Reciprocal cooperation can be stable when the presence of defectors selects for discrimination

Gilbert Roberts

Northumberland, UK

Reciprocity is a pillar of our understanding of how cooperation works. However, this paradigm has been challenged by Efferson et al (Efferson et al., 2024) who claim that repeated interaction is insufficient to support cooperation. The authors base their conclusion on a model in which reciprocators use a novel function to determine their response. They find that cooperation through repeated interaction breaks down. I point out here that the collapse of cooperation is not the result of the novel strategy set. Instead, it is a standard result in models of cooperation when a benign environment allows drift towards undiscriminating strategies. I show how and when introducing an element of defection into a cooperative environment provides selection for the maintenance of discrimination. Discriminating strategies can resist invasion by strategies that would exploit undiscriminating cooperation. In consequence, cooperation is stabilized. I conclude that Efferson et al.'s results can readily be explained by existing theory and that repeated interaction can support cooperation without additional processes such as between group competition.

Reciprocity is a classical explanation for cooperation involving repeated interactions where each move offers a temptation to defect (Trivers, 1971). It works when discriminating strategies, most famously 'Tit-for-Tat' (Axelrod & Hamilton, 1981), cooperate with cooperators but defect on defectors. However, this scenario has an Achilles heel. If cooperative strategies are profitable over multiple rounds, all agents will tend to evolve to cooperate through some form of reciprocation. In this cooperative utopia, the discrimination that makes strategies like Tit-for-Tat successful in an uncooperative environment becomes redundant and unconditionally cooperative strategies can drift in by mutation. This in turn paves the way for uncooperative strategies to invade by exploiting the underlying naivety (Bendor & Swistak, 1995). As a result, evolutionary dynamics of cooperation can take the form of 'boom and bust' cycles. Conditionally cooperative strategies pave the way for

unconditionally cooperative environments which in turn can be invaded by defection (Wahl & Nowak, 1999). These authors go as far as to note that such dynamics “are the underlying basis for every model of cooperation”. The drift of undiscriminating cooperators constitutes an ‘indirect invasion’ (van Veelen, 2012).

However, these dynamics are associated with rather unrealistic model assumptions. To make their models more realistic, researchers have assumed that a uniformly cooperative system is unlikely and that some defection is retained. This may be the result of genetic mutation or due to divergence between the genetic strategy and the actual behaviour employed, perhaps because of limited resources, or errors of perception or execution. Simulations have shown that introducing defection in the form of phenotypic defectors, which carry a cooperative genetic strategy yet defect due to being unable to cooperate, then selection acts to maintain discrimination amongst cooperators (Lotem et al., 1999; Sherratt & Roberts, 2001). The reason for this is simple. When defectors are present, discriminating strategies can do better than undiscriminating strategies. Thus, paradoxically, defection allows cooperation.

Now, Efferson et al. (Efferson et al., 2024) report uncovering a flaw in the paradigm of repeated interaction as a pillar of cooperation. They suggest that researchers have not considered the full strategy space available to agents when they can continuously vary their cooperative investments and so have overlooked a weakness in reciprocity. Efferson et al. explain the process involved in the collapse of cooperation in their model. In a cooperative population, individuals tend to start with high offers. This means there is little scope for escalation and consequently the parameters determining the degree of escalation in their model have little effect. Relieved from selection, the escalation parameters drift into a range where they are vulnerable to what they call ‘ambiguous’ strategies. These novel strategies can short-change the resident when playing the role of second mover. As a result, cooperation declines to low levels.

Here, I investigate whether the breakdown of cooperation observed by Efferson et al. could be prevented by introducing defecting strategies. I consider three ways in which defection might reasonably occur.

(1) Phenotypic defectors. As noted above, the logic behind introducing phenotypic defectors into the model is that not every individual that carries a cooperative genotype may have the resources or ability to perform cooperative acts.
(2) Mutation to zero. Efferson et al. included only a local mutation process. Supplementing this with a process in which individuals fully defected by giving zero allows strategies of cooperation to be tested against defectors rather than just those that give a little bit more or less.
(3) Errors in which give zero. Efferson et al. included small errors but including errors in which individuals give nothing means that strategies are tested against full defection.

