Punishment and reputation in spatial public goods games

Hannelore Brandt¹, Christoph Hauert^{1,2*} and Karl Sigmund^{1,3}

¹*Institut für Mathematik, Universität Wien, Strudlhofgasse 4, A-1090 Vienna, Austria*

²*Department of Zoology, University of British Columbia, University Boulevard 6270, Vancouver, British Columbia V6T 1Z4, Canada*

³*International Institute for Applied Systems Analysis, A-2361 Laxenburg, Austria*

The puzzle of the emergence of cooperation between unrelated individuals is shared across diverse fields of behavioural sciences and economics. In this article we combine the public goods game originating in economics with evolutionary approaches traditionally used in biology. Instead of pairwise encounters, we consider the more complex case of groups of three interacting individuals. We show that territoriality is capable of promoting cooperative behaviour, as in the case of the Prisoner's Dilemma. Moreover, by adding punishment opportunities, the readiness to cooperate is greatly enhanced and asocial strategies can be largely suppressed. Finally, as soon as players carry a reputation for being willing or unwilling to punish, highly cooperative and fair outcomes are achieved. This group-beneficial result is obtained, intriguingly, by making individuals more likely to exploit their co-players if they can get away with it. Thus, less-cooperative individuals make more-cooperative societies.

Keywords: evolutionary game theory; spatial games; cooperation; punishment; reputation

The theory of many person games may seem to stand to that of two-person games in the relation of sea-sickness to a headache.

(Hamilton 1975, p. 151)

1. INTRODUCTION

In economics, the public goods game (Binmore 1994; Kagel & Roth 1995) and the ultimatum game (Güth *et al.* 1982; Bolton & Zwick 1995) are well established as paradigms for discussing altruistic cooperative behaviour both in theory as well as in experimental settings. In biology and psychology, another closely related game, the Prisoner's Dilemma (Axelrod & Hamilton 1981), has attracted most attention in this context. Recent papers reflect a convergence of the two traditions (e.g. Nowak *et al.* 2000; Fehr & Gächter 2002; Hauert *et al.* 2002; Milinski *et al.* 2002). Obviously, a fruitful interdisciplinary cooperation is emerging.

On the one hand, in its simplest form, the public goods game corresponds to an extension of the pairwise interactions in the Prisoner's Dilemma to an arbitrary number of players. On the other hand, an ultimatum game reduces in its strategic essentials to a Prisoner's Dilemma game followed by a punishment round (Sigmund *et al.* 2001). Humans do not need to be told about the ubiquity of punishment in social interactions. In non-human biology, punishment is also frequently observed (Clutton-Brock & Parker 1995).

A large number of experiments document that most humans display a pronounced readiness to cooperate and a strong preference for fair strategies (e.g. Wedekind &

Milinski 1996; Fehr & Gächter 2000; Henrich *et al.* 2001). These findings are clearly at odds with predictions from both classical and evolutionary game theory stating that asocial behaviour should invariably become established. In economics this represents a major challenge to the selfish behaviour prescribed to *homo oeconomicus*.

To account for the interaction networks in human and animal societies, we consider spatially extended systems where players do not interact randomly, as in a well-mixed population, but interact with their neighbours only. For pairwise interactions it is well known that the inclusion of a neighbourhood structure has decisive effects on the evolution of cooperation (e.g. Nowak & May 1992; Killingback *et al.* 1999; Hauert 2001). The spatial arrangement of players enables cooperators to form clusters and thereby to minimize exploitation by asocial players. In this paper, we consider public goods games played in groups of three players arranged on a hexagonal lattice. We show that spatial extension again has significant effects on the equilibrium frequencies of cooperators and defectors. Moreover, by including punishment and by introducing reputation, we obtain scenarios where defectors hardly stand a chance and highly cooperative outcomes emerge that put a strong emphasis on fair strategies.

The step from two- to three-player interactions may seem a small one, but, in certain cases, it can lead to crucial differences. This is so, in particular, with the iterated public goods game. In two-player interactions (i.e. the familiar Prisoner's Dilemma), cooperators can punish defectors by withholding their contribution in the next round (in a tit-for-tat kind of move). In groups of three or more, this is no longer an efficient move, because withholding one's contribution hurts cooperators and defectors alike: tit-for-tat does not discriminate between the good and the bad. Economic experiments show very clearly that this is not just a theoretical issue but highlights

*Author for correspondence (hauert@zoology.ubc.ca).

an essential difference between Prisoner's Dilemma experiments involving two and three or more players.

