z > c/wbq\bar{\varepsilon}$ (see Fig. 10).

This means in particular that the mixture of discriminating and undiscriminating altruists given by $z = c/wbq\bar{\varepsilon}$ and $y = 0$ corresponds to a fixed point of the replicator dynamics. A cooperative population of two types of altruists can exist, if the average level of information within the population is sufficiently high. We note that this equilibrium is stable. However, it is not asymptotically stable, since it belongs to a segment of fixed points.

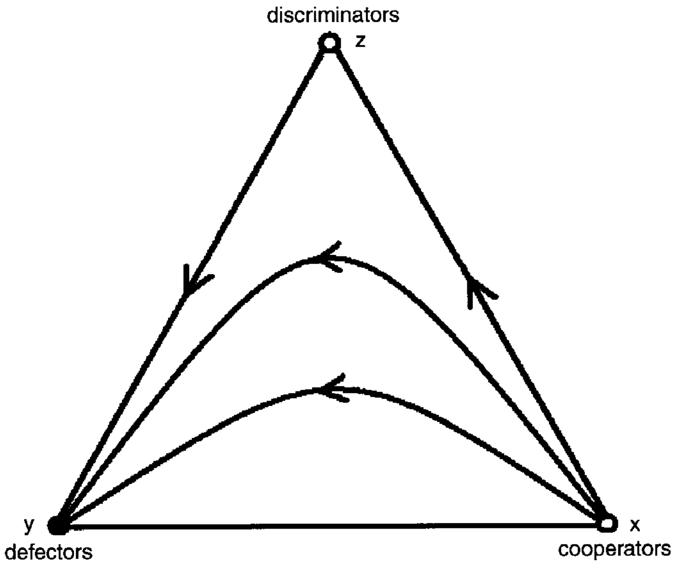

Fig. 9. The replicator dynamics of indirect reciprocity for $\varepsilon > 0$ if q, the information about the co-player's last move, is small. The same dynamics holds if discriminators are expected to donate whenever they received a donation in the previous round.

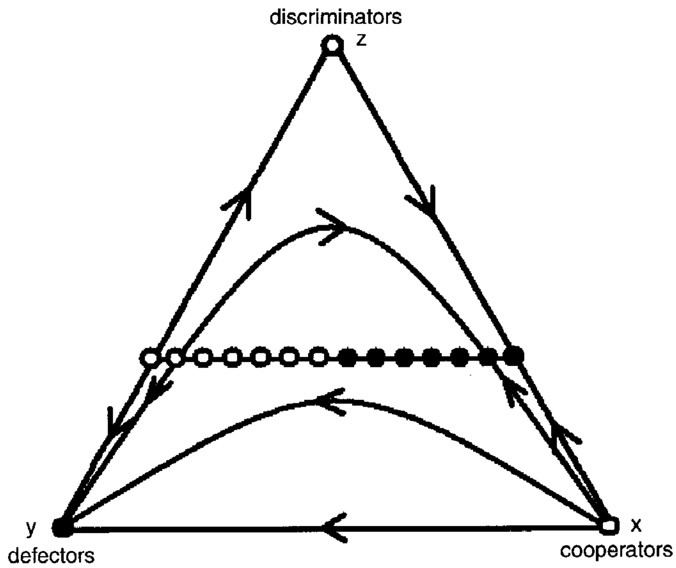

Fig. 10. The replicator dynamics of indirect reciprocity for $\varepsilon > 0$ if q, the information about the co-player's last move, is sufficiently large. The dynamics looks the same if one assumes synchronous rounds, and $w < 1$.

The dynamic behaviour in the vicinity of the segment of Nash equilibria is interesting. One part of the segment is transversally stable, in the sense that small perturbations away from the segment are counteracted by the dynamics. In the other part, small perturbations are amplified by the dynamics. A small deviation to higher z-values will lead, first to an increase and then to a decrease of discriminators, and thus eventually back to the stable part of the segment. By contrast, in the unstable part of the fixed points segment, a small deviation to lower z-values leads to the fixation of defectors.

