Reciprocity and communication of partner quality

O. LEIMAR

Department of Zoology, Stockholm University, S-106 91 Stockholm, Sweden (`olof.leimar@zool.su.se`)

SUMMARY

In a cooperative exchange, the size of a partner's contribution is likely to depend both on the partner's ability to supply help and on the partner's need for help in return. Referring to such needs and abilities as aspects of partner quality, it follows that variation in the amount of help offered in a relationship could transmit information about partner quality. A plausible behaviour might then be to vary the investment in a partner according to available information about partner quality and to invest little in a partner who offers little in return. Thus, regulation of a relationship through communication of partner quality would tend to follow the principle of reciprocity. In an analysis of an iterated game where players have private information about their needs and abilities, I verify this possibility by describing an evolutionarily stable state space strategy, referred to as 'state-dependent reciprocity', entailing communication of partner quality. Although the evolution of cooperation has been studied in great detail, there has been no previous analysis of communication of needs and abilities in a relationship. It may well be that such communication is of major importance for the evolution of cooperative behaviour in nature.

1. INTRODUCTION

Trivers (1971) suggested that mutual aid giving could prevail in evolution through a behavioural mechanism called reciprocal altruism. Reciprocity operates when an initially present inclination to help a partner is reduced by the partner's failure to offer help in return. Most previous evolutionary analyses of reciprocity, usually formalized as a repeated Prisoner's Dilemma (PD) game (e.g. Axelrod & Hamilton 1981; Boyd & Lorberbaum 1987; Boyd 1989; Nowak & Sigmund 1992, 1993, 1994; Bendor & Swistak 1995; Leimar 1997), have ignored a possibly very important cause of behavioural variability in cooperative interactions. As noted by Boyd (1992), in all likelihood there will be variation in the need for help and in the ability to provide help, both from one individual to the next and over time for a given individual, although the characteristics of a potential partner may be poorly known initially. Here, I present the simplest possible cooperative evolutionarily stable strategy (ESS) for an iterated game where players have private information about their needs and abilities. With this ESS, the actions performed in the game can communicate private information and an offer of help may be viewed as motivated by a combination of a player's own needs and abilities and the available information about the partner's characteristics.

In the terminology of Wasser (1982), an individual B with a good ability to supply cooperative services of interest to another individual A has high associate quality, meaning that it would be beneficial for A to have B as a partner. Similarly, it would be beneficial for A to have a partner with a strong need for the services A can supply, and we can regard a high benefit of receiving help and a low cost of giving help as two aspects of partner quality. Both these aspects are likely to influence cooperative behaviour. For instance, from an analysis of a repeated game where players differ in partner quality but know each other's pay-offs, Boyd (1992) concluded that there can be ESSs where the player with high partner quality provides help more frequently than the player with low partner quality. My intention here is to study a situation where players know their own needs and abilities, but can gain information about a partner's characteristics only through the behaviour performed.

Given variation in partner quality, we should expect an individual to regulate cooperative contributions according to its needs and abilities and, as a consequence of the partner acting in the same way, to also regulate its contributions according to available information about the partner's characteristics. Since the partner's actions provide information about partner quality, such regulation of a relationship tends to follow the principle of reciprocity. Receiving little or no help is an indication of low partner quality, making further investments in such a partner less profitable. Straightforward as this conclusion may seem, communication of partner quality has played little role in the very extensive literature on evolutionary models of cooperative interactions, perhaps because of the potential complexity of an analysis. Nevertheless, the most basic aspects of the phenomenon can be dealt with quite simply. In the following, I first discuss a game without variation in partner quality and I then show how the analysis can be extended to include variation in needs and abilities.

2. THE ALTERNATING PRISONER'S DILEMMA

For a game theoretical analysis, one needs to prescribe some format for the exchange of help. A pair of players interacting over a number of rounds could, for instance, either move simultaneously or in an alternating fashion. I will assume alternating moves, because this allows a very simple implementation of reciprocity. However, the choice of format is not really crucial.