In long-lasting interactions and noisy environments leading to erroneous decisions, cooperative outcomes may still be achieved even for more than two players (Hauert & Schuster 1997). However, it appears that synchronization and error-correction capabilities are the essential mechanisms rather than retaliation and reciprocity. It is generally accepted (and formally proven by Boyd & Richerson (1988), albeit for a restricted set of strategies) that reciprocating strategies are highly unlikely to evolve in interactions within larger groups if players cannot discriminate and channel their retaliation to specific wrong-doers. This problem occurs already in three-player interactions. An early analysis of a 'three-person Prisoner's Dilemma', together with a variant where two cooperators can gang together to punish a single defector, has been published by Hamilton (1975).

2. SPATIAL PUBLIC GOODS GAMES

In a typical public goods game, an experimenter asks n players to invest some money into a common pool. All players know that the total amount is multiplied by a factor r , with $1 < r < n$, and divided equally among the n players, irrespective of their contribution. If all players cooperate and contribute fully, they increase their initial capital by $(r-1)c$, where c denotes the cost of cooperation, i.e. the invested money. However, players face the temptation to defect by withholding their share and exploiting the other players' contributions. Obviously, such selfish behaviour yields a higher pay-off, irrespective of the co-players' actions, because the investment of each player returns only a fraction $r/n < 1$ to the investor. Consequently, a group of selfish players will not gain anything. In spite of this argument, in human experiments, the readiness to cooperate is surprisingly high. However, if the game is repeated, the contributions usually decrease from round to round. This decline soon approaches (but never exactly reaches) the economic stalemate predicted by theory (Fehr & Gächter 2000).

In most such experiments, stylized as they are and restricted to the bare essentials, there remains a wealth of possible strategic behaviour, e.g. concerning investment levels or conditional actions. To achieve a better theoretical understanding, we allow for binary choices only: all players must simultaneously decide whether to cooperate by contributing a fixed amount c or to defect by investing nothing.

In real life, individuals rarely interact with randomly chosen members of the population at large: most of their co-players reside in their immediate neighbourhood. Such situations can be approximated by confining each player to a lattice site and restricting interactions to nearest neighbours only. We consider the case of $n=3$ where the players are arranged on a hexagonal lattice. The lattice is updated over time in discrete steps, i.e. in a synchronized fashion, which models populations with non-overlapping generations. Each generation experiences two stages: an interaction stage followed by an imitation/reproduction stage. In the interaction stage every player participates in six games with two neighbours each. When numbering the neighbours clockwise, the first game takes place with

neighbours 1 and 2, the second with 2 and 3, etc. until the last game is with neighbours 6 and 1. Players use the same strategy in all six games.

The score of a player corresponds to the accumulated pay-off over all six encounters. This score determines the player's success in the subsequent imitation/reproduction stage, either through replication and displacement or through neighbours imitating and adopting a more successful strategy. Note that the two interpretations reduce to the same dynamics. In the imitation picture, players compare their score with the scores of each of their neighbours and adopt a neighbour's strategy with a probability proportional to the difference in scores, provided that the neighbour's score is greater and with probability zero otherwise. This update rule represents a spatial analogue of the replicator dynamics in well-mixed populations (Hofbauer & Sigmund 1998).

The equilibrium frequencies of cooperators and defectors are depicted in figure 1 as a function of the multiplication factor r that determines the value of the public goods. For r below a threshold $r_c \approx 2.07$ cooperators are doomed and defectors dominate. Conversely, for higher $r \geq 2.48$ cooperators manage to displace and eliminate all defectors. Intermediate factors $r_c \leq r \leq 2.48$ result in dynamic equilibria where cooperators and defectors coexist. Cooperators typically form clusters whose size increases with r . By contrast, in well-mixed populations, defectors would invariably dominate and achieve fixation, regardless of r . These results are consistent with findings for pairwise interactions in spatially extended systems.