In the limiting case $\varepsilon = 0$ (no errors), the edge $y = 0$ consists of fixed points, of which those with $z \geqslant c/wbq$ are Nash equilibria. The line with $z = c/wbq$ consists of fixed points, too. Below this line, all orbits converge to $y = 1$. Above the line, each orbit converges to a Nash equilibrium on $y = 0$. (see Fig. 11).

In the limiting case $q = 1$ (full information) the edge $x = 0$ consists of fixed points, of which those with $z < c/wb\bar{\varepsilon}$ are Nash equilibria. The line with $z = c/wb\bar{\varepsilon}$ consists of fixed points which are all stable. Hence the dynamics is as shown in Fig. 12.

If both $q = 1$ and $\varepsilon = 0$ the edges $x = 0$ and $y = 0$ both consist of fixed points. In the interior of S_3, all orbits

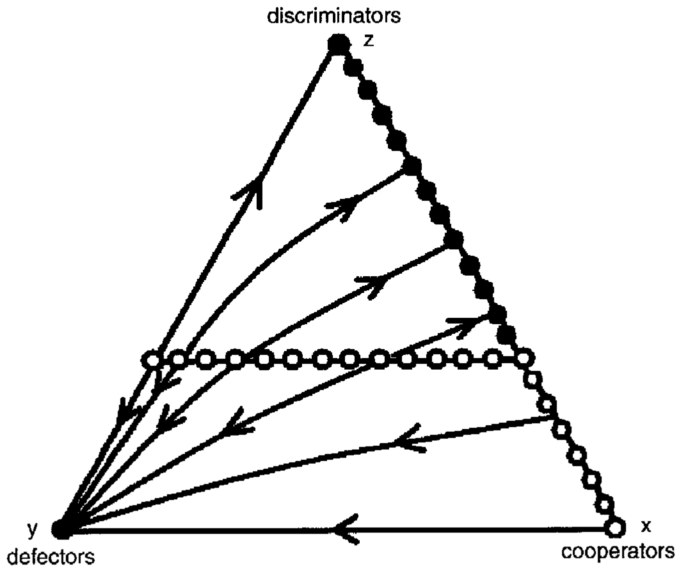

Fig. 11. The replicator dynamics of indirect reciprocity for $\varepsilon = 0$ (no errors) if $q < 1$ but sufficiently large. The dynamics looks the same if one assumes synchronous rounds, $\varepsilon = 0$, and $w < 1$.

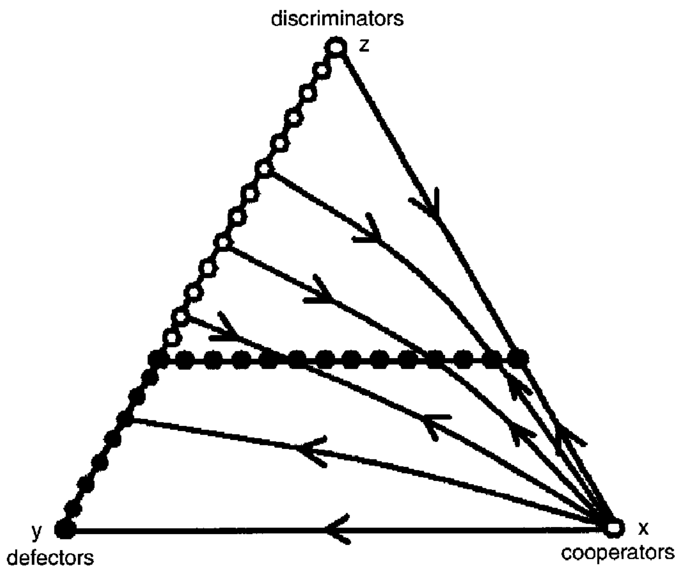

Fig. 12. The replicator dynamics of indirect reciprocity if $\varepsilon > 0$ and $q = 1$.

