

LETTERS

Winners don't punish

Anna Dreber^{1,6*}, David G. Rand^{1,2*}, Drew Fudenberg³ & Martin A. Nowak^{1,4,5}

A key aspect of human behaviour is cooperation^{1–7}. We tend to help others even if costs are involved. We are more likely to help when the costs are small and the benefits for the other person significant. Cooperation leads to a tension between what is best for the individual and what is best for the group. A group does better if everyone cooperates, but each individual is tempted to defect. Recently there has been much interest in exploring the effect of costly punishment on human cooperation^{8–23}. Costly punishment means paying a cost for another individual to incur a cost. It has been suggested that costly punishment promotes cooperation even in non-repeated games and without any possibility of reputation effects¹⁰. But most of our interactions are repeated and reputation is always at stake. Thus, if costly punishment is important in promoting cooperation, it must do so in a repeated setting. We have performed experiments in which, in each round of a repeated game, people choose between cooperation, defection and costly punishment. In control experiments, people could only cooperate or defect. Here we show that the option of costly punishment increases the amount of cooperation but not the average payoff of the group. Furthermore, there is a strong negative correlation between total payoff and use of costly punishment. Those people who gain the highest total payoff tend not to use costly punishment: winners don't punish. This suggests that costly punishment behaviour is maladaptive in cooperation games and might have evolved for other reasons.

The essence of cooperation is described by the Prisoner's Dilemma. Two players have a choice between cooperation, C, and defection, D. If both players cooperate they get more than if both defect, but defecting against a cooperator leads to the highest payoff, while cooperating with a defector leads to the lowest payoff. One way to construct a Prisoner's Dilemma is by assuming that cooperation implies paying a cost for the other person to receive a benefit, whereas defection implies taking something away from the other person (Fig. 1).

Without any mechanism for the evolution of cooperation, natural selection favours defection. However, several such mechanisms have been proposed, including direct and indirect reciprocity⁷. Direct reciprocity means there are repeated encounters between the same two individuals, and my behaviour depends on what you have done to me^{1–6}. Indirect reciprocity means there are repeated encounters within a group; my behaviour also depends on what you have done to others.

Costly (or altruistic) punishment, P, means that one person pays a cost for another person to incur a cost. People are willing to use costly punishment against others who have defected^{8–18}. Costly punishment is not a mechanism for the evolution of cooperation⁷ but requires a mechanism for its evolution^{19–23}. Like the idea of reputation effects²⁴, costly punishment is a form of direct or indirect reciprocity: if I punish you because you have defected against me, direct reciprocity is used; if I punish you because you have defected with others,

indirect reciprocity is at work. The concept of costly punishment suggests that the basic game should be extended from two possible behaviours (C and D) to three (C, D and P). Here we investigate the consequences of this extension for the repeated Prisoner's Dilemma.

A total of 104 subjects participated in repeated Prisoner's Dilemma experiments at the Harvard Business School Computer Lab for Experimental Research. Participants interacted anonymously in pairwise encounters by means of computer screens. Subjects did not know how long each interaction would last, but knew that the probability of another round was 0.75 (as in ref. 25). In any given round, the subjects chose simultaneously between all available options, which were presented in a neutral language. After each round, the subjects were shown the other person's choice as well as both payoff scores. At the end of the interaction, the participants were presented with the final scores and then randomly rematched for another interaction.

We performed two control experiments (C1 and C2) and two treatments (T1 and T2). In the control experiments, people played a standard repeated Prisoner's Dilemma. In each round they could either cooperate or defect. Cooperation meant paying 1 unit for the other person to receive 2 units (in C1 and T1) or 3 units (in C2 and T2). Defection meant gaining 1 unit at a cost of 1 for the other person. In the treatments, people had three options in every round: cooperate, defect or punish. Punishment meant paying 1 unit for the other person to lose 4. We used a 4:1 punishment technology because it has been shown to be more effective in promoting cooperation than

a	You get Other gets			
	C	-c +b		
	D	+d -d		
Your move	P	-α -β		
c		C D P		
	C	1 -2 -5		
	D	2 0 -3		
Your move	P	1 -2 -5		
b		C D P		
	C	b-c -d-c -β-c		
	D	b+d 0 -β+d		
Your move	P	b-α -d-α -β-α		
d		C D P		
	C	2 -2 -5		
	D	4 0 -3		
Your move	P	2 -2 -5		

Figure 1 | Payoff values. **a**, The game is formulated in terms of unilateral moves. There is the choice between cooperation (C), defection (D) and costly punishment (P). Cooperation means paying a cost c for the other person to get a benefit b . Defection means earning a payoff d at a cost d to the other person. Punishment means paying a cost α for the other person to incur a cost β . **b**, The payoff matrix is constructed from these unilateral moves. **c, d**, The actual payoff matrices of our experiments: C1 and T1 (**c**); C2 and T2 (**d**).

