

The effect of group size and frequency-of-encounter on the evolution of cooperation

Steve Phelps¹, Gabriel Nevarez², and Andrew Howes²

1. Centre for Computational Finance and Economic Agents (CCFEA),
University of Essex, UK
sphelps@essex.ac.uk
2. Manchester Business School, University of Manchester, UK

Abstract. We introduce a model of the evolution of cooperation in groups which incorporates both conditional direct-reciprocity (“tit-for-tat”), and indirect-reciprocity based on public reputation (“conspicuous altruism”). We use ALife methods to quantitatively assess the effect of changing the group size and the frequency with which other group members are encountered. We find that for moderately sized groups, although conspicuous altruism plays an important role in enabling cooperation, it fails to prevent an exponential increase in the level of the defectors as the group size is increased, suggesting that economic factors may limit group size for cooperative ecological tasks such as foraging.

1 The Model

Pairs of agents $(a_i, a_j) : i \neq j$ are drawn at random from $A = \{a_1, a_2, \dots, a_n\}$, and engage in bouts of grooming at different time periods $t \in \{0, 1, \dots, N\}$. We refer to n as the *group size* and N as the *frequency-of-encounter*.

At each time period t the groomer a_i may choose to invest a certain amount of effort $u_{(i,j,t)} \in [0, U] \subset \mathbb{R}$ in grooming their partner a_j , where $U \in \mathbb{R}$ is a parameter determining the maximum rate of grooming. This results in a negative fitness payoff $-u$ to the groomer, and a positive fitness payoff ku to the partner a_j :

$$\begin{aligned}\phi_{(j,t+1)} &= \phi_{(j,t)} + k \cdot u_{(i,j,t)} \\ \phi_{(i,t+1)} &= \phi_{(i,t)} - u_{(i,j,t)}\end{aligned}$$

where $\phi_{(i,t)} \in \mathbb{R}$ denotes the fitness of agent a_i at time t , and $k \in \mathbb{R}$ is a constant parameter.

In an ecological context, the positive fitness payoff ku might represent, for example, the fitness gains from parasite elimination, whereas the fitness penalty $-u$ would represent the opportunity cost of foregoing other activities, such as foraging, during the time u allocated for grooming.

Since we are interested in the *evolution* of cooperation, we analyse outcomes in which agents choose values of u that maximise their own fitness ϕ_i . Provided

that $k > 1$, over many bouts of interaction it is possible for agents to enter into reciprocal relationships that are mutually-beneficial, since the groomer’s initial cost u may be reciprocated with ku yielding a net benefit $ku - u = u(k - 1)$. Provided that agents reciprocate, they can increase their net benefit by investing larger values of u . However, by increasing their investment they put themselves more at risk from exploitation, since just as in the alternating prisoner’s dilemma [5], defection is the dominant strategy if the total number of bouts N is known: the optimal behavior is to accept the benefits of being groomed without investing in grooming in return. In the case where N is *unknown*, and the number of agents is $n = 2$, it is well known that conditional reciprocation is one of several evolutionary-stable solutions in the form of the so-called *tit-for-tat* strategy which copies the action that the opposing agent chose in the preceding bout at $t - 1$ [4]. However, this result does not generalise to larger groups $n > 2$.

Nowak and Sigmund [7] demonstrate that reciprocity can emerge *indirectly* in large groups, provided that information about each agent’s history of actions is summarised and made publicly available in the form of a reputation or “image-score” $r_{(i,t)} \in [r_{\min}, r_{\max}] \subset \mathbb{Z}$. The image-score r_i summarises the propensity-to-cooperate of agent a_i . As in the Nowak and Sigmund model, image scores in our model are initialised $\forall_i r_{(i,0)} = 0$ and are bound at $r_{\min} = -5$ and $r_{\max} = 5$. An agent’s image score is incremented at $t + 1$ if the agent invests a non-zero amount at time t , otherwise it is decremented:

$$r_{(i,t+1)} = \begin{cases} \min(r_{(i,t)} + 1, r_{\max}) & : u_{(i,x,t)} > 0 \\ \max(r_{(i,t)} - 1, r_{\min}) & : u_{(i,x,t)} = 0 \end{cases}$$

and agents invest conditionally on their partner’s image score:

$$u_{(i,j,t)} = \begin{cases} \gamma & : r_{(j,t)} \geq \sigma_i \\ 0 & : r_{(j,t)} < \sigma_i \end{cases}$$

where σ_i is a parameter determining the threshold image score above which agent a_i will cooperate, and $\gamma \in \mathbb{R}$ is a global parameter (as in [7] we use $\gamma = 10^{-1}$ and $k = 10$).