Note that these processes are quite different from introducing symmetrical small errors or 'noise' in that they are unidirectional and result in complete defection as opposed to small variations around a central parameter value.

Methods

Simulation methods followed those described in Efferson et al. and were coded independently in Microsoft Visual C. Simulations were initialized by setting up a population of 100 individuals. Note that Efferson et al used groups of 24 but I used larger groups as these are more appropriate when looking for the effects of rare events and to limit the effects of stochastic effects due to small population sizes. Play involved selecting every possible pair in turn (100*99/2 = 4950 pairings) and then for each pair, playing an iterated asynchronous continuous Prisoner's Dilemma game for 100 rounds. On the first move, the leader was determined randomly. The leader donated part of its 1 unit of endowment according to its genetic strategy *x*, which was in increments of 0.01 between 0 and 1.

The follower donated according to its response rule $(d\text{-}a)x$, where a and d were the left and right intercepts of a linear function of x. In subsequent rounds, both partners used this function to determine their response based on what their partner had donated. At the end of the full set of interactions, reproduction was proportional to payoffs and resulted in replacing the population of N=100 individuals. During reproduction, mutation occurred whereby the genetic strategies increased or decreased in steps of 0.01 with probability 0.02 within the bounds of 0 to 1. Simulations began with full cooperation given by x=1, a=0, d=1. Simulations ran for 100,000 generations of strategy evolution.

Within the simulations, defection (donating zero) was introduced in the three ways described above by doing the following.

(1) Phenotypic defectors (PDs). For every individual, their genetic strategy was copied to a phenotypic strategy. In the case of the 5% PDs, this was then overwritten by a strategy of giving zero on every move. Reproduction was by the genetic strategy (so all strategy fitnesses will reflect the fact that a small proportion of individuals defected on every move).

(2) Genetic mutation to zero. Mutation in Efferson et al. was implemented by adding or subtracting small increments to a strategy. This means that in a cooperative environment, cooperators are not tested against defectors. To make the model more biologically realistic I added a mutation process in which there was a 0.02 probability of genetic strategies mutating to zero.

(3) Errors in which individuals give zero. Here there is a probability 0.02 that instead of giving the amount determined by the genetic strategy an individual gives zero. Note that this mechanism differs from the others in that defection occurs on a single move whereas in the other two cases individuals which defect do so on every move with every other player.

Results

I first replicated the collapse of cooperation reported by Efferson et al. (Figure 1a). However, after introducing defection in the form of PDs, I found that cooperation was readily maintained at a stable level (Figure 1b). Comparing the levels of cooperation that had evolved at generation 100,000, the median donation from an endowment of 1, averaged across 20 simulations was 0.01 for the replication of Efferson et al. rising to 0.78 when 5% phenotypic defectors were introduced. Mutations to zero also caused a major increase in cooperation, with a median of 0.59. However, errors of giving zero did not increase cooperation, the median being 0.01 (Figure 2).

Discussion

From these results it is apparent that introducing a small proportion of individuals that always defect has the paradoxical effect of increasing cooperation. This effect does not occur when individuals only defect occasionally. In this case, withholding cooperation in response can only harm the relationship. In the case of permanent defectors, partners should respond by withholding cooperation. Analytically, this can readily be understood by extension from the binary model. In that case, discriminators do better against those who always defect because they suffer less exploitation that unconditional cooperators (Sherratt & Roberts, 2001). In the continuous model, there should be a strong positive relationship between how much a partner has previously given and how much you donate in response. In Efferson et al.'s 3-parameter model this essentially means having a low *a* and a high *d*.

I have shown that introducing defection can facilitate the maintenance of stable cooperation in a model with continuous investment and an extended strategy set including 'ambiguous reciprocators'. This is consistent with previous studies based on binary choice models and is well understood within the theoretical framework of reciprocity theory. Therefore Efferson et al.'s bold conclusion that cooperation does not work with repeated interaction should be qualified by saying that this only holds when cooperative strategies evolve to be unconditional. Their claim that repeated interaction is

insufficient to maintain cooperation depends upon the assumption that all individuals always have the resources to cooperate to a maximal extent and that mutation to full defection does not occur. The existence of a small proportion of individuals that never give anything is sufficient to mean that Efferson et al's central conclusion does not stand.