Consider now the alternating PD game, without any variation in partner quality (Nowak & Sigmund 1994; Leimar 1997). In this game, players move in alternating rounds as specified by two roles: the leader moves in rounds 1, 3, etc. and the follower in rounds 2, 4, etc. At the start, a coin toss determines who should be leader. When moving, a player has the options to either give help or not and these actions are denoted C and D ('cooperate' and 'defect'). With a small probability, a player makes a mistake when attempting to execute an action; for instance, defects when intending to cooperate. The benefit of receiving help is b and the cost of giving help is c. Not helping has no cost and confers no benefit. After each round, the interaction continues with probability ω and ends with probability $1-\omega$, where $0<\omega<1$.

A strategy for the game should prescribe a choice of action (or, for a randomized strategy, a distribution over the available actions) for each possible situation where a player is about to move. In principle, it can take into account the player's role and the sequence of past moves. For the analysis of repeated games, one commonly limits consideration to a much smaller set of strategies, typically those where the choice of action only depends on the immediate past. However, this particular restriction is somewhat too strong, since it excludes very natural strategies. A better approach is to consider strategies that can be represented using a simple state space (Leimar 1997). These may be viewed as behavioural or psychological mechanisms: in a given motivational state some particular action is performed, as specified by an action rule, and the actions can cause the state to change over time. An account of the actions performed in the most recent rounds can be interpreted as a state, but there are additional possibilities. For instance, a state could represent whether or not the interaction is proceeding in a satisfactory manner and if not, who is responsible for a disturbance. The strategy 'contrite tit-for-tat' (Sugden 1986; Boyd 1989) for the repeated PD with simultaneous moves is one such example.

For the alternating PD, there is a particularly simple and natural implementation of reciprocal altruism. This strategy, which I have called 'reciprocity' (Leimar 1997), has a state space with two states, 1 and 2, which we may label 'satisfied' and 'frustrated'. The action rule for reciprocity is to cooperate when satisfied and defect when frustrated. To understand how the states are determined it is helpful to look at a situation where both players use reciprocity. Initially, the leader is satisfied. In the following, a player becomes frustrated when the partner was satisfied but defected (by mistake) in the previous round; otherwise the player is satisfied with the course of play.

The strategy reciprocity serves to regulate behavioural noise in the interaction. For instance, consider the following sequence of states and actions:

$$\begin{array}{lllllll} \text{leader :} & 1\mathrm{C} & & 2\mathrm{D} & & 1\mathrm{C} & \\ \text{follower :} & & 1\underline{\mathrm{D}} & & 1\mathrm{C} & & 1\mathrm{C}. \end{array}$$

The underlined D represents a mistaken defection by the follower in the second round, which causes the leader to become frustrated and to defect in the third round. However, the follower remains satisfied in the fourth round, because the leader's defection was justified, and the exchange of help is resumed. It is worth noting that the states are not determined solely by the action in the previous round, nor are they fully determined by the actions in some fixed number of previous rounds. The state frustrated acts as a memory of the partner's unjustified failure to help and this memory will remain until the partner starts offering help.

When two players of reciprocity meet, a player's state is determined by the partner's state and action in the previous round. There is actually an even more intimate link between the states of the players. Note that a player is satisfied precisely when defection by the player would cause the partner to become frustrated, so the states can be viewed as summarising available information about the partner's future behaviour. I have called a state space strategy reflexive when it entails such a close fit between the players' behavioural mechanisms. Reflexivity turns out to be crucially important for evolutionary stability (Leimar 1997).

For games with a time structure, genetic variation in strategies and mistaken actions create different kinds of behavioural variability, in a way that matters for evolutionary stability. For instance, the concept of reflexivity has no particular relevance if some unspecified genetic variation lies behind most behavioural variability. On the other hand, reflexivity is essential if deviating actions are mostly noise in a particular behavioural mechanism, since one should then expect individuals to be adapted to each other's behavioural mechanisms.