Nowak and Sigmund [6] demonstrate that widespread defection is avoided if, and only if, the initial proportion of agents using a discriminatory¹ strategy is above a critical value, implying that strategies based on indirect reciprocity via reputation are an essential prerequisite for the evolution of cooperation in large groups.

We are interested in the effect of group size n and interaction frequency N on the evolution of cooperation. The analytical model of Nowak and Sigmund [6] assumes: a) that the group size n is large enough relative to N that strategies based on private history, such as *tit-for-tat*, are irrelevant (since the probability of encountering previous partners is very small); and b) that we do not need to take into account the fact that an agent cannot cooperate with itself when calculating the probability with which any given agent is likely to encounter a

¹ discriminatory strategies cooperate only if their partner’s image score is non-negative, that is: $\sigma_i = 0$.

particular strategy. However, in order to model changes in group size, and hence interaction in smaller groups, it is necessary to drop both of these assumptions. The resulting model is more complicated, and it is difficult to derive closed-form solutions for the equilibrium behaviour. Therefore we use ALife simulation to estimate payoffs, and numerical methods to compute asymptotic outcomes, as described in the next section.

2 Methodology

In order to study the evolution of populations of agents using the above strategies, we use both ALife methods and mathematical modeling based on evolutionary game-theory. However, rather than considering pairs of agents, our analysis concerns interactions amongst *groups* of size $n > 2$ assembled from a larger population of individuals. The resulting game-theoretic analysis is complicated by the fact that this results in a many-player game, which presents issues of tractability for the standard methods for computing the equilibria of normal-form games. A popular ALife approach to this issue is to use *Co*-evolutionary algorithms [3, 4]. In a co-evolutionary optimisation, the fitness of individuals in the population is evaluated relative to one another in joint interactions (similarly to payoffs in a strategic game), and it is suggested that in certain circumstances the converged population is an approximate Nash solution to the underlying game; that is, the stable states, or equilibria, of the co-evolutionary process are related to the evolutionary stable strategies (ESS) of the corresponding game. However, there are many caveats to interpreting the equilibrium states of standard co-evolutionary algorithms as approximations of game-theoretic equilibria, as discussed in detail by Sevan Ficici [1, 2]. In order to address this issue, we adopt a methodology called *empirical game-theory* [10, 12], which uses a combination of simulation and rigorous game-theoretic analysis. The empirical game-theory method uses a *heuristic* payoff matrix which is calibrated by running many simulations, as detailed below.

We can make one important simplification by assuming that the game is symmetric, and therefore that the payoff to a given strategy depends only on the *number* of agents within the group adopting each strategy. Thus for a game with j strategies, we represent entries in the payoff matrix as vectors of the form $\mathbf{p} = (p_1, \dots, p_j)$ where p_i specifies the number of agents who are playing the i^{th} strategy. Each entry $\mathbf{p} \in P$ is mapped onto an outcome vector $\mathbf{q} \in Q$ of the form $\mathbf{q} = (q_1, \dots, q_j)$ where q_i specifies the expected payoff to the i^{th} strategy.

For each entry in the payoff matrix we estimate the expected payoff to each strategy by running 10^5 ALife simulations and taking the mean² fitness.

With estimates of the payoffs to each strategy in hand, we are in a position to model the evolution of populations of agents using these strategies. In our evolutionary model, we do not restrict reproduction to within-group mating; rather, we consider a larger population which temporarily forms groups of size

² We take the average fitness of every agent adopting the strategy for which we are calculating the payoff, and then also average across simulations.