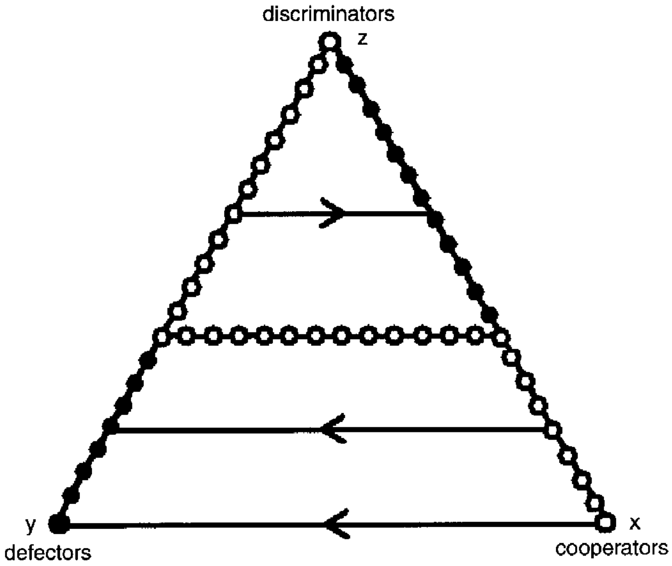

Fig. 13. The replicator dynamics of indirect reciprocity if $\varepsilon = 0$ and $q = 1$.

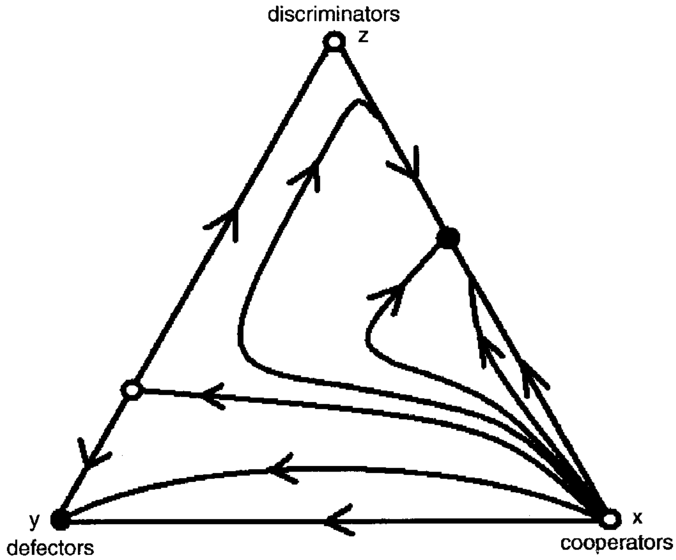

Fig. 14. The replicator dynamics of indirect reciprocity if one assumes synchronous rounds, $\varepsilon > 0$, and a number of rounds which is fixed, or Poisson distributed.

remain at parallel to the $z = 0$ edge. Those with $z > c/wb$ point from left to right (the defectors vanish), while those with $z < c/wb$ point from right to left (the undiscriminating altruists vanish). Again, the line $z = c/wb$ consists of Nash equilibria. The dynamics is shown in Fig. 13.

As an aside, let us turn to previous models of indirect reciprocity which were based on the assumption that all players experience their rounds in a synchronized way. In the case of no errors and a fixed number of rounds, this leads to a dynamics as in Fig. 1. Without errors and with a constant probability $w < 1$ for another round we obtain a dynamics as in Fig. 11 (Nowak and Sigmund, 1998b). With errors and a constant probability $w < 1$ for another round, the dynamics looks as in Fig. 10 (Panchanathan and Boyd, 2003). With errors and a number of rounds which is fixed in advance or Poisson distributed, the dynamics is bistable and displays an attractor consisting of a stable mixture of discriminating and undiscriminating altruists, as shown in Fig. 14 (Fishman, 2003; Brandt and Sigmund, 2004). Let us note that synchronized games can easily be set up in experiments, but they seem unlikely to occur under natural circumstances.

Finally, let us briefly consider the case when players are motivated by a general feeling of indebtedness, and discriminators decide to give whenever they have received support in the previous round. If we denote by h the probability that a player has received support in the previous round, we see that $h = \bar{\varepsilon}(x + hz)$.