In sum, Efferson et al.'s point is that their models demonstrate the fragility of cooperation through repeated interactions alone and this forms the basis for advancing between-group competition as a decisive process in human social evolution. However, it turns out that Efferson et al.'s result is not novel, nor is it a product of their shift to a continuous investment model with a novel strategy set. This is important because the results have been interpreted as demonstrating the role of cultural group selection in human uniqueness (Mathew, 2024).

Acknowledgments

This preprint is a companion paper to a critique first submitted by T Sherratt and GR in June 2024. This preprint focuses on the issues of drift, defection and discrimination whilst the companion paper focuses on an alternative modelling framework. P Barclay and co-authors have independently discussed similar ideas.

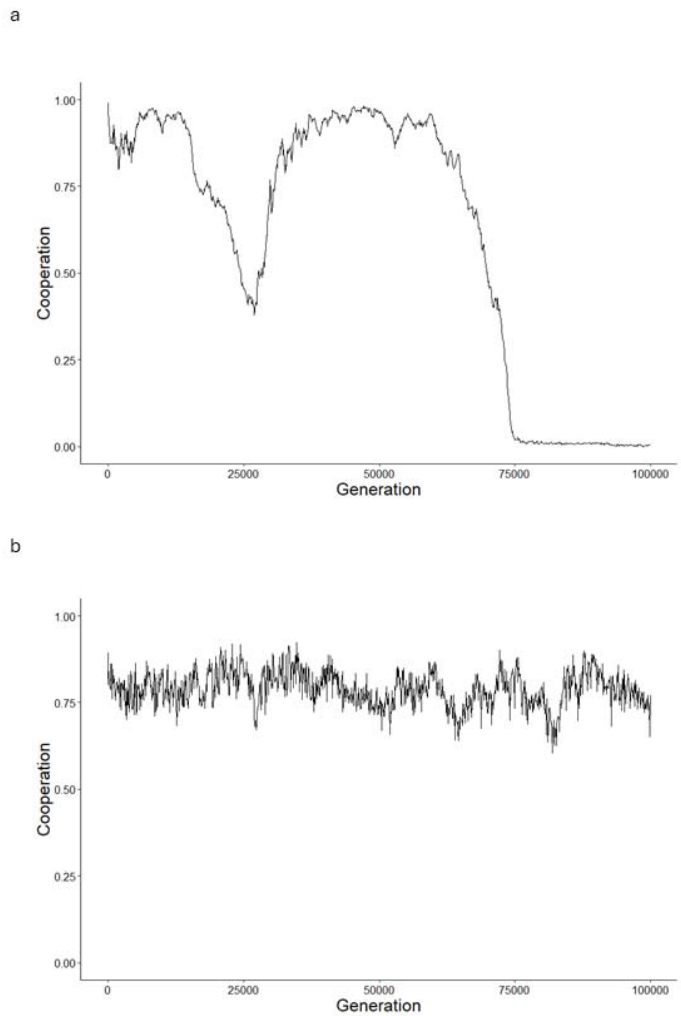

Figure 1. Evolutionary dynamics. (a) A replication of Efferson *et al.*. (b) With 5% phenotypic defectors (PDs). In each case, the first of 20 simulations is presented. Cooperation is calculated as the proportion of the endowment donated to the partner averaged over the ((100*99/2) = 4950 pairings * 100 rounds * 2 in a pair) 990,000 moves in each of the 100,000 generations.