When dealing with behavioural noise, one feasible modelling approach would be to explicitly assume certain mistake probabilities for the various situations a player may be in and then to apply Maynard Smith's (1982) criteria for an ESS. However, the most important effects of such noise will be present in the limit of small mistake probabilities and an attractive alternative is then to consider this limit. A systematic and powerful treatment of evolutionary stability when players are slightly fallible has been developed by Selten (1983), leading to the concept of a limit ESS. For state space strategies, the most important effect of behavioural noise is that all motivational states occur on a regular basis during interactions, so that actions and state transitions ought to be optimal for each state. This makes dynamic

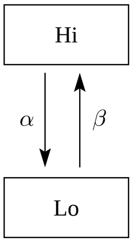

Figure 1. The stochastic dynamics of a player's private state z. There are two alternatives, $z = \text{Hi}$ and $z = \text{Lo}$. Private state changes occur randomly, and α and β are round-to-round state transition probabilities, with $\alpha > 0$, $\beta > 0$ and $\alpha + \beta \leqslant 1$.

programming a valuable tool for finding limit ESSs of repeated games. For the alternating PD, there are, in fact, many cooperative limit ESSs, although reciprocity is the simplest among them. Using dynamic programming one can show that reciprocity is a limit ESS when $\omega > c/b$ (Leimar 1997).

3. VARIATION IN PARTNER QUALITY

(*a*) *Private states*

While keeping the format of the alternating PD game, suppose now that players may vary in the costs and benefits of cooperation. They each have private knowledge of their own characteristics, represented as a private state variable and denoted by z. The benefit of receiving help is $b(z)$ and depends on the recipient's private state z. The cost of giving help is $c(z)$ and depends on the donor's private state z. As a simplification, assume that there are only two levels of partner quality, $z = \text{Lo}$ and $z = \text{Hi}$, where Hi could be interpreted as a high need for help and/or a high ability to provide help, and conversely for Lo. Having only two private states gives a rather crude picture of variation in costs and benefits. For instance, an individual in high need of help might not have a high ability to provide help; if anything, one might expect the opposite to be the case. Nevertheless, partner quality is a function of the cost of giving help in relation to the benefit of receiving help, and the two private states are intended to represent that kind of variation. It may be helpful to think of variation only in the benefits or only in the costs.

Various events cause a player's private state to change. Although receiving help might lessen the need for additional help, and giving help might increase the cost of giving more, as a first approximation, private state changes are taken to be unrelated to the actions performed during the game (figure 1). The changes may be fairly frequent, with several transitions expected during an interaction, or they can be quite rare. In either case, a player has no prior information about a partner's private state at the start of the game.

With the private state dynamics shown in figure 1, the probability of $z = \text{Hi}$ at statistical equilibrium is

$$p_{\text{eq}} = \frac{\beta}{\alpha + \beta}.$$

Table 1. *Public state space dynamics for state-dependent reciprocity*

(There are two public states, $y = 1$ and $y = 2$, having the interpretation satisfied and frustrated, and a player has a public state only in every second round, when the player moves. The leader's public state in the very first round is $y = 1$.)

player's state	partner's next state[a] C	D	player's next state[b] CC	CD	DC	DD
1	1	2	1	2	1	1
2	1	1	1	2	1	2

[a] Single-round state transition: the partner's next public state, after either a C or a D by the player.

[b] Two-round state transition: the player's next public state after the action sequence CC, CD, DC or DD. The two-round transition results from the composition of two single-round transitions.