In round n the payoff values for unconditional altruists, defectors and discriminators are $P_x(n) = -c\bar{\varepsilon} + hb$, $P_y(n) = hb$ and $P_z(n) = -ch\bar{\varepsilon} + hb$. If we assume that a discriminator, in the first round, always donates, we get $P_z(0) = -c\bar{\varepsilon} + bh$. After normalizing the total payoff values such that $P_y = 0$, we obtain, up to the factor $(1 - w)^{-1}$,

$$P_x = -c\bar{\varepsilon}, \quad P_z = P_x[1 - w(1 - h)]. \tag{35}$$

The dynamics looks as in Fig. 9: the defector's corner with $y = 1$ is a global attractor. This still holds if the error rates are modified, or if one assumes that the discriminators defect in the first round, etc. In particular, letting $\varepsilon \to 0$ or $w \to 1$ changes nothing. It is all the more surprising that some experiments (and, indeed, everyday introspection) show that indirect reciprocation based on a generalized feeling of indebtedness is not rare (Engelmann and Fischbacher, 2002; Dufenberg et al., 2001). To the best of our knowledge, a theoretical explanation for this is still lacking (cf. Boyd and Richerson, 1989).

5. Discussion

Direct and indirect reciprocation are obviously closely related. For instance, Nowak and Sigmund (1998b) pointed out that the discriminating strategy used in their treatment of indirect reciprocity is nothing but the 'Observer TFT' discussed by Pollock and Dugatkin (1992) in the context of the iterated Prisoner's Dilemma. There is a line of papers by economists, most of it antedating the work by evolutionary biologists, which discuss indirect reciprocation in populations of rational players (see e.g. Rosenthal, 1979; Kandori, 1992; Ellison, 1994). In particular, the so-called folk theorem on repeated games states that every feasible pair of payoff values for the two players is obtainable by strategies in equilibrium (i.e. such that no player has an incentive to deviate). It must only be assumed that (a) the payoff is larger than the security level that players can guarantee themselves (in our case, this is 0), and (b) that the probability for another round is sufficiently large (see e.g. Fudenberg and Maskin, 1986). The equilibrium can be achieved by so-called 'trigger strategies' which switch to defection as soon as the other player defects. Intuitively, it makes no sense to exploit the co-player in one round if one, thereby, forfeits all chances for mutual cooperation in further rounds. The argument works for direct as well as for indirect reciprocation. The difference between personal enforcement, in the former case, and community enforcement, in the latter, is irrelevant to the sequence of payoffs encountered by an individual player. The folk theorem assumes rational players having full information, but it can be considerably extended.

In both direct and indirect reciprocity, the presence of unconditional altruists weakens the stability of the cooperation. Although the dynamics in Figs. 2 and 10 look very different, in each case the defector's corner $y = 1$ has a basin of attraction which can, and will, be ultimately reached if the population is subject to arbitrarily small random shocks for a sufficiently long time. As long as this is not the case, the three strategies coexist in a stable, but not in an asymptotically stable way. In one case, periodic oscillations around the Nash equilibrium can increase, and in the other case, the state can wander along a continuum of Nash equilibria, until the defector's basin is threateningly close.

There are devices leading out of this fundamental unstability, of course. One variation working for both direct and indirect reciprocity is provided by the concept of 'good standing' (Sugden, 1986). A player who cooperates is in good standing. A player failing to donate to a recipient in good standing will acquire a 'bad standing'. But a player refusing to donate to a player in bad standing will keep a good standing. Essentially, this means that discriminators have to distinguish between justified and unjustified defections. The corresponding strategy in the context of direct reciprocity is called contrite TFT, in the context of indirect reciprocity 'standing strategy'. In both cases, the strategy (or family of strategies, to be precise) is more stable than the discriminating strategy we have considered in this paper (see e.g. Boerlijst et al., 1997; Leimar and Hammerstein, 2001; Panchanathan and Boyd, 2003). On the other hand, it requires higher cognitive capabilities, and suffers from errors in perception (rather than implementation). Experimental evidence for the standing strategy seems to be disappointing (Milinski et al., 2001).

