

Costly punishment does not always increase cooperation

Jia-Jia Wu^a, Bo-Yu Zhang^{a,b}, Zhen-Xing Zhou^b, Qiao-Qiao He^a, Xiu-Deng Zheng^a, Ross Cressman^{c,1}, and Yi Tao^{a,1}

^aKey Laboratory of Animal Ecology and Conservation Biology, Centre for Computational Biology and Evolution, Institute of Zoology, Chinese Academy of Sciences, Beijing 100101, China; ^bSchool of Mathematical Sciences, Beijing Normal University, Beijing 100875, China; and ^cDepartment of Mathematics, Wilfrid Laurier University, Waterloo, ON, Canada N2L 3C5

Edited by Brian Skyrms, University of California, Irvine, CA, and approved August 13, 2009 (received for review May 29, 2009)

In a pairwise interaction, an individual who uses costly punishment must pay a cost in order that the opponent incurs a cost. It has been argued that individuals will behave more cooperatively if they know that their opponent has the option of using costly punishment. We examined this hypothesis by conducting two repeated two-player Prisoner's Dilemma experiments, that differed in their payoffs associated to cooperation, with university students from Beijing as participants. In these experiments, the level of cooperation either stayed the same or actually decreased when compared with the control experiments in which costly punishment was not an option. We argue that this result is likely due to differences in cultural attitudes to cooperation and punishment based on similar experiments with university students from Boston that found cooperation did increase with costly punishment.

antisocial punishment | cultural effects | experimental outcome | Prisoner's Dilemma repeated game | reputation

Costly (altruistic) punishment means paying a cost for another individual to incur a cost (1) and has been advanced as a key mechanism to explain cooperative behavior in human societies (1, 2). A great deal of research, both through experimental studies and theoretical models (see also refs. 3–13), has examined whether the option of costly punishment promotes cooperative behavior in one-shot multiplayer public goods games where non cooperative free riders gain at no cost the same group benefits from the “good” as those individuals who contributed to it. For treatments that include costly punishment, players may punish other group members after they are informed what contribution each person made. Empirically, this option has consistently increased the level of cooperation as measured by average individual contribution, especially in typical experiments designed so subjects either do not interact more than once or do not know that they have. Of particular relevance for us is the additional experimental finding (1, 14) that different societal norms alter the prevalence of costly punishment and its effectiveness in promoting cooperative behavior.

To a lesser extent, costly punishment has also been examined in repeated games and in other games where players can base their behavior at later stages of an interaction on earlier actions of their opponents. For instance, players are now able to retaliate against an individual who previously punished them and reputation effects play a role (6). Experimental evidence based on public goods games with counterpunishment (12, 13) again show costly punishment promotes cooperative behavior. Instead, our experiments are based on the repeated two-player Prisoner's Dilemma (PD) game with and without the option of costly punishment. In this regard, Dreber et al. (15) recently performed repeated PD experiments where, in each round, participants chose between cooperation, defection and possibly costly punishment. The subjects participating in their experiments were college/university students in the Boston area. Along with the result that people who gain the highest payoff in this repeated game tend not to use costly punishment (i.e., winners don't

punish), they found that the prevalence of cooperation increases with the option of costly punishment.

To determine whether societal norms play a role in these results, we replicated their experiments with the same experimental design (see *Methods*) at the Computer Lab in the School of Life Sciences, Beijing Normal University with university students from Beijing as subjects. That is, following Dreber et al. (15), we performed two control experiments (C1 and C2) and two treatments (T1 and T2). In the control experiments, two subjects chosen at random play a standard repeated PD game without knowing how many rounds there will be in this two-player game but knowing that, after one round is complete, there is a 0.75 probability of another round. There are only two options in every round: cooperate and defect. Cooperation (C) means paying 1 unit for the other person to receive 2 units (in C1 and T1) or 3 units (in C2 and T2). Defection (D) means gaining 1 unit at a cost of 1 unit for the other person. In the treatments, people have three options in every round: cooperate, defect, or punish. Punishment (P) means paying 1 unit for the other person to lose 4 units. Subjects were told the payoff consequences of each strategy choice but not the name we associated with it to avoid any predisposition on their part to certain labels.

That is, the payoff matrices for the experiments C1 and T1 are given by

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} 1 & -2 \\ 3 & 0 \end{pmatrix} \end{array} \quad \text{and} \quad \begin{array}{ccc} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{pmatrix} 1 & -2 & -5 \\ 3 & 0 & -3 \\ 1 & -2 & -5 \end{pmatrix}, \end{array}$$

respectively, and the payoff matrices for the experiments C2 and T2 are given by

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} 2 & -2 \\ 4 & 0 \end{pmatrix} \end{array} \quad \text{and} \quad \begin{array}{ccc} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{pmatrix} 2 & -2 & -5 \\ 4 & 0 & -3 \\ 2 & -2 & -5 \end{pmatrix}, \end{array}$$

respectively. These payoff matrices are identical to those in Dreber et al. (15).