Table 2. *Action rule for state-dependent reciprocity*

(A player's total state, $x = (y, z)$, is a combination of a public state y and a private state z, and the player offers help only when in state $x = (1, \text{Hi})$.)

public state	private state	action
1	Hi	C
1	Lo	D
2	Hi	D
2	Lo	D

Given that $z = \text{Hi}$ in one round, the probability of $z = \text{Hi}$ in the next round is $1 - \alpha$, and given that $z = \text{Lo}$ in one round, the probability of $z = \text{Hi}$ in the next round is β. For the case $\alpha + \beta = 1$, these are both equal to p_{eq}, so information about the current z does not help to predict a future z, and the private states in different rounds are independent random variables. For $\alpha + \beta < 1$, there is instead a gradual decay of information, e.g. given that $z = \text{Hi}$ in one round, the probability of $z = \text{Hi}$ in subsequent rounds gradually approaches p_{eq}, starting from $1 - \alpha$, and the approach is slower for smaller $\alpha + \beta$ (note also that a case with $\alpha + \beta > 1$ would give an oscillatory approach, which is of no interest to us here).

(*b*) *State-dependent reciprocity*

For certain kinds of variation in costs and benefits, an individual in the Lo private state ought not to offer any help; for instance, this is obvious if $c(\text{Lo})$ exceeds $b(\text{Lo})$ and the rate β of transition to the Hi private state is very small. Thus, the act of defection by a partner could communicate to a player that the partner's private state is Lo. There may, however, be other circumstances causing the partner to defect, such as the player's own previous defection, and an efficient strategy must use available information

Table 3. *Examples of interaction sequences for SDR, showing the public and private states and the alternating actions of the two players*

((*a*) The follower drops to Lo in round 4 and stops helping, thus communicating the private state change to the leader. This frustrates the leader, who stops helping although still in private state Hi. In round 8, the follower switches back to Hi and resumes the giving of help, which is reciprocated by the satisfied leader. (*b*) The leader initially has $z = \text{Lo}$ and the exchange of help does not start until round 5, when the leader communicates a switch to Hi by an offer of help, thus satisfying the follower.)

	1	2	3	4	5	6	7	8	9	10	11	12
(*a*)	1HiC	Hi	1HiC	Hi	2HiD	Hi	2HiD	Hi	1HiC	Hi	1HiC	···
	Hi	1HiC	Hi	1LoD	Lo	1LoD	Lo	1HiC	Hi	1HiC	Hi	···
(*b*)	1LoD	Lo	1LoD	Hi	1HiC	Hi	1HiC	Hi	1HiC	···		
	Hi	2HiD	Hi	2HiD	Hi	1HiC	Hi	1HiC	Hi	···		

to distinguish the various circumstances in the best possible way. In general, a player can obtain information about a partner's private state only through the sequence of actions performed so far. Since these actions are 'public', i.e. known to both players, it is natural to attempt to summarize the information in a public state. It turns out that the states 'satisfied' and 'frustrated' can serve as such public states. A player's choice of action could then depend on both the public and the private state.

The strategy 'state-dependent reciprocity' (SDR), where state-dependent is meant to indicate a dependence on the private state, consists of a public state space with state transitions (dynamics), as shown in table 1 together with the action rule in table 2. The public states and their transitions are identical to the ones for reciprocity. I previously described these in a context where both players use reciprocity, and this description corresponds to the single-round transitions in table 1. A more basic description, making no reference to the partner's strategy or states, is provided by the two-round transitions in table 1. SDR differs from reciprocity only in the action rule: whereas reciprocity prescribes cooperation when satisfied, SDR prescribes cooperation only when both satisfied and in private state Hi.

One can interpret SDR as a psychological mechanism. The public states, satisfied and frustrated, encode a player's experience of previous actions in the game, filtering out all but one crucial circumstance: whether the partner gave help when such help was expected. The public state y and the private state z can be regarded as two interacting components of a player's motivation to help the partner (table 2). In order to illustrate the operation of the mechanism, two examples of interaction sequences where both players use SDR are shown in table 3.

Let us now examine more closely how a player's public state encodes information about the partner's private state, given that both use SDR. To simplify the argument, assume mistake probabilities to be small enough to ignore the possibility of mistaken actions. Also, assume $\alpha + \beta < 1$, implying a gradual decay of information about the private state.