There are other ways of boosting the stability of reciprocation. For instance, the so-called Pavlov strategy (which prescribes to donate if both players, in the previous round, made the same decision) can lead to stable cooperation in direct reciprocation whenever $b > 2c$ (Fudenberg and Maskin, 1990; Nowak and Sigmund, 1993). In indirect reciprocation, the assumption that the social information of each player grows during his or her lifetime can also lead to an asymptotically stable mixture of discriminating and indiscriminating altruists (Brandt and Sigmund, 2005).

To return to Maynard Smith, let us note that the last chapter of 'The Major Transitions in Evolution' (Maynard Smith and Szathmary, 1997) discusses the evolution of human cooperation between non-relatives by placing particular emphasis on the social contract and the public goods game. It can be argued that indirect reciprocity occupies a place in between direct reciprocity and public goods games (see also Milinski et al., 2002a, b; Panchanathan and Boyd, 2004). It describes interactions between two players only (donor and recipient), but it involves reputation, and hence communication, acting within a larger group. Moreover, like the social contract, indirect reciprocity requires at least a rudimentary form of a theory of mind. Indeed, empathy is clearly required for a bystander to form a moral judgement about an action taking place between two other players. Reputation also plays an essential role in the theory proposed by Sigmund et al. (2001) to explain the role of punishment in public goods games. Thus reciprocation, both direct and indirect, constitutes certainly an important element for understanding the last of the major transitions on the list of Maynard Smith, the one that led to human cooperation.

Acknowledgements

We wish to thank Lorens Imhof and Martin A. Nowak for helpful discussions.

References

Alexander, R.D., 1987. The Biology of Moral Systems. Aldine de Gruyter, New York.

Axelrod, R., Hamilton, W.D., 1981. The evolution of cooperation. Science 211, 1390–1396.

Boerlijst, M.C., Nowak, M.A., Sigmund, K., 1997. The logic of contrition. J. Theor. Biol. 185, 281–294.

Boyd, R., Richerson, P.J., 1989. The evolution of indirect reciprocity. Soc. Networks 11, 213–236.

Brandt, H., Sigmund, K., 2004. The logic of reprobation: action and assessment rules in indirect reciprocity. J. Theor. Biol. 231, 475–486.

Brandt, H., Sigmund, K., 2005. Indirect reciprocity, image scoring, and moral hazard. Proc. Natl Acad. Sci. USA 102, 2666–2670.

Dufenberg, M., Gneezy, U., Gueth, W., van Damme, E., 2001. Direct vs indirect reciprocation—an experiment. Homo Oeconomicus 18, 19–30.

Dugatkin, L.A., 1997. Cooperation among Animals: an Evolutionary Perspective. Oxford University Press, Oxford.

Ellison, G., 1994. Cooperation in the Prisoner's Dilemma with anonymous random matching. Rev. Econ. Stud. 61, 567–588.

Engelmann, D., Fischbacher, U., 2002. Indirect reciprocity and strategic reputation-building in an experimental helping game, Working paper, University of Zürich, Zürich.

Fehr, E., Fischbacher, U., 2003. The nature of human altruism. Nature 425, 785–791.

Fishman, M.A., 2003. Indirect reciprocity among imperfect individuals. J. Theor. Biol. 225, 285–292.

Frean, M.R., 1994. The Prisoner's Dilemma without synchrony. Proc. R. Soc. London (B) 257, 75–79.

Fudenberg, D., Maskin, E., 1986. The folk theorem in repeated games with discounting or with incomplete information. Econometrica 50, 533–554.

Fudenberg, D., Maskin, E., 1990. Evolution and cooperation in noisy repeated games. New Dev. Econ. Theory 80, 274–279.

Gintis, H., Bowles, S., Boyd, R., Fehr, E., 2003. Explaining altruistic behavior in humans. Evolution Hum. Behav. 24, 153–172.