¹Program for Evolutionary Dynamics, ²Department of Systems Biology, ³Department of Economics, ⁴Department of Mathematics, and ⁵Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, Massachusetts 02138, USA. ⁶Department of Economics, Stockholm School of Economics, 11358 Stockholm, Sweden.

*These authors contributed equally to this work.

other ratios¹³. The resulting payoff matrices are shown in Fig. 1. See Supplementary Information for more details.

Figure 2 shows some examples of games that occurred in the treatments T1 and T2. A number of games were all-out cooperation. Sometimes cooperation could be maintained by forgiving an opponent's defection. At other times, defection in response to defection was able to restore cooperation. Typically, costly punishment did not re-establish cooperation. In some cases, costly punishment provoked counter-punishment, thereby assuring mutual destruction. Giving people the option of costly punishment can also lead to unprovoked first strikes, with disastrous consequences.

Comparing the two control experiments, C1 and C2, we find that the frequency of cooperation increases as the benefit-to-cost ratio increases. In C1, 21.2% of decisions are cooperation, in contrast with 43.0% in C2. For both parameter choices, cooperation is a sub-game perfect equilibrium. Comparing each control experiment with its corresponding treatment, we find that punishment increases the frequency of cooperation. In T1 and T2, 52.4% and 59.7% of all decisions are cooperation.

Punishment, however, does not increase the average payoff. In T1 and T2, we observe that 7.6% and 5.8% of decisions are punishment, P. We find no significant difference in the average payoff when comparing C1 with T1 and C2 with T2. Punishment therefore has no benefit for the group, which makes it hard to argue that punishment might have evolved by group selection²².

Examining the data of experiments T1 and T2 at the individual level, we find no correlation between the use of cooperation or defection and payoff, but a strong negative correlation between the use of punishment and payoff (Fig. 3). In experiment T1, the five top-ranked players, who earned the highest total payoff, have never used costly punishment. In both experiments, the players who end up with the lowest payoff tend to punish most often. Hence, for maximizing

the overall income it is best never to punish: winners don't punish (Fig. 3).

It might be that the winners of our experiment were merely lucky in that they were paired with people against whom punishment was not necessary. To test this hypothesis, we analysed the correlation between payoff and the first-order conditional strategies used by people. Figure 4 illustrates a strong negative correlation between payoff and the probability to use punishment, P, after the opponent has defected, D. Winners tend to respond by using D against D, whereas losers use P against D. The response to another person's defection is the only strategic feature that is clearly correlated with winning or losing the game. Winners play a tit-for-tat-like strategy^{2,4}, whereas losers use costly punishment.

It could be that using costly punishment becomes more beneficial as the game progresses. To test this possibility, we separately analysed the data from the last one-quarter of all interactions. Again, it remains true that there is a strong negative correlation between an individual's payoff and that individual's use of costly punishment.

In previous experiments, punishment was usually offered as a separate option after one or several rounds of a public goods game. The public goods game is a multi-person Prisoner's Dilemma, in which each player can invest a certain sum into a common pool,

Decisions	Payoff in this interaction	Final rank
a Nice people finish first		
C C C C	8	1
C C C C	8	2
b Punish and perish		
C P P P	-10	25
D D D D	-9	22
c Defection restores cooperation		
C D D C D C	10	15
D C C C C C	4	9
d Turning the other cheek		
C C C C C C	2	6
D D C C C C	14	19
e Mutually assured destruction		
C P P P D D	-20	30
D D P P P P	-14	25
f Revenge is not so sweet		
C C C P D D P P	-6	24
C D D D D D D D	-4	22
g A 'pre-emptive strike'		
C P D	2	29
C C D	-4	24

Figure 2 | Games people played. There were 1,230 pairwise repeated interactions, each lasting between one and nine rounds. Some examples are given (**b**, **e** and **g** are from T1; the others are from T2). The two players' moves, the cumulative payoff of that interaction and the final rank of each player (sorted from highest to lowest payoff) are shown. **a**, All-out cooperation between two top-ranked players. **b**, Punish and perish. **c**, Defection for defection can sometimes restore cooperation. **d**, Turning the other cheek can also restore cooperation. **e**, Mutual punishment is mutual destruction. **f**, Punishment does not restore cooperation. Player 1 punishes a defection, which leads to mutual defection. Then player 1 is unsatisfied and deals out more punishment. **g**, "Guns don't kill people, people kill people." (Punishment itself is not destructive, only the people who use it.) Here, an unprovoked first strike destroys cooperation. The option to punish allows irrational people to inflict harm on the undeserving.