3. PUNISHMENT

Let us now introduce the possibility of imposing fines onto other players. After every public goods interaction, the players have the opportunity to punish defecting co-players. However, punishment is costly, i.e. in experiments, the fees for punishing as well as the fines must be paid to the experimenter. Punishment is therefore an unselfish behaviour. Nevertheless it turns out to be an efficient way to increase contributions to the public good, i.e. to promote and stabilize cooperative behaviour. In experiments, humans demonstrate a pronounced and high readiness to punish defectors. This readiness increases with increasing temptation to defect (Yamagishi 1988) and remains efficient even in the absence of future interactions (Fehr & Gächter 2002).

Obviously, punishment strategies may be very complex and depend for instance on the number of defectors (as in Hamilton's three-person Prisoner's Dilemma), but again, for the sake of simplicity, we consider binary options only: after every public goods game each player decides whether or not to punish all non-cooperating co-players. To the punisher, this involves the cost γ per defecting co-player. For the defectors, this means an imposed fine β per punishing co-player. This results in four possible strategies:

- (i) G_1 , cooperate, and punish defectors—this reflects 'social behaviour';
- (ii) G_2 , defect, and punish defectors—a rather paradoxical strategy, which performs badly when playing against its like;

Figure 1. (a) Equilibrium frequencies of cooperators (solid line) and defectors (dotted line) as functions of the multiplication factor r with $c = 1$. For sufficiently large $r > r_c \approx 2.07$, cooperators thrive by forming clusters and coexist in dynamic equilibrium with defectors. The average cluster size increases with r until $r \approx 2.49$, where defectors vanish. Initially the strategies are randomly distributed with equal frequencies on a 600×600 lattice. The system is then relaxed over 3000 generations before a measurement period of 2000 generations. (b) The snapshot shows a typical configuration of cooperators (white) and defectors (black) in the region of coexistence ($r = 2.2$). Grey shades indicate players that have recently changed their strategy.

- (iii) G_3 , defect, but do not punish—this is the asocial ‘selfish’ strategy, maximizing the short-term personal profits; and
- (iv) G_4 , cooperate, but do not punish—this is a ‘mild’ strategy, which is free-riding because it avoids the costs of punishing defectors.

The two punishing strategies, G_1 and G_2 , condition their punishment activities on the strategic behaviour of their fellow players in the preceding public goods game. The non-punishing strategies, G_3 and G_4 , correspond to those studied in § 2.

According to classic game theory, players should adopt the asocial strategy G_3 irrespective of the value of the public good, i.e. the multiplication factor r . Evolutionary game theory in well-mixed populations yields this result for the long-term behaviour (Sigmund *et al.* 2001; Hauert *et al.* 2003).

In contrast to public goods interactions without punishment (compare figures 1 and 2), coexistence of cooperators and defectors is no longer possible. The punishment opportunity renders the system bistable, i.e. the initial frequencies and distributions of the strategies determine whether the system converges to a homogeneous state of asocial G_3 or a cooperative mixture of G_1 and G_4 (see figure 2). Note that any configuration of those two strategies is frozen because in the absence of defectors they are indistinguishable. The paradoxical strategy G_2 consistently vanishes within the first few generations.

In addition to these substantial changes in the qualitative dynamics, the threshold value r_c is considerably reduced to $r_c \approx 1.35$. In theory, for $c = \gamma = 1$, $\beta = 1.5$ and $r \geq 1.2$ straight boundaries of the social strategy G_1 are impenetrable by G_3 . But, at the same time, clusters of G_1 are readily invaded at corners and along rugged boundaries. Only above the threshold r_c are the two processes equally efficient.

For r slightly above r_c , the bistability of the system is reflected in the fact that only sufficiently large clusters of G_1 are capable of growing and eventually displacing G_3 . Thus, in the vicinity of r_c , figure 2 essentially depicts the

probability of (i) finding a suitable cluster after initialization and (ii) evolving towards a homogeneous G_1 or G_3 state, respectively. Note that the mild G_4 consistently vanishes because it requires the concerted punishment activities of G_1 to defeat G_3 at such a low r . Obviously, the presence of suitable G_1 clusters strongly correlates with the system size and with the initial frequencies and distributions of the strategies. In the limiting case of an infinite system size, the transition is expected to become sharp and independent of the initialization.

