

# Stability in negotiation games and the emergence of cooperation<sup>†</sup>

Peter D. Taylor\* and Troy Day

Department of Mathematics and Statistics, Queen's University, Kingston, Ontario K7L 3N6, Canada

Consider a two-player game in which each player contributes a costly resource to the common good of the pair. For such contests, the Nash equilibrium contribution,  $x^*$ , is one for which neither player can increase its pay-off by unilaterally altering its contribution from  $x^*$ . We study an elaboration of this game, which allows the players to exchange  $x$ -offers back and forth in a negotiation phase until they converge to a final pair of contributions,  $\hat{x}_1$  and  $\hat{x}_2$ . A significant feature of such negotiation games, hitherto unrecognized, is the existence of a set of neutrally stable equilibrium points in negotiation phase space. To explore the long-term evolutionary outcome of such games, we simulate populations containing various mixtures of negotiation strategies and, contrary to previous results, we often find convergence to a contribution that is more cooperative than the Nash equilibrium. Mathematical analysis suggests why this might be happening, and provides a novel and robust explanation for cooperation, that negotiation can facilitate the evolution of cooperative behaviour.

**Keywords:** cooperation; negotiation strategy; evolutionary stability

## 1. INTRODUCTION

The development of a theory to explain the evolution of cooperative behaviour between unrelated individuals has proved to be a demanding task. What is needed is a mechanism through which cooperative individuals can bestow the benefits of altruism on one another without being taken advantage of by selfish individuals. One early idea (Wilson 1975) is that this might occur when the effects of altruism itself tend to group altruistic individuals together more often than might occur by chance. Another idea (Axelrod & Hamilton 1981) hinges on reciprocity: individuals cooperate if there is some chance of meeting again in the future, to repay the good deeds and punish the bad. More recent developments have focused on the evolution of cooperation through indirect reciprocity by mechanisms such as reputation and image scoring (Nowak & Sigmund 1998; Lotem *et al.* 1999; Riolo *et al.* 2001; Milinski *et al.* 2002).

One feature common to all the above approaches is an assumption that an individual's action during an encounter is fixed and irreversible once chosen. Cooperative behaviour is very often observed during relatively prolonged and complex interactions, however, and therefore it is perhaps more reasonable to assume that some form of 'negotiation' between individuals takes place before the fitness-determining actions of each player are settled upon. During the negotiation, each individual observes its opponent and alters its own actions accordingly.

Focusing on the evolution of the negotiation strategy itself rather than the evolution of the action settled upon by an individual opens the door to the possibility of a natural mechanism for the evolution of cooperation. In this context, any player's final action will depend on how its negotiation strategy interacts with that of its opponent. Perhaps natural selection can drive the evolution of

negotiation strategies to a point where most individuals tend to negotiate cooperative outcomes with one another, but where such strategies nevertheless cannot be taken advantage of because they negotiate more guarded outcomes when interacting with individuals attempting to exploit them.

## 2. ANALYSIS

Analysing such games requires a fundamental change in the game-theoretic approach currently used in evolutionary biology (McNamara *et al.* 1999). In particular, an individual's negotiation strategy can be viewed as a reaction norm for its behaviour as a function of its opponent's behaviour (Agrawal 2001). One then focuses on characterizing the evolutionary stability of behavioural reaction norms rather than the behaviours themselves. The observed behaviours are then an outcome of the interaction of evolutionarily stable behavioural reaction norms. Here, we investigate a simple approach for modelling such negotiation games and use it to provide a novel explanation for the emergence of cooperative behaviour. Our approach also reveals a fundamental feature of such games: the existence of a curve of neutral stability.