Hammerstein, P. (Ed.), 2003. Genetic and Cultural Evolution of Cooperation. MIT Press, Cambridge.
Hofbauer, J., Sigmund, K., 1998. Evolutionary Games and Population Dynamics. Cambridge University Press, Cambridge.
Imhof, L., Fudenberg, D., Nowak, M.A., 2005. Evolutionary cycles of cooperation and defection. PNAS USA 102, 10797–10800.
Kandori, M., 1992. Social norms and community enforcement. Rev. Econ. Stud. 59, 63–80.
Leimar, O., Hammerstein, P., 2001. Evolution of cooperation through indirect reciprocation. Proc. R. Soc. London B 268, 745–753.
Lotem, A., Fishman, M.A., Stone, L., 1999. Evolution of cooperation between individuals. Nature 400, 226–227.
Maynard Smith, J., 1982. Evolution and the Theory of Games. Cambridge University Press, Cambridge.
Maynard Smith, J., 1984. Game theory and the evolution of behaviour. Behavioural Brain Sci. 7, 95–125.
Maynard Smith, J., 1988. Games Sex and Evolution. Harvester-Wheatsheaf, London.
Maynard Smith, J., Szathmary, E., 1997. The Major Transitions in Evolution. Oxford University Press, Oxford.
Milinski, M., Semmann, D., Bakker, T.C.M., Krambeck, H.J., 2001. Cooperation through indirect reciprocity: image scoring or standing strategy? Proc. R. Soc. London B 268, 2495–2501.
Milinski, M., Semmann, D., Krambeck, H.J., 2002a. Donors in charity gain in both indirect reciprocity and political reputation. Proc. R. Soc. London B 269, 881–883.
Milinski, M., Semmann, D., Krambeck, H.J., 2002b. Reputation helps solve the 'Tragedy of the commons'. Nature 415, 424–426.
Nowak, M.A., Sigmund, K., 1990. The evolution of stochastic strategies in iterated games. Acta Applicandae Math. 20, 247–265.
Nowak, M.A., Sigmund, K., 1993. Win-stay, lose-shift outperforms tit-for-tat. Nature 364, 56–58.
Nowak, M.A., Sigmund, K., 1994. The alternating Prisoner's Dilemma. J. Theor. Biol. 168, 219–226.
Nowak, M.A., Sigmund, K., 1998a. Evolution of indirect reciprocity by image scoring. Nature 282, 462–466.
Nowak, M.A., Sigmund, K., 1998b. The dynamics of indirect reciprocity. J. Theor. Biol. 194, 561–574.
Ohtsuki, H., Iwasa, Y., 2004. How should we define goodness? Reputation dynamics in indirect reciprocity. J. Theor. Biol. 231, 107–120.
Panchanathan, K., Boyd, R., 2003. A tale of two defectors: the importance of standing for evolution of indirect reciprocity. J. Theor. Biol. 224, 115–126.
Panchanathan, K., Boyd, R., 2004. Indirect reciprocity can stabilise cooperation without the second-order free-rider problem. Nature 432, 499–502.
Pollock, G.B., Dugatkin, L.A., 1992. Reciprocity and the evolution of reputation. J. Theor. Biol. 159, 25–37.
Rosenthal, R.W., 1979. Sequences of games with varying opponents. Econometrica 47, 1353–1366.
Seabright, P., 2004. The Company of Strangers. Princeton University Press, Princeton, NJ.
Selten, R., Hammerstein, P., 1984. Gaps in Harley's argument on evolutionarily stable learning rules and in the logic of tit for tat. Behavioural Brain Sci. 7, 115–116.
Sigmund, K., Hauert, C., Nowak, M.A., 2001. Reward and punishment. Proc. Natl Acad. Sci. USA 98, 10757–10763.
Sugden, R., 1986. The Economics of Rights, Cooperation and Welfare. Basil Blackwell, Oxford.
Trivers, R., 1971. The evolution of reciprocal altruism. Q. Rev. Biol. 46, 35–57.
Wedekind, C., Braithwaite, V.A., 2002. The long-term benefits of human generosity in indirect reciprocity. Curr. Biol. 12, 1012–1015.
Wedekind, C., Milinski, M., 2000. Cooperation through image scoring in humans. Science 288, 850–852.
Zeeman, E.C., 1980. Population dynamics from game theory. In: Global Theory of Dynamical Systems. Lecture Notes in Mathematics, vol. 819. Springer, New York.