The proportions of G_1 and G_4 depend on the initial conditions and on r . For lower r , the mild strategy G_4 usually vanishes or remains at very low frequencies. But as r increases, the fraction of G_4 reaches almost 80%. This has two causes. First, it is costly to punish and defeat G_3 , which leads to a disadvantage of G_1 compared with G_4 . Second, for larger r the clustering advantage facilitates cooperation even in the absence of punishment (cf. figure 1).

4. REPUTATION

The previous scenarios assumed players operating under full anonymity. However, in more realistic scenarios relating to higher organisms and in particular to humans, players may accumulate information about their environment and specifically about potential future interaction partners. Similar to the conditioning of the punishment activity, each player may then condition his cooperative effort on the punishing behaviour of his fellows in other interactions. In particular, a cooperator who knows he is matched with two non-punishers could be tempted to take advantage of the situation by temporarily switching to defection without having to fear punishment. In that sense, all players carry some sort of reputation reflecting their strategic character. Through observations of third-party interactions and gossip, a player's reputation may become known to others. Therefore, we assume that, with a probability μ , a cooperator learns about the punishing behaviour of its co-players and at the same time succumbs to the temptation when faced with two non-punishers. In

Figure 2. (a) Frequencies of the four strategies social G_1 (solid line), paradoxical G_2 (vanished), asocial G_3 (dotted line) and mild G_4 (dashed line) as functions of the multiplication factor r with $c = 1$ and the punishment parameters $\gamma = 1$, $\beta = 1.5$. Punishment opportunities result in bistable dynamics. Above the significantly lower threshold $r_c \approx 1.35$ (cf. figure 1) coexistence of cooperators and defectors is no longer possible. In the long run either the asocial G_3 dominates or a frozen cooperative mixture of G_1 and G_4 is established with a sensitive dependence on the initial configuration. For higher r mild players abound, free-riding on the efforts of G_1 to ban asocial behaviour. The paradoxical strategy always vanishes within a few generations. Note that the discontinuities in the frequencies of G_1 and G_4 near $r \approx 1.8$ and $r \approx 2.25$ arise (i) from changes in the fate of certain local strategy configurations and (ii) from the synchronized lattice update. The 200×200 lattice was initialized with equal frequencies of all four strategies and then allowed to evolve until either the homogeneous G_3 or the frozen G_1 and G_4 state was reached. To eliminate the effects of particular initial configurations, the process was repeated and averaged over 100 realizations. Near r_c and the discontinuities the system size was increased to 400×400 . (b) The snapshot depicts a typical frozen mixture of G_1 (white) and G_4 (grey) for $r = 2.2$.

well-mixed populations with random encounters, reputation can promote and stabilize the social strategy G_1 (Sigmund *et al.* 2001; Hauert *et al.* 2003). A complementary case occurs if, with a probability ν , defectors who learn that they are up against punishers are sufficiently intimidated and cooperate. We shall not consider this effect here, because it turns out to be less important.

For $\mu > 0$, interactions between G_1 and G_4 are no longer neutral. Indeed, G_4 performs worse because any G_1 or G_4 player matched with two G_4 players will occasionally defect and this lowers the overall score of G_4 players.

Reputation preserves the bistability introduced by punishment and further increases the range of r feasible for cooperation by slightly lowering the threshold to $r_c \approx 1.25$ (see figure 3). As before, the paradoxical G_2 strategy quickly vanishes and, for r in the vicinity of r_c , the time evolution sensitively depends on the initial configuration, i.e. on the presence of a sufficiently large G_1 cluster. Actually, the value of r_c is essentially determined by the performance of G_1 against G_3 . Reputation strengthens the position of G_1 because these players now occasionally refrain from cooperation when matched with two G_3 s. In contrast to these minor changes near r_c , significant changes are observed for higher r . Reputation clearly promotes the social strategy G_1 and reduces the mild players to a small minority, so that invading defectors are reliably punished and quickly eliminated.

5. DISCUSSION

In spatial settings with three interacting players, cooperators may thrive simply by forming clusters and thereby reducing interactions with defectors. This is in agreement with related work on pairwise interactions in the spatial Prisoner's Dilemma (e.g. Nowak & May 1992; Killingback *et al.* 1999; Hauert 2001). However, persistent

cooperative behaviour requires relatively high multiplication factors r , i.e. above the threshold $r_c \approx 2.07$. For $r > r_c$ cooperators and defectors coexist in a dynamic equilibrium. The average cluster size of cooperators increases with r until they eventually displace all defectors.