Suppose that each player contributes a costly resource to the common good of the pair. If player  $i$  contributes  $x_i$ , the fitness of player  $i$  will be

$$F_i(x_1, x_2) = B_i(x_1 + x_2) - K_i(x_i), \quad (2.1)$$

where the benefits  $B_i(z)$  are assumed to depend on the total contribution  $z = x_1 + x_2$  and  $K_i(x_i)$  is the cost to player  $i$  of contributing  $x_i$ . We assume that benefit increases with total contribution with diminishing returns, and that costs increase with contribution in an accelerating manner. For the game without negotiation (Houston & Davies 1985), the Nash equilibria are determined by the two conditions

$$\partial F_i / \partial x_i = 0 \quad (i = 1, 2). \quad (2.2)$$

<sup>†</sup>Both authors contributed in equal part to this work.

\* Author for correspondence (taylorp@post.queensu.ca).




Figure 1. The negotiation process is illustrated for players with linear response rules (player 1,  $x_1$ , has  $\lambda_1 = 2/3$  and  $\rho_1 = 0.4$ , and player 2,  $x_2$ , has  $\lambda_2 = 1/2$  and  $\rho_2 = 0.3$  in equation (2.1)). Note that the  $x_2$ -response line uses the  $x_1$ -axis as the abscissa, and it projects any  $x_1$  offer onto the  $x_2$ -axis. Similarly, the  $x_1$ -response line uses the  $x_2$ -axis as the abscissa, and projects any  $x_2$  offer onto the  $x_1$ -axis. In the negotiation illustrated,  $x_1$  begins, and its successive offers are given by successive vertical lines moving from left to right. The responses by  $x_2$  are given by successive horizontal lines moving from top to bottom. The final contributions are at the intersection of the response lines, at  $\hat{x}_1 = 0.3$  and  $\hat{x}_2 = 0.15$ .

To incorporate negotiation, we employ response rules (McNamara *et al.* 1999). In our context, player  $i$ 's response rule  $r_i(x)$  specifies its resource 'offer', given that its opponent has offered resource level  $x$  (i.e.  $r_i(x)$  is its behavioural reaction norm). The negotiation phase involves an iteration of back-and-forth responses between the two players (each using their fixed response rule) and the equilibrium values of  $x$  are obtained for both players where the two response rules intersect (figure 1). Instead of seeking the Nash equilibrium level of resources offered (i.e.  $x$ ), we now seek the Nash equilibrium response rules,  $r_i(x)$  (i.e. the Nash equilibrium behavioural reaction norms). Using a local analysis at a fixed equilibrium, McNamara *et al.* (1999) argue that no player can do better than to use a linear response rule of the form

$$r_i(x) = \rho_i - \lambda_i x, \tag{2.3}$$

where the *responsiveness*  $\lambda_i$  measures the degree to which an individual responds to a change in  $x$  from its partner. In this version, a player's strategy is determined by the pair  $(\lambda, \rho)$ . If player  $i$  plays  $(\lambda_i, \rho_i)$ , its fitness is

$$W_i(\lambda_1, \rho_1, \lambda_2, \rho_2) = F_i(\hat{x}_1, \hat{x}_2), \tag{2.4}$$

where  $\hat{x}_1$  and  $\hat{x}_2$  are the solutions to the equations

$$\begin{aligned} x_1 &= \rho_1 - \lambda_1 x_2, \\ x_2 &= \rho_2 - \lambda_2 x_1, \end{aligned} \tag{2.5}$$

and this gives us

$$\hat{x}_1 = \frac{\rho_1 - \rho_2 \lambda_1}{1 - \lambda_1 \lambda_2} \quad \text{and} \quad \hat{x}_2 = \frac{\rho_2 - \rho_1 \lambda_2}{1 - \lambda_1 \lambda_2}, \tag{2.6}$$




Figure 2. The  $(\lambda, \rho)$  equilibrium curve for the functions  $B(z) = z/z + 1$  and  $K(x) = x^2$ . This curve has general equation  $F_1 - \lambda F_2 = 0$  which in this case becomes  $(1 - \lambda)(1 + \lambda)^3 = 2\rho(2\rho + 1 + \lambda)^2$ . The neutral line is drawn for the point  $\lambda^* = 0.5$ ,  $\rho^* = 0.223$ . Note that the Nash equilibrium for the game without negotiation (Houston & Davies 1985) is the value of  $x^*$  at which the  $F$ -level curve crosses the diagonal with slope zero, and therefore it corresponds to the  $\lambda = 0$  point of the equilibrium curve (where it crosses the  $\rho$ -axis). This value of  $x$  is signalled by a horizontal line in figure 3.

provided  $|\lambda_i| < 1$  for  $i = 1, 2$  (an assumption we now make).