Results

Our first experimental results, depicted in Fig. 1A, show that costly punishment does not increase cooperation. Specifically, in the two control experiments, C1 and C2, the ratio of cooperation

Author contributions: R.C. and Y.T. designed research; J.-J.W., B.-Y.Z., Z.-X.Z., Q.-Q.H., and X.-D.Z. performed research; J.-J.W., B.-Y.Z., Z.-X.Z., and X.-D.Z. analyzed data; and R.C. and Y.T. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence may be addressed. E-mail: yitao@ioz.ac.cn or rcressman@wlu.ca.

This article contains supporting information online at www.pnas.org/cgi/content/full/0905918106/DCSupplemental.




Fig. 1. Average frequency of cooperation in each session. *A* is from our experiments, and *B* is from Dreber et al. (15) (see their supplementary figure 2A). Error bars represent the standard error from the mean.

to all decisions is 24.62% in C1, and 23.95% in C2. In T2, the frequency of cooperation is 26.45%, which is slightly higher than in C2 but there is no significant difference ($P = 0.2124$, Fisher's two-tailed exact test). Furthermore, in T1, the frequency of cooperation, 18.25%, is substantially smaller than in C1 ($P = 0.0544$). That is, costly punishment decreases the frequency of cooperation in T1. In summary, comparing each control experiment with its treatment, we find that the option of costly punishment either makes no significant difference or actually decreases the amount of cooperation in our experiments.

These results are unexpected because previous experiments on the repeated PD game [see Fig. 1*B*, which is reproduced with permission from Dreber et al. (15)] and on public goods games consistently show that costly punishment increases the frequency of cooperation. In particular, our result for T1 is completely opposite to that of previous experiments. Our study seems to be the first experiment that exhibits this latter outcome, although theoretical models that analyze evolutionary dynamics under direct reciprocity (16) also suggest costly punishment may not promote cooperation.

Although Dreber et al. (15) find empirically that there is significantly more cooperation in each treatment than in its corresponding control experiment, there is no significant difference in the average payoff (Fig. 2*B*) because any extra payoff due to more subjects mutually cooperating is offset by lower payoffs for both the punishers and those punished. That is, in their experiments, the option of costly punishment is neutral in regards to advantage to the group. Because cooperation does not increase in our treatment, we obtain the expected result that average payoff decreases (Fig. 2*A*). This decrease is more significant in C1 vs. T1 (Mann–Whitney test: $P = 5.0324 \times 10^{-6}$ and $z = -4.5634$) than in C2 vs. T2 ($P =$




Fig. 2. Average payoff per round in each session. *A* is from our experiments, and *B* is from Dreber et al. (15) (see their supplementary figure 2*B*). Also shown is the payoff per round for mutual cooperation (ALLC) in each session. Clearly, punishment does not provide any advantage for the group, and all control and treatment payoffs are lower than the ALLC payoff.

0.2092 and $z = 1.2558$). Thus, our experiments argue against group selection in cooperative games as a mechanism to explain the appearance of costly punishment in human societies (7). This reinforces the conclusion (15) that the use of costly punishment evolved in human societies for other reasons than to promote cooperation.

The most likely explanation for the different results in the two studies lies in cultural effects in China compared with the U.S. (see also the discussion at the end of *Methods* where we rule out that these differences are based on the level of monetary rewards). In particular, Chinese and Americans have different cultural attitudes toward reputation and authority. In China, a person's reputation is enhanced by establishing and maintaining a network of two-person dyadic relationships, known as “guanxi,” with more people who themselves have good reputations (17). Reputation through guanxi networks is clearly related to indirect reciprocity effects [i.e., when my behavior toward you depends on how you behave with others (11)]. Along with this emphasis on collectivism compared with individualism, Chinese society is hierarchical with deference given to authority figures (18). Because neither indirect reciprocity nor a predetermined authority figure are possible in the experimental setup of our repeated two-player PD game, these cultural characteristics do not provide a reason for our participants in Beijing to increase the level of cooperation in the presence of costly punishment. However, direct reciprocity (i.e., when my behavior toward you depends on how you have behaved to me) is relevant in our repeated games (16) and forms the basis of a player's reputation (in the sense of being able to predict how opponents will react to each other's actions) during the course of several

rounds between the same two players. This type of dyadic reputation [which is more important in Western society (17)], can be used to give theoretical predictions close to the outcome of the experiments in Boston (16). Heuristically, the levels of cooperation increase and defection decrease when individuals who have a reputation of defecting run the risk of being punished.