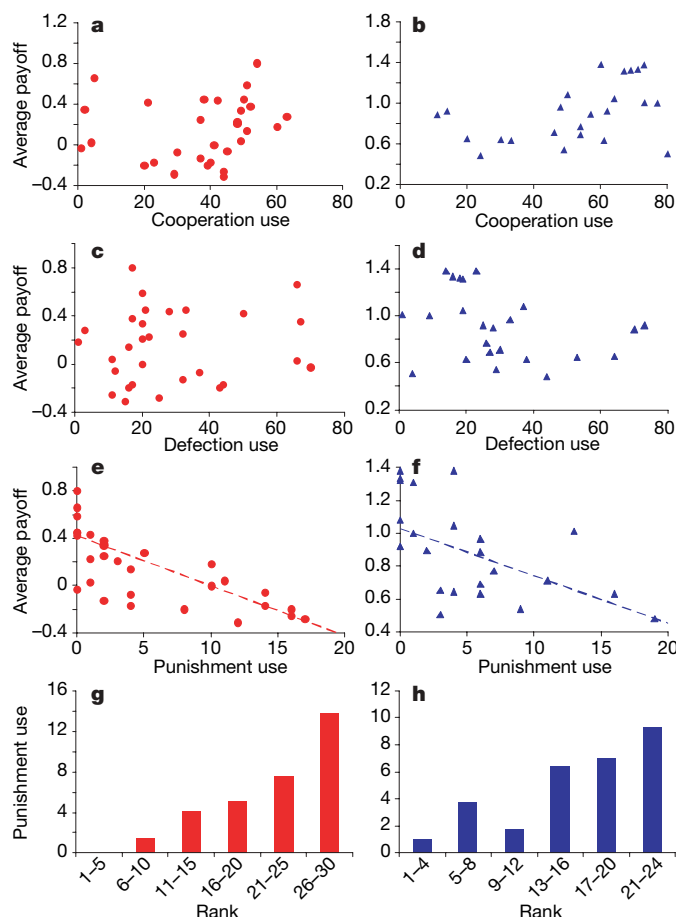


Figure 3 | Punish and perish. In both treatments T1 (red; $b/c = 2$) and T2 (blue; $b/c = 3$), there is no correlation between average payoff per round and use of cooperation (quantile regression; **a**, $P = 0.33$; **b**, $P = 0.21$) or between average payoff per round and use of defection (**c**, $P = 0.66$; **d**, $P = 0.36$). However, there is a significant negative correlation between average payoff per round and punishment use in both treatments (**e**, slope = -0.042 , $P < 0.001$; **f**, slope = -0.029 , $P = 0.015$). Punishment use is the overriding determinant of payoff. The x axis in **a-f** shows the total number of moves of the given type made over the course of the experiment. **g**, **h**, Ranking players according to their total payoff shows a clear trend: players with lower rank (higher payoffs) punish less than players with higher rank (lower payoff).




Figure 4 | Tit-for-tat prevails over costly punishment. Lower payoffs are correlated not only with punishment use, but also specifically with choosing to punish after the opponent has defected. The probability of punishing immediately after a co-player's defection is negatively correlated with the average payoff per round, both in T1 (**a**; $b/c = 2$) and in T2 (**b**; $b/c = 3$) (quantile regression; **a**, slope = -0.81 , $P < 0.001$; **b**, slope = -0.94 , $P = 0.015$). Thus, the lower payoffs of punishers were not caused by the bad luck of interacting with defectors. Winners use a tit-for-tat-like approach (D for D), whereas losers use costly punishment (P for D).

which is then multiplied by a factor and equally divided between all players irrespective of whether they have invested or not²⁶. After the public goods game, people are asked if they want to pay money for others to lose money. People are willing to use this option to punish those who have invested nothing or only very little, and the presence of this option has been found to increase contributions^{8,10}.