In  $(\lambda, \rho)$  phase space, the Nash conditions for player  $i$  are

$$\begin{aligned} \frac{\partial W_i}{\partial \rho_i} &= \frac{1}{1 - \lambda_1 \lambda_2} \left[ \frac{\partial F_i}{\partial x_i} - \lambda_j \frac{\partial F_i}{\partial x_j} \right] = 0, \\ \frac{\partial W_i}{\partial \lambda_i} &= \frac{-\hat{x}_j}{1 - \lambda_1 \lambda_2} \left[ \frac{\partial F_i}{\partial x_i} - \lambda_j \frac{\partial F_i}{\partial x_j} \right] = 0, \end{aligned} \tag{2.7}$$

where  $j \neq i$  and the partial derivatives of the  $F_i$  are evaluated at  $(\hat{x}_1, \hat{x}_2)$ . We see that these two conditions both yield the same equation:

$$\frac{\partial F_i}{\partial x_i} - \lambda_j^* \frac{\partial F_i}{\partial x_j} \Big|_{\mathbf{x}=\hat{\mathbf{x}}^*} = 0 \quad (i = 1, 2, j \neq i), \tag{2.8}$$

and thus the equilibrium conditions for both players give us two equations in  $(\lambda_1, \rho_1, \lambda_2, \rho_2)$  and this specifies a two-dimensional equilibrium surface in four-space. To simplify matters here we will ignore any differences between the two players in features such as their inherent quality (see, for example, McNamara *et al.* 1999) and restrict attention to a symmetric version of the game. In this case we have only one function  $F(x_1, x_2) = B(x_1 + x_2) - K(x_1)$ , and we let  $W(\tilde{\lambda}, \tilde{\rho}, \lambda, \rho) = F(\hat{x}_1, \hat{x}_2)$  be the fitness of a  $(\tilde{\lambda}, \tilde{\rho})$  mutant in a  $(\lambda, \rho)$  population. The Nash conditions (equation (2.8)) give us one equation in  $(\lambda, \rho)$  and this specifies an equilibrium curve in the  $(\lambda, \rho)$  plane. An example of such a curve is shown in figure 2. One can show (Appendix A) that the curve is neutrally stable in the sense that in a resident population at a point on the curve, no mutant strategy has greater than resident fitness,




Figure 3. Change over time from simulation studies of the negotiation game under the joint forces of selection and mutation, starting with a uniform distribution in a rectangular region of  $(\lambda, \rho)$  strategy space. The population distribution of contributions  $x$  is plotted against log time (generations). (a) With the benefit and cost functions used for the graphs of this paper (figure 2 caption). (b) As in (a) except with linear benefit  $B(z) = 2z$ . The Nash equilibrium contribution for the game without negotiation (equation (2.2)) is shown as a horizontal line. We see convergence to a more cooperative contribution, the effect being more striking in (b). (c) The population of (a) recorded in  $(\lambda, \rho)$  space under the joint forces of selection and mutation. The region  $-1 \leq \lambda \leq 1$ ,  $0 \leq \rho \leq 0.6$  is partitioned as a  $30 \times 40$  grid. The population size is 45 000 individuals. Each generation, each individual has one random encounter and mutates with probability 0.005. The mutation range is  $\delta\rho = 0.01$ ,  $\delta\lambda = 0.025$ .

though there are mutant strategies that have resident fitness (see below).