n in order to perform some ecological task. Thus we use the standard replicator dynamics equation [11] to model how the frequency of each strategy in the larger population changes over time in response to the within-group payoffs:

$$\dot{m}_i = [u(e_i, \mathbf{m}) - u(\mathbf{m}, \mathbf{m})] m_i \quad (1)$$

where \mathbf{m} is a mixed-strategy vector, $u(\mathbf{m}, \mathbf{m})$ is the mean payoff when all players play \mathbf{m} , and $u(e_i, \mathbf{m})$ is the average payoff to pure strategy i when all players play \mathbf{m} , and \dot{m}_i is the first derivative of m_i with respect to time. Strategies that gain above-average payoff become more likely to be played, and this equation models a simple co-evolutionary process of adaptation. Since mixed strategies represent population frequencies, the components of \mathbf{m} sum to one. The geometric corollary of this is that the vectors \mathbf{m} lie in the *unit-simplex* $\Delta^{j-1} = \{\mathbf{x} \in \mathbb{R}^j : \sum_{i=1}^j x_i = 1\}$. In the case of $j = 3$ strategies the unit-simplex Δ^2 is a *two-dimensional* triangle embedded in three-dimensional space which passes through the coordinates corresponding to pure strategy mixes: $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$. We shall use a two dimensional projection of this triangle to visualise the population dynamics in the next section³.

In our experiments we solve this system numerically: we choose 10^3 randomly sampled initial values which are chosen uniformly from the unit simplex [9], and for each of these initial mixed-strategies we solve Equation 1 as an initial value problem using MATLAB's `ode15s` solver [8]. This results in 10^3 trajectories which either terminate at stationary points, or enter cycles.

We consider $j = 5$ strategies:

1. *C* which cooperates unconditionally ($\sigma_i = r_{min}$);
2. *D* which defects unconditionally ($\sigma_i = r_{max} + 1$);
3. *S* which cooperates conditionally with agents who have a good reputation ($\sigma_i = 0$) but cooperates unconditionally in the first round (when reputations have not yet been established);
4. *Sd* which cooperates conditionally ($\sigma_i = 0$) but defects unconditionally in the first round of play;
5. *T4T* which cooperates conditionally with agents who have cooperated in previous rounds, and cooperates unconditionally against unseen opponents.

3 Results

Initially we restrict attention to $j = 3$ strategies and 10^2 initial values, allowing us to more easily visualise the population dynamics and to compare our results with that of Nowak and Sigmund [6] (who assume a large group size n relative to N).

When we exclude discriminators and consider only cooperator (*C*), defectors (*D*) and tit-for-tat (*T4T*) we find that defection is the dominant strategy regardless of N , implying that conditional reciprocity cannot sustain group cooperation in the absence of reputation.

³ See [11, pp. 3–7] for a more detailed exposition of the geometry of mixed-strategy spaces.





Fig. 1. Direction field for $n = 10$ agents and $N = 13$ pairwise interactions per generation (above) compared with $N = 100$ (below). C denotes unconditional altruists, D unconditional defectors and S discriminators who cooperate in the first round. Each line represents a trajectory whose termination is represented by an open circle. The arrows show the direction of change.




Fig. 2. Mean frequency of each strategy in equilibrium as the number of pairwise interactions per generation N is increased relative to n . The error bars show the confidence interval for $p = 0.05$. C denotes the proportion of unconditional altruists; D unconditional defectors; S discriminators who cooperate in the first round; Sd discriminators who defect in the first round; $T4T$ the tit-for-tat strategy

We obtain more subtle results when we introduce reputation-based strategies. Figure 1 shows the phase diagram for the population frequencies when we analyse the interaction between cooperators (C), defectors (D) and discriminators (S) when we have a small group of $n = 10$ agents. As in [6] we find that a minimum initial frequency of discriminators (y axis) is necessary to prevent widespread convergence to the defection strategy (in the bottom right of the simplex). However, the results of our model differ in two important respects.

Firstly, when the critical threshold of discriminators is reached, our model results in various stationary mixes of discriminators and cooperators with a total absence of defection, and no limit cycles. This is in contrast to [6] where the population cycles endlessly between all three strategies if the critical threshold is exceeded.

Secondly, the behaviour of our model is sensitive to the number of pairwise interactions per generation N : as N is increased from $N = 13$ to $N = 100$ we see that the basin of attraction of the pure defection equilibrium is significantly decreased, and correspondingly the critical threshold of initial discriminators necessary to avoid widespread defection. Our intuitive interpretation of these results is that defection is less likely⁴ as we increase the frequency of interaction relative to the group size.

⁴ Assuming that all points in the simplex are equally likely as initial values.