Careful analysis, however, has revealed that, in most cases, punishment does not increase the average payoff. In some experiments, punishment reduces the average payoff^{9,10,12,27}, whereas in others it does not lead to a significant change^{11,14,15}. Only once has punishment been found to increase the average payoff¹³. The higher frequency of cooperation is usually offset by the cost of punishment, which affects both the punisher and the punished. Our findings are in agreement with this observation: the option of costly punishment does not increase the average payoff of the group. It is possible, however, that in longer experiments and for particular parameter values punishment might increase the average payoff.

It is sometimes argued that costly punishment is a mechanism for stabilizing cooperation in anonymous, one-shot games. But whether or not this is the case seems to be of little importance, because most of our interactions occur in a context where repetition is possible and reputation matters. For millions of years of human evolution, our ancestors have lived in relatively small groups in which people knew each other. Interactions in such groups are certainly repeated and open ended. Thus, our strategic instincts have been evolving in situations where it is likely that others either directly observe my actions or eventually find out about them. In addition, in modern life most of our interactions occur with people whom we meet frequently. Typically, we can never rule out 'subsequent rounds'. Therefore, if costly punishment is important for the evolution of human cooperation, then it must have a beneficial role in the setting of repeated games. Our findings do not support this claim.

We also believe that our current design has some additional advantages over previous ones. In our setting, costly punishment is always one of three options. Hence, there is an opportunity cost for using punishment, because the subject forfeits the opportunity to cooperate or to defect. Our design also minimizes the experimenter and participant demand effects²⁸, because there are always several options²⁷. In many previous experiments retaliation for punishment is not possible^{9-16,27}, but it is a natural feature of our setting.

Thus, our data show that costly punishment strongly disfavours the individual who uses it and hence it is opposed by individual selection in cooperation games in which direct reciprocity is possible. We conclude that costly punishment might have evolved for reasons other than promoting cooperation, such as coercing individuals into submission and establishing dominance hierarchies^{20,29}. Punishment might enable a group to exert control over individual behaviour. A stronger individual could use punishment to dominate weaker ones.

People engage in conflicts and know that conflicts can carry costs. Costly punishment serves to escalate conflicts, not to moderate them. Costly punishment might force people to submit, but not to cooperate. It could be that costly punishment is beneficial in these other games, but the use of costly punishment in games of cooperation seems to be maladaptive. We have shown that in the framework of direct reciprocity, winners do not use costly punishment, whereas losers punish and perish.

METHODS SUMMARY

A total of 104 subjects (45 women, 59 men, mean age 22.2 years) from Boston-area colleges and universities participated voluntarily in a modified repeated Prisoner's Dilemma game at the Harvard Business School Computer Lab for Experimental Research (CLER). The lab consists of 36 computers, which are visually partitioned. The participants interacted anonymously through the software z-Tree³⁰ and were from a number of different schools and a wide range of fields of study; it was therefore unlikely that any subject would know more than one other person in the room. We asked subjects for their sex and major field of study. No significant difference in level of cooperation, punishment use or payoff was found between males and females, or between economics majors and non-economics majors (Mann-Whitney test, $P > 0.05$ for all sessions). Subjects were not allowed to participate in more than one session of the experiment. In all, four sessions were conducted in April and May 2007, with an average of 26 participants playing an average of 24 interactions, for an average of 79 total rounds per subject.

Each experiment was begun by reading instructions (included in the Supplementary Information), answering two test questions to verify understanding of the payoffs, and playing a practice interaction against another subject. At the start of each new interaction, subjects were unaware of the previous decisions of the other player. After each round, the subjects were shown the other person's choice as well as both payoff scores. At the end of the interaction, the participants were presented with the final scores and then randomly rematched for another interaction.

In each session, the subjects were paid a \$15 show-up fee. Each subject's final score summed over all interactions was multiplied by \$0.10 to determine additional earned income. Thus, one game unit corresponded to \$0.10. To allow for negative incomes while maintaining the \$15 show-up fee, \$5 was added to each subject's earned income at the end of the session. Subjects were informed of this extra \$5 at the beginning of the session. The average payment per subject was \$26 and the average session length was 1.25 h.

Received 13 November 2007; accepted 15 January 2008.