Introducing punishment alters the dynamics in a fundamental way. The system becomes bistable and coexistence of cooperators and defectors is no longer possible. At the same time, punishment opportunities render cooperation favourable for a wider parameter range by significantly reducing the threshold to $r_c \approx 1.35$. In any finite system near r_c the initial configuration, i.e. the presence of a sufficiently large cluster of social players, determines whether the system evolves into a cooperative or a defecting homogeneous state. Cooperative states are frozen mixtures of the social and mild strategies (in the absence of defectors the two strategies perform equally well). For higher r the mild strategy, which avoids punishing (as well as being punished), spreads and reaches above 80%. This leaves the costly task of defeating asocial behaviour to fewer and fewer social individuals.

Intriguingly, when increasing the fines β , the threshold r_c is lowered until eventually $r_c < 1$. Thus, cooperative outcomes occur even for $r < 1$, but note that the net benefit for mutual cooperation is $(r - 1)c$, i.e. it becomes negative. Consequently, punishment can force a population into a cooperative state that leaves everyone worse off than in an asocial population—resulting in another dilemma.

Finally, overwhelmingly social and fair outcomes are achieved when combining punishment and reputation, i.e. if players may learn about the strategical character of their fellow players and adjust their behaviour accordingly. Reputation alters the dynamics in a subtler way. The system remains bistable and the threshold is only slightly lowered to $r_c \approx 1.25$. But, more importantly, reputation resolves the frozen mixtures of social and mild players.

Figure 3. (a) Frequencies of G_1 (solid line), G_2 (vanished), G_3 (dotted line) and G_4 (dashed line) as functions of the multiplication factor r with $c = 1$, $\gamma = 1$, $\beta = 1.5$ and a probability $\mu = 0.1$ of cooperators defecting if they can get away with it. Knowing and exploiting the co-players' propensity not to punish further lowers the threshold $r_c \approx 1.25$ (cf. figure 2) but, more importantly, results in overwhelmingly fair outcomes for $r > r_c$. The social strategy G_1 dominates, allowing for only a small minority of non-punishing cooperators G_4 , i.e. solely isolated G_4 players may persist. Those G_4 do equally well because G_1 keeps cooperating because of the other G_1 co-player. Simulation parameters as in figure 2. (b) The snapshot indicates the dominance of the fair G_1 strategy (white) interspersed with G_4 (grey) for $r = 2.2$.

The social strategy now thrives by occasionally defecting when faced with two mild non-punishers. Consequently, only single and isolated mild players may persist.

Note that in well-mixed populations the introduction of punishment alone does not change much. In the long run the asocial strategy inevitably spreads and outcompetes all other strategies. Only with the additional introduction of reputation does the system become bistable and social outcomes attainable (Sigmund *et al.* 2001). By contrast, in the spatial extended scenario, punishment alone leads to bistable dynamics. The long-term evolution of any finite system near the threshold value r_c sensitively depends on the initial configuration of the system. Adding reputation does not lead to similar fundamental changes, i.e. it leaves the bistability untouched and slightly lowers r_c . However, only reputation allows for overwhelmingly social and fair outcomes by defeating the mild strategy.

For an interactive tutorial we refer the reader to Hauert (2003). This allows visitors to vary conditions and parameters, and to verify that the essence of the findings presented here remains valid even if the scenarios are altered, for instance for asynchronous lattice updates or varying group sizes on different lattice structures, etc.

To conclude, we have studied three effects, each boosting cooperation in triplets of interacting individuals: (i) localized interactions; (ii) punishment directed against exploiters; and (iii) reputation unmasking non-punishers. All three effects tend to reduce the anonymity. In a crowd of faceless individuals randomly milling around, cooperation in sizeable groups is indeed hard to achieve. But if (i) players know the addresses of their co-players, (ii) they can trace defection to the perpetrators and (iii) they are aware of individual reputations, cooperation is much more easily achieved.

Paradoxically, the highly beneficial effect of reputation, which almost guarantees cooperation and fairness, is obtained by allowing cooperators to defect if they believe that they can get away with it. Thus, higher levels of fairness and cooperation are achieved by a lowering of the morals.