As discussed in Section 1, if we increase N relative to n we need to consider the effect of strategies that take into account private interaction history as well as strategies that are based on public reputation. Figure 2 shows the mean frequency in equilibrium of each strategy when we analyse all five strategies and systematically vary N while holding n fixed. We plot the equilibrium population frequency against $\frac{N}{n}$, and obtain the same graph for both $n = 10$ and $n = 20$ agents. This suggests that the ratio $\frac{N}{n}$ determines the asymptotic behaviour.

4 Discussion

The frequency of both unconditional cooperation and discrimination increases with N , and these strategies become more prevalent than discriminators for $\frac{N}{n} > 2$. As we would intuitively expect, for $\frac{N}{n} > 1$ discriminators become more prevalent as we increase group size or decrease frequency of interaction. However, this is not sufficient to prevent free-riding. Most striking is that the likelihood of defection decreases exponentially as we increase the number of interactions per generation N . Correspondingly, as we increase the group size n we observe an *exponential increase* in the level of defection.

If we consider the possibility of inter-group competition and hence group selection, then since the expected frequency of defectors in equilibrium determines the per-capita fitness of the agents in the overall population⁵, we can interpret our results as showing how group fitness changes as a function of group size (n) and frequency-of-encounter (N). That is, for any given N we can determine the optimum group size n .

Our results indicate that the stylised cooperation task described by our model introduces a very strong selection pressure for smaller group sizes. Of course, this task is not the only ecological task which influences per-capita fitness for any given species. For example, our model could be used in conjunction with optimal foraging models to derive a comprehensive model of optimum group size for a particular species in a particular niche. Our main contribution is to highlight that economic factors play a significant role in determining optimal group size, when other ecological tasks such as foraging favour group sizes that are relatively small compared with the frequency-of-encounter.

5 Conclusion

Our model predicts that neither reputation nor conditional punishment are sufficient to prevent free-riding as the group size increases. Although both types of strategy play an important role, as the group size increases the level of conspicuous altruism based on reputation rises, but defection rises faster. Thus, when other tasks already favour smaller groups, we predict that economic factors will limit the maximum group size independently of the group size favoured by other niche-specific tasks.

⁵ In the absence of defection all other strategies are able to gain the maximum available surplus.

Acknowledgements

We are grateful for the comments of the anonymous reviewers, and for financial support received from the UK EPSRC through the *Theory for Evolving Socio-cognitive Systems* project (EP/D05088X/1).

References

1. S. G. Ficci and J. B. Pollack. Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In *Proceedings of ALIFE-6*, 1998.
2. S. G. Ficci and J. B. Pollack. A game-theoretic approach to the simple coevolutionary algorithm. In Hans-Paul Schwefel Marc Schoenauer, Kalyanmoy Deb, Günter Rudolph, Xin Yao, Evelyne Lutton, Juan Julian Merelo, editor, *Parallel Problem Solving from Nature — PPSN VI 6th International Conference*, Paris, France, 16-20 2000. Springer Verlag.
3. W. D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. In Langton et al., editor, *Proceedings of ALIFE-2*, pages 313–324. Addison Wesley, 1992.
4. J. H. Miller. The coevolution of automata in the repeated Prisoner’s Dilemma. *Journal of Economic Behavior and Organization*, 29(1):87–112, Jan. 1996.
5. M. A. Nowak and K. Sigmund. The alternating prisoner’s dilemma. *Journal of theoretical Biology*, 168:219–226, 1994.
6. M. A. Nowak and K. Sigmund. The dynamics of indirect reciprocity. *Journal of Theoretical Biology*, 194(4):561–574, Oct. 1998.
7. M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 383:537–577, 1998.
8. L. F. Shampine and M. W. Reichelt. The MATLAB ODE suite. http://www.mathworks.com/access/helpdesk/help/pdf_doc/otherdocs/ode_suite.pdf, 2009.
9. R. Stafford. Random vectors with fixed sum. <http://www.mathworks.com/matlabcentral/fileexchange/9700>, Jan. 2006.
10. W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent games. In *AAAI-02 Workshop on Game Theoretic and Decision Theoretic Agents*, 2002. <http://wewalsh.com/papers/MultiAgentGame.pdf>.
11. J. W. Weibull. *Evolutionary Game Theory*. MIT Press, First MIT Press edition, 1997.
12. M. P. Wellman. Methods for Empirical Game-Theoretic Analysis. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, pages 1152–1155, 2006.