With the action rule in table 2, a partner's action when satisfied communicates the private state, but there is no such communication when the partner is frustrated. Thus, considering the one-round transitions in table 1, when a player is frustrated the partner must have been satisfied but failed to offer help in the previous round, so the partner's private state must have been Lo in that round. From figure 1 we then see that the probability of the partner having $z = \text{Hi}$ in a round where the player is frustrated is equal to β, which is less than p_{eq}. On the other hand, when a player is satisfied, the partner's private state must have been Hi the last time it was communicated, which means that the probability of $z = \text{Hi}$ for the partner must be greater than or equal to p_{eq}. Equality holds only when the partner's private state has not yet been communicated (e.g. for the leader player in rounds 1, 3 and 5 in table 3*b*); otherwise, a satisfied player knows that the partner was Hi either one round ago (e.g. the follower player in round 2 in table 3*a*), or several rounds ago (e.g. five rounds ago for the follower player in round 8 in table 3*a*). In conclusion, satisfied and frustrated correspond to high and low values of the probability that the partner's private state is Hi.

(*c*) *Evolutionary stability of SDR*

As an aid to understanding, I have described SDR as a psychological mechanism, but for the question of evolutionary stability a more formal treatment is desirable. The original alternating PD is a game of perfect information, entailing that all circumstances potentially influencing a player's decisions are know to both players, which greatly simplifies the determination of best replies to a strategy. The presence of private states adds the complication of having to deal with conditional probability distributions over a partner's private states. As is well known, when an optimization problem has a state structure but the states are not completely observable, dynamic programming can be performed using the conditional probability distribution over the state space as a 'probabilistic state' (Whittle 1983). In view of the reflexivity of the public dynamics in table 1 (Leimar 1997), this means that one can find best replies to SDR using dynamic programming on a player's total state space, extended to also include the probability of the partner currently having z = Hi, given that the partner uses SDR.

The possibility of mistaken actions must also be taken into account, and the natural question to pose is whether SDR is a limit ESS. A point of concern with regard to SDR might be whether a mistaken C from a frustrated partner provides any information about the partner's private state (cf. table 2). I have used the simple assumption that the probability of such a mistake is independent of the partner's private state (but assuming the mistake to be more likely for a partner with z = Hi would actually not change the stability conditions for SDR).

With arguments similar to those in Leimar (1997), one can show that SDR is a limit ESS when the dynamic programming equation for SDR has strict optimal actions equal to those specified by SDR. In fact, for small, positive mistake probabilities, strict optimal actions imply that SDR is a strict best reply to itself and thus an ESS by way of Maynard Smith's first condition. As pointed out above, for a frustrated player, the partner has z = Hi with probability β (in the limit of small mistake probabilities), whereas for a satisfied player, there is a range of different values, from p_{eq} up to $1 - \alpha$, and there should be strict optimal actions agreeing with SDR for each of these values. The resulting stability conditions are illustrated in figure 2. For C to be a strict optimal action at $(y, z) = (1, \mathrm{Hi})$, the benefit $b(\mathrm{Hi})$ must be large enough to be covered by the inclined surface. For D to be a strict optimal action at $(1, \mathrm{Lo})$, the cost $c(\mathrm{Lo})$ must be above the inclined surface. In addition, D is a strict optimal action at $(2, z)$ as soon as $c(z) > 0$. In general, the requirements on the parameters for SDR to be a limit ESS are fairly lenient. A high cost of giving help for individuals of low partner quality (i.e. in private state Lo) is perhaps the most important stabilizing factor (figure 2).