McNamara *et al.* (1999) did not formulate the Nash conditions (equation (2.7)) in  $(\lambda, \rho)$  space and they therefore overlooked this key observation. As a result, their conclusion, that there exists a single evolutionarily stable strategy (ESS) at which the level of care is lower than that of non-negotiated outcomes, is incorrect. Rather, there is a continuum of outcomes that are all ESSs, spanning a range of levels of care that goes both lower and higher than the level found in the non-negotiated outcome. The above results illustrate that the existence of an equilibrium surface (or curve in the case that the game is symmetric) might be quite a general property of negotiation games.

It is useful to note that in a pure  $(\lambda, \rho)$  population the common contribution  $\hat{x} = \rho/(1 + \lambda)$  is the slope of the line (figure 2) drawn from  $(-1, 0)$  to the point  $(\lambda, \rho)$ . It follows from this interpretation of  $\hat{x}$  that points on the left side of the equilibrium curve are more cooperative than points on the right. Interestingly, although the above results demonstrate that all points on the equilibrium curve are weakly evolutionarily stable, the dynamics of evolutionary change actually tend to produce a relatively cooperative outcome, as we detail next.

### 3. THE EMERGENCE OF COOPERATION

We ran simulations of this game with individual variation in the response rule, and with an individual's reproductive success (and thus the representation of its response rule in the next generation) given by equation (2.1). Using two different benefit functions, our results reveal the evolution of negotiation rules that yield cooperative behaviour (figure 3).

A mathematical analysis of the negotiation game helps to expose the mechanism through which this occurs (Appendix A). In a resident population playing any one of the Nash equilibrium rules, there is a line of 'mutant' response rules that, when played against the resident or against one another, yield the same negotiated outcome as the residents, and therefore have resident fitness (figures 2 and 4). Thus, none of the Nash equilibrium response rules is evolutionarily stable in the strict sense (Hofbauer & Sigmund 1998). Crucially, however, the neutral mutant reaction norms will produce a different negotiated outcome than the resident reaction norm when played against something other than the resident. To see this, take figure 5 and move the blue resident line and observe the change in its intersection with the dashed red mutant line. This exemplifies a key feature of the model; it provides a




Figure 4. (a) Some level curves of  $W$ . At points on the diagonal  $x_1 = x_2 = x$  between  $x = 0$  and  $x = x_{max}$  these curves have slope between  $-1$  and  $1$ . (b) The  $(x_1, x_2)$  configuration at the equilibrium point  $\lambda^* = 0.5, \rho^* = 0.223$  depicted in figure 2. The response lines for players 1 (red) and 2 (blue) both intersect the diagonal at  $\hat{x}^* = 0.1486$  and the level curve for the fitness  $F(x_1, x_2)$  is drawn through that point. The response line for player 2 is tangent to that curve, indicating that player 1 cannot increase its fitness by moving its response line (red) to another location. The negotiated contributions are determined at the intersections of the red and blue lines.

natural mechanism through which an individual can ‘recognize’ different types of opponent and the negotiated outcome can differ depending upon the type of opponent encountered.

Although the mathematical analysis demonstrates an infinite number of Nash equilibrium negotiation strategies, our simulations demonstrate some form of directionality to evolutionary change towards a band of strategies on the left side of the equilibrium curve that result in cooperative behaviour. This can be investigated analytically by using standard techniques for modelling the evolutionary dynamics of  $\lambda^*$  and  $\rho^*$  (Appendix B). For tractability, these techniques typically assume that there is very little genetic variation in the population at any given time. The results of Appendix B show that points on the




Figure 5. Any response line for player 1 (dashed) that intersects the diagonal at  $\hat{x}^*$  will have resident fitness. These are the response lines that correspond to the neutral line in figure 2.