It also appears that societal differences are the main factor for varying attitudes to punishment in experiments based on the multiplayer public goods game (1). The variation found by Herrmann et al. (1) is most clear for antisocial punishment behavior (i.e., punishing the high contributors to the public good because these are suspected of being the primary punishers of the low contributors), which was rare for participants from democratic societies with advanced market economies and common in more traditional societies (19). Interestingly, in these experiments conducted in 16 comparable participant pools of undergraduate university students, the results on antisocial punishment for the two pools in the U.S. and China (Boston and Chengdu, respectively) are quite similar [figure 1 in Herrmann et al. (1)]. One explanation of this result is that indirect reciprocity occurs in these multiplayer games and so reputation (in either of the senses discussed in the previous paragraph) becomes a factor in people's behavior. That is, antisocial punishment is rare in the U.S. and China due to reputation even though the basis of this reputation differs in the two societies.

Our experimental results provide further evidence, based on the two-player repeated PD game, that participant's attitude to punishment depends on their cultural background. Specifically, Fig. 3A plots the histogram of rounds in which costly punishment is first used. Of particular interest is that the frequency of first P use during round 1 in our experiments is 24.47% in T1, and 29.82% in T2. These rates are much higher than in Dreber et al. (15), where this frequency is <5% in both T1 and T2 (Panel B). Participants who use costly punishment in round 1 are indiscriminately punishing before they know anything about their opponent's behavior. Furthermore, the combined use of P and D in the first round of our treatments is much higher than those of Dreber et al. (15) (>69% in both T1 and T2 compared with <28% in Boston). The outcome of experiments in the U.S. is consistent with participants using direct reciprocity to foster higher payoffs through mutual cooperation by initially exhibiting cooperation. They are willing to take this risk due to a belief it will be compensated for later through a better reputation. As argued above, guanxi reputation effects are not relevant for Chinese participants in the repeated PD game. Punishment in the first round could then be an attempt to establish oneself as a dominant authority figure who is willing to punish in later rounds if dissatisfied with how the interaction proceeds. Such a strategy (summarized by the Chinese phrase "xia ma wei," which means to deal someone a head-on blow at first encounter) may serve as a strong admonition (perhaps even intimidation) to an opponent. It is shown in the SI that there is no trend in the level of costly punishment or of defection used by our subjects in the first round over the course of either treatment session. Thus, the high frequency of first P use in round 1 as shown in Fig. 3A is a stable phenomenon that does not decrease as our subjects interact with more opponents.

These empirical discrepancies, on the initial use of costly punishment and its ongoing effect on the overall prevalence of cooperation and average payoff for the repeated PD game, between experiments with subjects in the U.S. and in China using the identical experimental design raise many questions on how culture impacts strategy choice. To examine this further, we compared other results from our experiments to those reported elsewhere. These are given in the SI where we show that our experiments match other aspects of the effects of costly punishment found by Dreber et al. (15). In particular, we find that




Fig. 3. Histogram of rounds in which costly punishment (P) is first used by a subject in T1 and in T2. Panel A is from our experiments where costly punishment was used 23 times in the first round of the 94 interactions where it was used at least once in T1 and 17 times in the first round of the 57 interactions where it was used at least once in T2. Panel B is from Dreber et al. (15) (see their supplementary figure 6).

winners don't punish and that average payoff also decreases with the increased use of punishment after defection (i.e., winners were not merely lucky by being paired with people against whom punishment was unnecessary). We also find that defection increases in later rounds of this repeated game as typically occurs in such experiments (15, 20).

Methods

A total of 94 subjects (58 women, 36 men, mean age 22.8 years old) from Beijing Normal University, Beijing Jiaotong University, Beijing University of Posts and Telecommunication, and Institute of Zoology, Chinese Academy of Sciences participated voluntarily in a modified repeated Prisoner's Dilemma game at the computer lab in the School of Life Sciences, Beijing Normal University. The lab consists of 38 computers. We developed a computer program for the anonymous interactions of participants, which is similar to the software zTree used by Dreber et al. (15) in their experiments. Subjects did not know how many rounds there would be in one pairwise interaction, only that the probability of another round was 0.75. In each round of an interaction, the two participants chose simultaneously between all available options and, after the round, were shown the other person's choice and both payoff scores. At the end of each interaction, participants were told both their total scores and then randomly rematched for another interaction. The experiments (four sessions) were conducted in October 2008, with an average of 24 participants playing an average of 22 interactions, for an average of 79 total rounds per subject.

Each session began by reading the written instructions to the subjects on how the repeated game is played. These instructions in Chinese are the translation of those used in Dreber et al. (15). Each participant answered two