As seen from the first two examples in figure 2, the stability conditions are rather similar for low and moderately high rates of private state transitions. In the third example, which has $\alpha + \beta = 1$, information about partner quality in one round does not help to predict partner quality in the next round. One could then view the occurrence of z = Lo as a kind of behavioural noise; for instance, an accidental decrease in the ability to supply help, corresponding to a large $c(\mathrm{Lo})$. This illustrates that SDR contains elements of both noise regulation and assessment of partner quality, with the former dominating for $\alpha + \beta$ close to one and the latter dominating for small $\alpha + \beta$.

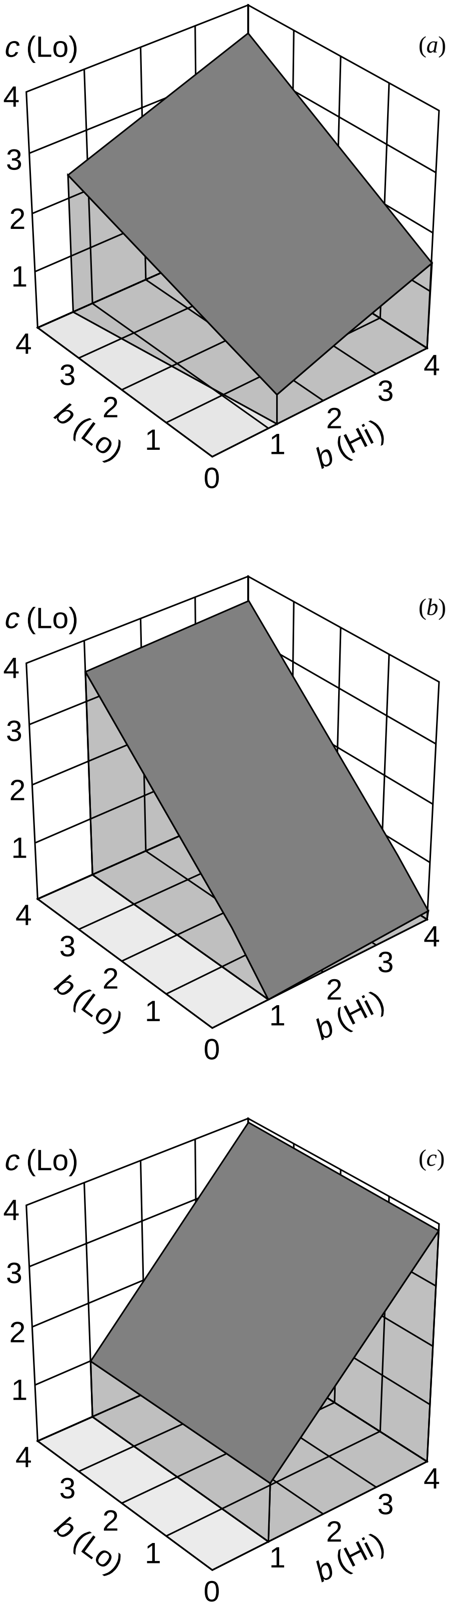

Figure 2. SDR is a limit ESS for benefit–cost combinations above the inclined surfaces. As a pay-off normalization, the cost $c(\mathrm{Hi})$ is put to one. The three examples have different transition probabilities α and β, but in all three the interaction lasts on average 100 rounds ($\omega = 0.99$): (*a*) a player is usually in z = Hi and state transitions are fairly frequent ($\alpha = 0.1$ and $\beta = 0.4$); (*b*) a player is usually in z = Hi and state transitions are relatively rare ($\alpha = 0.01$ and $\beta = 0.04$); (*c*) a player is almost always in z = Hi and the private states in different rounds are statistically independent ($\alpha = 0.01$ and $\beta = 0.99$).