- Trivers, R. The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57 (1971).
- Axelrod, R. & Hamilton, W. D. The evolution of cooperation. *Science* **211**, 1390–1396 (1981).
- Fudenberg, D. & Maskin, E. Evolution and cooperation in noisy repeated games. *Am. Econ. Rev.* **80**, 274–279 (1990).
- Nowak, M. A. & Sigmund, K. Tit for tat in heterogeneous populations. *Nature* **355**, 250–253 (1992).
- Binmore, K. & Samuelson, L. Evolutionary stability in repeated games played by finite automata. *J. Econ. Theory* **57**, 278–305 (1992).
- Colman, A. M. *Game Theory and its Applications in the Social and Biological Sciences* (Routledge, New York, 1995).
- Nowak, M. A. Five rules for the evolution of cooperation. *Science* **314**, 1560–1563 (2006).
- Yamagishi, T. The provision of a sanctioning system as a public good. *J. Pers. Soc. Psychol.* **51**, 110–116 (1986).
- Ostrom, E., Walker, J. & Gardner, R. Covenants with and without a sword: self-governance is possible. *Am. Pol. Sci. Rev.* **86**, 404–417 (1992).
- Fehr, E. & Gächter, S. Altruistic punishment in humans. *Nature* **415**, 137–140 (2002).
- Botelho, A., Harrison, G. W., Pinto, L. M. C. & Rutström, E. E. Social norms and social choice (Working Paper no. 05-23, Department of Economics, College of Business Administration, University of Central Florida, 2005).
- Egas, M. & Riedl, A. The economics of altruistic punishment and the demise of cooperation (IZA Discussion Paper no. 1646, Institute for the Study of Labor (IZA), Bonn, 2005).
- Nikiforakis, N. & Normann, H. T. A comparative statics analysis of punishment in public-good experiments (Discussion Paper in Economics no. 05/07, Royal Holloway Univ. London, London, 2005).
- Page, T., Putterman, L. & Unel, B. Voluntary association in public goods experiments: reciprocity, mimicry and efficiency. *Econ. J.* **115**, 1032–1053 (2005).
- Bochet, O., Page, T. & Putterman, L. Communication and punishment in voluntary contribution experiments. *J. Econ. Behav. Org.* **60**, 11–26 (2006).
- Gürrer, Ö., Irlenbusch, B. & Rockenbach, B. The competitive advantage of sanctioning institutions. *Science* **312**, 108–111 (2006).

17. Rockenbach, B. & Milinski, M. The efficient interaction of indirect reciprocity and costly punishment. *Nature* **444**, 718–723 (2006).
18. Denant-Boemont, L., Masclet, D. & Noussair, C. N. Punishment, counterpunishment and sanction enforcement in a social dilemma experiment. *Econ. Theory* **33**, 145–167 (2007).
19. Boyd, R. & Richerson, P. J. Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethol. Sociobiol.* **13**, 171–195 (1992).
20. Clutton-Brock, T. H. & Parker, G. A. Punishment in animal societies. *Nature* **373**, 209–216 (1995).
21. Sigmund, K., Hauert, C. & Nowak, M. A. Reward and punishment. *Proc. Natl Acad. Sci. USA* **98**, 10757–10762 (2001).
22. Boyd, R., Gintis, H., Bowles, S. & Richerson, P. J. The evolution of altruistic punishment. *Proc. Natl Acad. Sci. USA* **100**, 3531–3535 (2003).
23. Fowler, J. H. Altruistic punishment and the origin of cooperation. *Proc. Natl Acad. Sci. USA* **102**, 7047–7049 (2005).
24. Fudenberg, D. in *Advances in Economic Theory: Sixth World Congress* (ed. Laffont, J. J.) 89–131 (Cambridge Univ. Press, Cambridge, 1993).
25. Dal Bó, P. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *Am. Econ. Rev.* **95**, 1591–1604 (2005).
26. Hardin, G. The tragedy of the commons. *Science* **162**, 1243–1248 (1968).
27. Sefton, M., Shupp, R. & Walker, J. M. The effect of rewards and sanctions in provision of public goods (Economics CAEPR Working Paper no. 2006–005, Indiana Univ. Bloomington, 2006).
28. Carpenter, J., Liat, A. & Vickery, B. They come to play: supply effects in an economic experiment (Economics Discussion Paper no. 06–02, Middlebury College, Middlebury, VT, 2006).
29. Samuelson, L. Economic theory and experimental economics. *J. Econ. Lit.* **43**, 65–107 (2005).
30. Fischbacher, U. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* **10**, 171–178 (2007).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements Support from the John Templeton Foundation, the National Science Foundation (NSF)–National Institutes of Health joint program in mathematical biology, the Jan Wallander Foundation (A.D.) and an NSF grant (D.F.) is gratefully acknowledged. The Program for Evolutionary Dynamics is sponsored by J. Epstein.

Author Information Reprints and permissions information is available at www.nature.com/reprints. Correspondence and requests for materials should be addressed to M.A.N. (martin_nowak@harvard.edu).