H.B. and K.S. acknowledge the support of the Wissenschaftskolleg WK W008 Differential Equation Models in Science and Engineering of the Austrian FWF; C.H. acknowledges support from the Swiss National Science Foundation 8220-64682.

REFERENCES

- Axelrod, R. & Hamilton, W.D. 1981 The evolution of cooperation. *Science* **211**, 1390–1396.
- Binmore, K. G. 1994 *Playing fair: game theory and the social contract*. Cambridge, MA: MIT Press.
- Bolton, G. E. & Zwick, R. 1995 Anonymity versus punishment in ultimatum bargaining. *Games Econ. Behav.* **10**, 95–121.
- Boyd, R. & Richerson, P. J. 1988 The evolution of reciprocity in sizeable groups. *J. Theor. Biol.* **132**, 337–356.
- Clutton-Brock, T. H. & Parker, G. A. 1995 Punishment in animal societies. *Nature* **373**, 209–216.
- Fehr, E. & Gächter, S. 2000 Cooperation and punishment in public goods experiments. *Am. Econ. Rev.* **90**, 980–994.
- Fehr, E. & Gächter, S. 2002 Altruistic punishment in humans. *Nature* **415**, 137–140.
- Güth, W., Schmittberger, R. & Schwarze, B. 1982 An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Org.* **3**, 367–388.
- Hamilton, W.D. 1975 Innate social aptitudes of men: an approach from evolutionary genetics. In *Biosocial anthropology* (ed. R. Fox), pp. 133–155. London: Malaby Press.
- Hauert, C. 2001 Fundamental clusters in spatial 2×2 games. *Proc. R. Soc. Lond. B* **268**, 761–769. (DOI 10.1098/rspb.2000.1424.)
- Hauert, C. 2003 See <http://www.univie.ac.at/virtuallabs/RewardAndPunishment>
- Hauert, C. & Schuster, H. G. 1997 Effects of increasing the number of players and memory size in the iterated Prisoner's Dilemma: a numerical approach. *Proc. R. Soc. Lond. B* **264**, 513–519. (DOI 10.1098/rspb.1997.0073.)
- Hauert, C., De Monte, S., Hofbauer, J. & Sigmund, K. 2002 Volunteering as red queen mechanism for cooperation in public goods games. *Science* **296**, 1129–1132.

- Hauert, C., Haiden, N. & Sigmund, K. 2003 The dynamics of public goods. *Discrete Cont. Dynamical Syst. B.* (In the press.)
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H. & McElreath, R. 2001 Cooperation, reciprocity and punishment in fifteen small scale societies. *Am. Econ. Rev.* **91**, 73–78.
- Hofbauer, J. & Sigmund, K. 1998 *Evolutionary games and population dynamics*. Cambridge University Press.
- Kagel, J. H. & Roth, A. E. (eds) 1995 *The handbook of experimental economics*. Princeton University Press.
- Killingback, T., Doebeli, M. & Knowlton, N. 1999 Variable investment, the continuous Prisoner's Dilemma, and the origin of cooperation. *Proc. R. Soc. Lond. B* **266**, 1723–1728. (DOI 10.1098/rspb.1999.0838.)
- Milinski, M., Semmann, D. & Krambeck, H.-J. 2002 Reputation helps solve the 'tragedy of the commons'. *Nature* **415**, 424–426.
- Nowak, M. A. & May, R. M. 1992 Evolutionary games and spatial chaos. *Nature* **359**, 826–829.
- Nowak, M. A., Page, K. M. & Sigmund, K. 2000 Fairness versus reason in the ultimatum game. *Science* **289**, 1773–1775.
- Sigmund, K., Hauert, C. & Nowak, M. A. 2001 Reward and punishment. *Proc. Natl Acad. Sci. USA* **98**, 10 757–10 762.
- Wedekind, C. & Milinski, M. 1996 Human cooperation in the simultaneous and the alternating Prisoner's Dilemma: Pavlov versus generous tit-for-tat. *Proc. Natl Acad. Sci. USA* **93**, 2686–2689.
- Yamagishi, T. 1988 Seriousness of social dilemmas and the provision of a sanctioning system. *Social Psychol. Q.* **51**, 32–42.