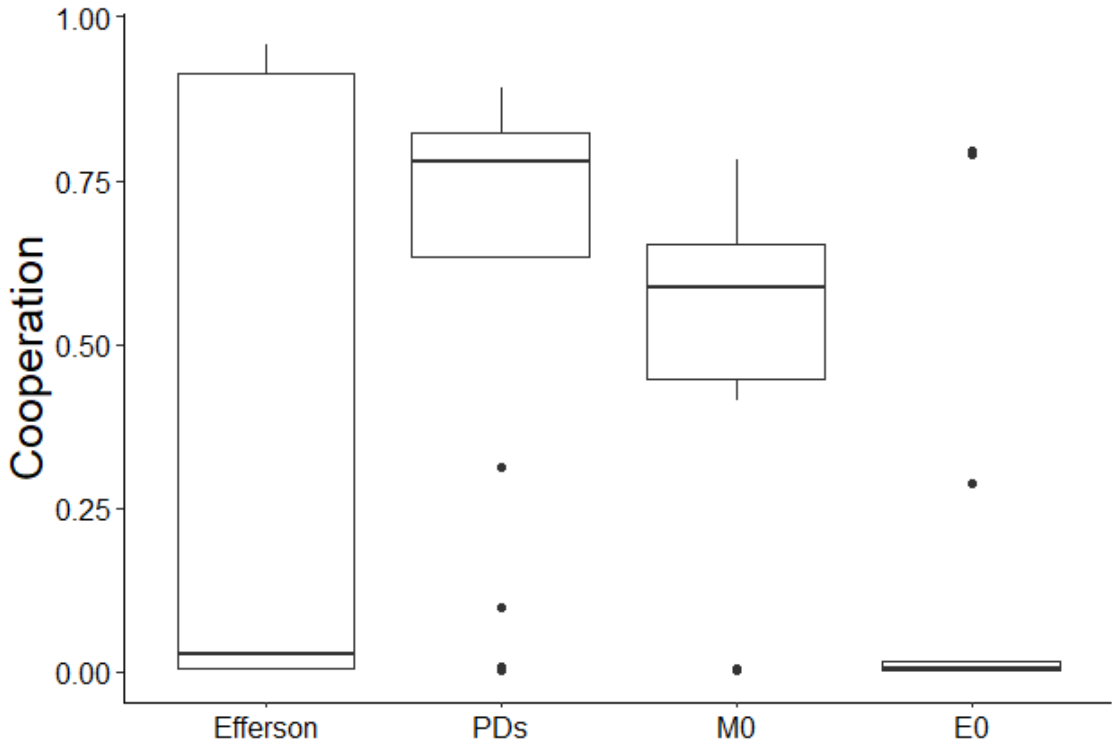

Figure 2. Cooperation in relation to introduced defection. The Figure compares mean cooperation in generation 100,000 where cooperation is the proportion of the 1 unit of endowment donated to the partner calculated across the 100*99/2 = 4500 pairings each with 2 individuals and 100 rounds. Boxplots show the distributions of these means over 20 simulations for each treatment. Efferson: replication of Efferson *et al.* model. PDs: phenotypic defectors were introduced at 5%. M0: in addition to the local mutation process which allows evolution between strategies, mutations resulting in individuals carrying a genetic strategy of giving 0 were introduced at a frequency of 2%. E0: on every move, responses were subject to errors at 2% in which individuals gave 0.

References

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

Bendor, J., & Swistak, P. (1995). Types of evolutionary stability and the problem of cooperation. *Proceedings of the National Academy of Sciences*, *92*(8), 3596-3600. https://doi.org/10.1073/pnas.92.8.3596

Efferson, C., Bernhard, H., Fischbacher, U., & Fehr, E. (2024). Super-additive cooperation. *Nature*, *626*(8001), 1034-1041. https://doi.org/10.1038/s41586-024-07077-w

Lotem, A., Fishman, M. A., & Stone, L. (1999). Evolution of cooperation between individuals. *Nature*, *400*(6741), 226-227.

Mathew, S. (2024). Why reciprocity is common in humans but rare in other animals. *Nature*, *626*, 955-956. https://doi.org/doi: 10.1038/d41586-024-00308-0

Sherratt, T. N., & Roberts, G. (2001). The role of phenotypic defectors in stabilizing reciprocal altruism. *Behavioral Ecology*, *12*, 313-317.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, *46*, 35-57.

van Veelen, M. (2012). Robustness against indirect invasions. *Games and Economic Behavior*, *74*(1), 382-393. https://doi.org/https://doi.org/10.1016/j.geb.2011.05.010

Wahl, L. M., & Nowak, M. A. (1999). The Continuous Prisoner's Dilemma: II. Linear Reactive Strategies with Noise. *Journal of Theoretical Biology*, *200*(3), 323-338. https://doi.org/https://doi.org/10.1006/jtbi.1999.0997