Figure 6. Curves of constant mutant fitness in a population with resident strategy  $\lambda^* = 0.5, \rho^* = 0.223$  at frequency  $p = 0.8$  and mutant strategy  $(\lambda, \rho)$  at frequency  $p = 0.2$ . The level curve  $\Delta W = 0$  (red) consists of the neutral line and a curve that starts at  $(-1, 0)$  on the left, rises above the equilibrium curve, and then crosses it twice, the second time at the equilibrium point  $(\lambda^*, \rho^*)$ . The region between this curve and the neutral line is the region of positive  $\Delta W$  and consists of the two regions marked A and B. The blue contours  $\Delta W = 0.005$  (outer) and  $\Delta W = 0.01$  (inner) are also included.

equilibrium curve are neutrally (and thus not strictly) convergence stable (Eshel 1983; Christiansen 1991). However, this analysis fails to reveal any source of directionality to the evolutionary dynamics.

The above game-theoretic analyses follow standard techniques by examining a two-strategy population with an equilibrium resident strategy  $(\lambda^*, \rho^*)$  of frequency  $1-p$  and a mutant strategy  $(\lambda, \rho)$  of frequency  $p$  and random interactions. One then calculates the fitness difference:  $\Delta W = W(\text{mutant}) - W(\text{resident})$  in the limit where  $p = 0$ . The fact that this approach reveals a line of neutral ( $\Delta W = 0$ ) mutant negotiation strategies, with all other mutant strategies having  $\Delta W < 0$ , suggests an examination of the case where  $p$  is positive. In this case it can be shown that there is always a region adjacent to the neutral mutant line in which  $\Delta W > 0$  (figure 6) for which the mutant can invade.

This finding provides the key to understanding the evolution of negotiation strategies that result in cooperative behaviour. In the presence of variation, there are mutants that tend to produce cooperative outcomes (and thus have higher fitness) when played against themselves (or similar negotiation strategies) but that produce an outcome, and thus fitness, roughly equivalent to that of the resident when played against the resident. Therefore, in the presence of variation, these mutants gain the benefits of cooperation without paying much of a cost when interacting with less cooperative negotiation strategies (figure 6, region A).

Interestingly, there are also mutants that can invade at a positive frequency that produces more selfish outcomes when played against themselves (or similar negotiation strategies; figure 6, region B). Analysis (P. D. Taylor and T. Day, unpublished results) reveals that these invade through a fundamentally different mechanism than the mutants mentioned above. Instead of gaining benefits when interacting among themselves and having approximately resident fitness when interacting with a resident (as above), these 'selfish' negotiation strategies have reduced fitness when interacting with themselves (because they are less cooperative) but they cause the resident negotiation strategy to have even lower fitness when the two interact. These mutants thereby gain their evolutionary advantage through a form of spite, by hurting themselves but hurting the resident type even more.

Overall, those negotiation strategies yielding a more cooperative outcome tend to prevail over the strategies yielding a more selfish outcome (figure 2). This is presumably because the rare cooperative negotiators gain their evolutionary advantage by helping themselves relative to the residents, whereas the rare selfish negotiators gain their evolutionary advantage by hurting the resident relative to themselves. Thus, the rare cooperative negotiators have little effect on the selfish negotiators' fitness, but the selfish negotiators inadvertently further enhance the fitness advantage of the cooperative negotiators by depressing resident fitness.

In a sense these results bring us full circle, back to the ideas of reciprocity, but at a different and potentially more powerful level for explaining the emergence of cooperative behaviour. Previous results on the evolution of cooperation through reciprocity have revealed that strategies such as tit-for-tat are evolutionarily successful because they are nice (by cooperating on first encounters), they are retaliatory (by being selfish if they re-encounter an opponent who was selfish in the past) and they are forgiving (by returning to cooperative behaviour if

previously selfish opponents become cooperative; Wilson 1975). In these previous results, these strategies play out over several successive meetings of opponents, where each meeting produces a fitness outcome. In negotiation games, a similar process occurs. The evolutionarily stable behavioural reaction norm for such games is somewhat analogous to the tit-for-tat strategy from previous theory in that it produces a cooperative outcome when interacting with similar individuals, but it becomes less cooperative if its opponent does so. All of this happens within the context of a single meeting, however, and therefore it potentially provides a robust explanation for cooperative behaviour that requires few assumptions. Under this hypothesis for the evolution of cooperation, we expect a positive relation between the level of cooperative behaviour observed between individuals and the scope that exists for negotiation between them before their fitness outcomes are realized.