4. DISCUSSION

Most previous modelling of reciprocity has either neglected all reasons for behavioural variability beyond genetic variation (Axelrod & Hamilton 1981; Boyd & Lorberbaum 1987; Bendor & Swistak 1995), or has also dealt with random errors in the execution of a strategy (Boyd 1989; Nowak & Sigmund 1992, 1993, 1994; Leimar 1997). The main biological interest of my analysis is that it includes variation in partner quality and communication of partner quality as factors behind behavioural variability in cooperative interactions. Put slightly differently, we can ask whether we should interpret reciprocity, insofar as it has been observed, as a safeguard against non-cooperative mutants, as a means to regulate a relationship in the face of noise or as an adaptation to deal with partner quality variation. These are all logically consistent possibilities and their relative importance is an empirical issue.

My game theoretical analysis provides some suggestions for an assessment of the importance of variation in partner quality. Since the particular assumptions I used in constructing the game might not correspond exactly to any real situation, one should focus on those properties of SDR that have a chance of being robust. Two general aspects of the strategy are quite basic to its operation. First, the action rule (table 2) entails that an individual's tendency to invest in a partner depends both on its own needs and abilities and on the information about a partner's needs and abilities. These factors interact in determining the investment and both must be favourable for substantial investment to occur. Second, the action rule together with the state space dynamics (table 1) imply that a partner's behaviour is an important source of information about partner quality.

Consequently, for an interacting pair where one has low and the other has high partner quality, the exchange of help ought to stop or be reduced to a low level after an initial higher investment by the individual with high partner quality. If both have low partner quality, little will be transacted, whereas if both have high partner quality, a more intense and lasting exchange is to be expected. Changes in partner quality during an ongoing interaction should influence the investments in similar ways. Since these types of reactions are along the lines of the traditional idea of reciprocity, one really needs to know the nature of partner quality in a given case, and perhaps to manipulate it, in order to assess its importance.

Some information about partner quality, such as partner size (Külling & Milinski 1992), could be transmitted separately from any exchange of help, as was assumed in the analysis by Boyd (1992). This type of information would influence the level of investment in a partner and could also be of great importance for the forming of relationships (Noë & Hammerstein 1994), but would not in itself lead to reciprocity. Other aspects of partner quality will become evident only during the course of an interaction. Variation in 'boldness', possibly as a result of variation in hunger (Godin & Crossman 1994), could be an example from the carefully studied phenomenon of joint predator inspection in fish (Milinski 1987; Dugatkin 1988; Milinski *et al.* 1990; Dugatkin & Alfieri 1991; Külling & Milinski 1992). For egg trading in simultaneous hermaphrodites (Fischer 1980; Sella 1985), a partner's egg load is an aspect of partner quality that sometimes cannot be assessed directly (Sella 1988). Failure to reciprocate an egg laying bout might be a reliable indication of a lack of eggs in the partner.

In the interspecific mutualism between lycaenid butterfly larvae and ants (Pierce 1987), where the ants give protection and the larvae provide a nutritious secretion, both a larva's perception of the risk of enemy attack and the nutritional status of an ant colony have a strong influence on the interaction (Leimar & Axén 1993; Axén *et al.* 1996). For instance, a larva will sharply increase its rate of secretion following a simulated enemy attack, and the change in larval behaviour then leads to higher ant attendance (Leimar & Axén 1993; Axén *et al.* 1996).

Thus, empirical observations suggest that cooperative behaviour is commonly influenced by variation in needs and abilities, so that communication of partner quality could be an important driving force behind reciprocity. Nevertheless, purely random variation in the outcome of behaviour will be ever present and cannot be neglected when considering the function of reciprocity, particularly for long-lasting relationships where partners know each other well.

I thank three anonymous reviewers for their very helpful comments. This work was supported by the Swedish Natural Science Research Council.

REFERENCES

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

Axén, A. H., Leimar, O. & Hoffmann, V. 1996 Signalling in a mutualistic interaction. *Anim. Behav.* **52**, 321–333.

Bendor, J. & Swistak, P. 1995 Types of evolutionary stability and the problem of cooperation. *Proc. Natn. Acad. Sci. USA* **92**, 3596–3600.

