Game Experiments on Cooperation Through Punishment and/or Reward

Ross Cressman
Department of Mathematics
Wilfrid Laurier, Waterloo, Canada

Jia-Jia Wu, Boyu Zhang, Zhen-Xing Zhou, Qiao-Qiao He, Xiu-Deng Zheng, Ross Cressman and Yi Tao "Costly punishment does not always increase cooperation", PNAS, 106, 17448-17451, 2009.

Y. Tao, Cong Li., J-J. Wu and R. Cressman. "The efficiency ratio of costly punishment", PNAS, 106, E136, 2009.

J-J. Wu, Cong Li, B. Zhang, Hong-De Xiao, Qiao-Qiao He, R. Cressman and Y. Tao "Cooperation through institutional reward and punishment: group versus individual incentives", mimeo, 2010.

OUTLINE

Game Experiments on Cooperation

Repeated Prisoner's Dilemma (PD) Game

Cooperation with(out) costly punishment?
Experimental results in Beijing
Comparison to other experimental results
Reputation: dyadic versus network
Reputation and cultural differences
First use of costly punishment

Repeated Public Goods Game (PGG)

Cooperation without institutional incentives
Comparison to other experimental results
Institutional reward and/or punishment
Experimental results in Beijing
Comparisons to peer incentives
Effects on group incentives
Individual relative incentives
Evolution of cooperation

Costly Punishment in the Repeated PD Game

Costly punishment (sometimes also called altruistic punishment) means to pay a cost for another individual to incur a cost.

Experiments based on the one-shot multi-player public goods game consistently show that the option of costly punishment (i.e. the possibility that players may punish other group members after being informed what contributions each person made to the public good) promotes co-operative behaviour in this game.

In the two-player PD game, costly punishment is included as

$$\begin{array}{cc} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} you \quad get \\ other \quad gets \end{array} & \begin{bmatrix} -c & d & -\alpha \\ b & -d & -\beta \end{bmatrix} \end{array}$$

The one-shot PD game (with costly punishment) has payoff matrix:

$$\begin{array}{cc} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{bmatrix} b-c & -d-c & -\beta-c \\ b+d & 0 & -\beta+d \\ b-\alpha & -d-\alpha & -\beta-\alpha \end{bmatrix} \end{array}$$

D is a strictly dominant strategy in this three-strategy game. When costly punishment is not an option, this is a simplified PD game.

The one-shot PD game (with costly punishment) has payoff matrix:

$$\begin{array}{cc} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{bmatrix} b-c & -d-c & -\beta-c \\ b+d & 0 & -\beta+d \\ b-\alpha & -d-\alpha & -\beta-\alpha \end{bmatrix} \end{array}$$

D is a strictly dominant strategy in this three-strategy game. When costly punishment is not an option, this is a simplified PD game.

In our experiments (that had a 75% chance the game would continue after each round), we took $c = 1, d = 1, \alpha = 1, \beta = 4$ and b equal to either 2 or 3.

$$\begin{array}{cc} \mathbf{b=2} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{bmatrix} 1 & -2 & -5 \\ 3 & 0 & -3 \\ 1 & -2 & -5 \end{bmatrix} \end{array} \qquad \begin{array}{cc} \mathbf{b=3} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{bmatrix} 2 & -2 & -5 \\ 4 & 0 & -3 \\ 2 & -2 & -5 \end{bmatrix} \end{array}$$

$$\begin{array}{cc} \mathbf{b=2} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{bmatrix} 1 & -2 & -5 \\ 3 & 0 & -3 \\ 1 & -2 & -5 \end{bmatrix} \end{array} \qquad \begin{array}{cc} \mathbf{b=3} & \begin{array}{ccc} C & D & P \end{array} \\ \begin{array}{c} C \\ D \\ P \end{array} & \begin{bmatrix} 2 & -2 & -5 \\ 4 & 0 & -3 \\ 2 & -2 & -5 \end{bmatrix} \end{array}$$