This project was supported by the Natural Sciences and Engineering Research Council of Canada. Technical work, the artwork and the execution of the simulations, were performed by Daniel Nagy. Ido Pen provided several valuable comments and verified the simulations.

## APPENDIX A: ANALYSIS OF THE GAME

Consider the family of level curves of  $F$  which have a slope between  $-1$  and  $1$  where they cross the diagonal at  $(x^*, x^*)$  (figure 4a). Take any one of these and let the tangent at  $(x^*, x^*)$  have slope  $-\lambda^*$  and  $x_2$ -intercept  $\rho^*$  (figure 4b). Because the level curve of  $F$  is concave-up (this follows from the assumptions that  $B' > 0$ ,  $B'' \leq 0$ ,  $K' > 0$ ,  $K'' > 0$ ) no rare mutant strategy in a  $(\lambda^*, \rho^*)$  population can have greater than normal fitness. This provides an infinite curve of negotiation strategies that are Nash equilibria: this is a type of population-wide neutrality. In figure 2, this curve is displayed in  $(\lambda, \rho)$  space. In the original game of Houston & Davies (1985) without negotiation the ESS contribution  $x^*$  corresponds to the point at which the level curve of  $F$  crosses the diagonal horizontally (figure 4a). This corresponds to the point at which the equilibrium curve of figure 2 crosses the vertical axis.

There is also a type of mutant neutrality that occurs at any of the above-mentioned Nash equilibrium negotiation strategies. Specifically, at any such strategy there are mutant strategies having fitness identical to the resident fitness. These have a response rule passing through  $(x^*, x^*)$  (figure 5). In figure 2, this set of strategies appears as a line in  $(\lambda, \rho)$  space. Indeed, all points on this line have the same value of  $\rho/(1 + \lambda)$  and if any two such points are put into the  $\hat{x}$ -equation (2.6) they yield  $\hat{x}_1 = \hat{x}_2 = x^* = \rho^*/(1 + \lambda^*)$ .

## APPENDIX B

If the strategy mix in the population is closely concentrated about its mean value  $(\bar{\lambda}, \bar{\rho})$ , then the fitness of a  $(\lambda, \rho)$  individual can be reasonably approximated by  $W(\lambda, \rho, \bar{\lambda}, \bar{\rho})$  and the direction and speed of evolutionary change of the population mean is modelled with the standard evolutionary dynamic (Abrams *et al.* 1993;




Figure 7. The vector field of the dynamic equations measures the direction and speed of evolutionary change at any  $(\rho, \lambda)$ . The vectors are always perpendicular to the line drawn from the point  $(-1, 0)$  and point up when below the equilibrium curve and otherwise down. At points on the equilibrium curve, the field is zero.

Dieckmann & Law 1996; Geritz *et al.* 1998; Hofbauer & Sigmund 1998):

$$\frac{d\bar{\lambda}}{dt} = k \frac{\partial W}{\partial \lambda} \Big|_{\substack{\lambda=\bar{\lambda} \\ \rho=\bar{\rho}}} = -k \frac{\bar{\rho}}{1+\bar{\lambda}} \frac{F_1 - \bar{\lambda}F_2}{1-\bar{\lambda}^2},$$

$$\frac{d\bar{\rho}}{dt} = k \frac{\partial W}{\partial \rho} \Big|_{\substack{\lambda=\bar{\lambda} \\ \rho=\bar{\rho}}} = k \frac{F_1 - \bar{\lambda}F_2}{1-\bar{\lambda}^2},$$

where  $F_i$  is the partial derivative of  $F$  with respect to  $x_i$  evaluated at  $x_1 = x_2 = \bar{\rho}/1 + \bar{\lambda}$ .