Boyd, R. 1989 Mistakes allow evolutionary stability in the repeated Prisoner's Dilemma game. *J. Theor. Biol.* **136**, 47–56.

Boyd, R. 1992 The evolution of reciprocity when conditions vary. In *Coalitions and alliances in humans and other animals* (ed. A. H. Harcourt & F. B. M. de Waal), pp. 473-489. Oxford University Press.

Boyd, R. & Lorberbaum, J. 1987 No pure strategy is evolutionarily stable in the repeated Prisoner's Dilemma game. *Nature* **327**, 58–59.

Dugatkin, L. A. 1988 Do guppies play TIT FOR TAT during predator inspection visits? *Behav. Ecol. Sociobiol.* **23**, 395–399.

Dugatkin, L. A. & Alfieri, M. 1991 Tit-For-Tat in guppies (*Poecilia reticulata*): the relative nature of cooperation and defection during predator inspection. *Evol. Ecol.* **5**, 300–309.

Fischer, E. A. 1980 The relationship between mating system and simultaneous hermaphroditism in the coral reef fish, *Hypoplectrus nigricans* (Serranidae). *Anim. Behav.* **28**, 620–633.

Godin, J.-G. J. & Crossman, S. L. 1994 Hunger-dependent predator inspection and foraging behaviours in the threespine stickleback (*Gasterosteus aculeatus*) under predation risk. *Behav. Ecol. Sociobiol.* **34**, 359–366.

Külling, D. & Milinski, M. 1992 Size-dependent predation risk and partner quality in predator inspection of sticklebacks. *Anim. Behav.* **44**, 949–955.

Leimar, O. 1997 Repeated games: a state space approach. *J. Theor. Biol.* **184**, 471–498.

Leimar, O. & Axén, A. H. 1993 Strategic behaviour in an interspecific mutualism: interactions between lycaenid larvae and ants. *Anim. Behav.* **46**, 1177–1182.

Maynard Smith, J. 1982 *Evolution and the theory of games*. Cambridge University Press.

Milinski, M. 1987 TIT FOR TAT in sticklebacks and the evolution of cooperation. *Nature* **325**, 433–437.

Milinski, M., Külling, D. & Kettler, R. 1990 Tit for Tat: sticklebacks (*Gasterosteus aculeatus*) 'trusting' a cooperating partner. *Behav. Ecol.* **1**, 7–11.

Noë, R. & Hammerstein, P. 1994 Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism and mating. *Behav. Ecol. Sociobiol.* **35**, 1–11.

Nowak, M. A. & Sigmund, K. 1992 Tit for tat in heterogeneous populations. *Nature* **355**, 250–253.

Nowak, M. A. & Sigmund, K. 1993 A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **364**, 56–58.

Nowak, M. A. & Sigmund, K. 1994 The alternating Prisoner's Dilemma. *J. Theor. Biol.* **168**, 219–226.

Pierce, N. E. 1987 The evolution and biogeography of associations between lycaenid butterflies and ants. *Oxf. Surv. Evol. Biol.* **4**, 89–116.

Sella, G. 1985 Reciprocal egg trading and brood care in a hermaphroditic polychaete worm. *Anim. Behav.* **33**, 938–944.

Sella, G. 1988 Reciprocation, reproductive success, and safeguards against cheating in a hermaphroditic polychaete worm, *Ophryotrocha diadema* Åkesson, 1976. *Biol. Bull.* **175**, 212–217.

Selten, R. 1983 Evolutionary stability in extensive two-person games. *Math. Soc. Sci.* **5**, 269–363.

Sugden, R. 1986 *The economics of rights, co-operation and welfare*. Oxford: Blackwell.

Trivers, R. L. 1971 The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57.

Wasser, S. K. 1982 Reciprocity and the trade-off between associate quality and relatedness. *Am. Nat.* **119**, 720–731.

Whittle, P. 1983 *Optimization over time*, vol. II. New York: Wiley.

Received 27 March 1997; accepted 1 May 1997