This dynamic has a simple geometric interpretation (figure 7). Because the line through  $(-1, 0)$  and  $(\bar{\lambda}, \bar{\rho})$  has slope  $\bar{\rho}/1 + \bar{\lambda}$ , the vector field of the dynamic is orthogonal to this line at any  $(\bar{\lambda}, \bar{\rho})$  and is directed up when below the equilibrium curve ( $F_1 > \bar{\lambda}F_2$ ), down when above the equilibrium curve ( $F_1 < \bar{\lambda}F_2$ ) and is zero on the equilibrium curve ( $F_1 = \bar{\lambda}F_2$ ). In particular, the equilibrium curve is exactly the set of equilibrium points of the dynamic. Not surprisingly, the system is neutrally stable at each such point. [An equilibrium point of the dynamical system is stable if the matrix

$$A = \begin{bmatrix} \dot{\lambda}_\lambda & \dot{\lambda}_\rho \\ \dot{\rho}_\lambda & \dot{\rho}_\rho \end{bmatrix}_{\substack{\lambda=\bar{\lambda} \\ \rho=\bar{\rho}}}$$

is a stability matrix, where the dot denotes the time derivative and the subscripts denote partial differentiation. This will be the case if  $\text{trace}(A) < 0$  and  $\text{det}(A) > 0$ . A straightforward calculation shows that the first holds but that  $\text{det}(A) = 0$ .]

It is clear from figure 7 that the dynamic points in opposite directions on either side of the equilibrium curve, and thus it is not clear in which direction the combined effects of local mutation and selection might cause the strategy mix to move.

## REFERENCES

- Abrams, A., Harada, H. & Matsuda, H. 1993 On the relationship between quantitative genetic and ESS models. *Evolution* **47**, 982–985.
- Agrawal, A. A. 2001 Phenotypic plasticity in the interactions and evolution of species. *Science* **294**, 321–326.
- Axelrod, R. & Hamilton, W. D. 1981 The evolution of cooperation. *Science* **211**, 1390–1396.
- Christiansen, F. B. 1991 On conditions for evolutionary stability for a continuously varying character. *Am. Nat.* **138**, 37–50.
- Dieckmann, U. & Law, R. 1996 The dynamical theory of coevolution: a derivation from stochastic ecological processes. *J. Math. Biol.* **34**, 579–612.
- Eshel, I. 1983 Evolutionary and continuous stability. *J. Theor. Biol.* **103**, 99–111.
- Geritz, S. A. H., Kisdi, E., Meszina, G. & Metz, J. A. J. 1998 Evolutionarily singular states and the adaptive growth and branching of the evolutionary tree. *Evol. Ecol.* **12**, 35–57.
- Hofbauer, J. & Sigmund, K. 1998 *Evolutionary games and population dynamics*. Cambridge University Press.
- Houston, A. I. & Davies, N. B. 1985 The evolution of cooperation and life history in the dunnock *Prunella modularis*. In *Behavioural ecology: the ecological consequences of adaptive behaviour* (ed. R. M. Sibly & R. H. Smith), pp. 471–487. Oxford: Blackwell Scientific.
- Lotem, A., Fishman, M. A. & Stone, L. 1999 Evolution of cooperation between individuals. *Nature* **400**, 226–227.
- McNamara, J. M., Gasson, C. E. & Houston, A. I. 1999 Incorporating rules for responding into evolutionary games. *Nature* **401**, 368–371.
- Milinski, M., Semmann, D. & Krambeck, H.-J. 2002 Reputation helps solve the ‘tragedy of the commons’. *Nature* **415**, 424–426.
- Nowak, M. A. & Sigmund, K. 1998 Evolution of indirect reciprocity by image scoring. *Nature* **393**, 573–577.
- Riolo, R. L., Cohen, M. D. & Axelrod, R. 2001 Evolution of cooperation without reciprocity. *Nature* **414**, 441–443.
- Wilson, D. S. 1975 A theory of group selection. *Proc. Natl. Acad. Sci. USA* **72**, 143–146.