Experimental Results

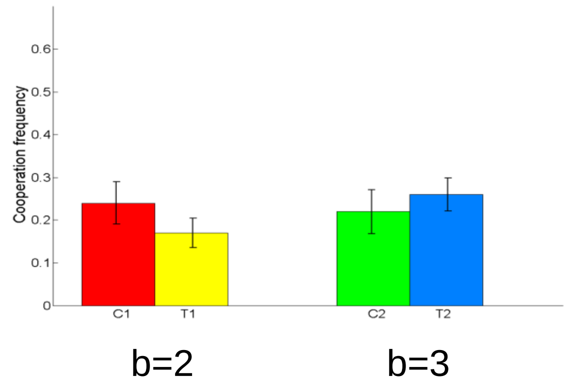

1. Increasing b did not increase the level of cooperation in the control experiments.
2. Adding costly punishment in the treatments did not significantly increase the level of cooperation.

Dreber, Rand, Fudenberg and Nowak (Winners don't punish, Nature 452, 348-351, 2008) carried out the same experiment and obtained the left-hand results.

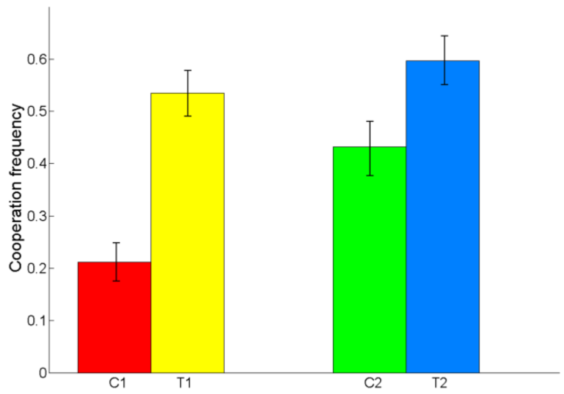

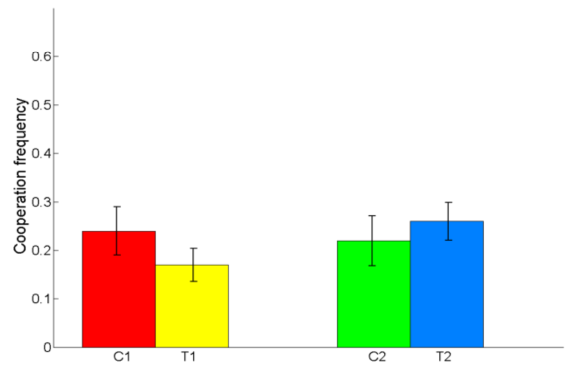

We claim the discrepancies between these two replications of the same experiments are based on cultural differences in attitudes to cooperation and to costly punishment.

The left-hand results arise from experiments conducted in Boston using college and university students as subjects whereas the right-hand experiments had subjects from universities in Beijing.

Reputation Effects

Dyadic Reputation: My behaviour toward you depends on your reputation from previous interactions between us (i.e. on how you have behaved toward me). This is related to direct reciprocity.

Guanxi Reputation: Reputation based on having a “guanxi” network formed by maintaining good relationships with more people who have good reputations themselves. Now my behaviour to you depends on how you behave with others, which forms the basis of indirect reciprocity.

There is empirical evidence that guanxi reputation is more important in China than dyadic reputation and that the opposite is true in Western culture.

Dyadic reputation is relevant for any repeated two-player game. It is then argued that levels of cooperation increase with the option of costly punishment since individuals who have a dyadic reputation of defecting run the risk of being punished.

On the other hand, guanxi networks are not possible to form in the repeated PD game. Thus, based on guanxi reputation, Chinese subjects have no reason to increase cooperation in the presence of costly punishment.

Dyadic reputation is relevant for any repeated two-player game. It is then argued that levels of cooperation increase with the option of costly punishment since individuals who have a dyadic reputation of defecting run the risk of being punished.

On the other hand, guanxi networks are not possible to form in the repeated PD game. Thus, based on guanxi reputation, Chinese subjects have no reason to increase cooperation in the presence of costly punishment.

On the other hand, Herrmann et al. (2008) found cooperation did increase with costly punishment in both the U.S. and China in the multi-player public goods game. However, in these games, both types of reputation discussed above are relevant.

More important for us is their result that cooperation (i.e. contribution levels) were similar in the U.S. and China for the multi-player public goods game (PGG).

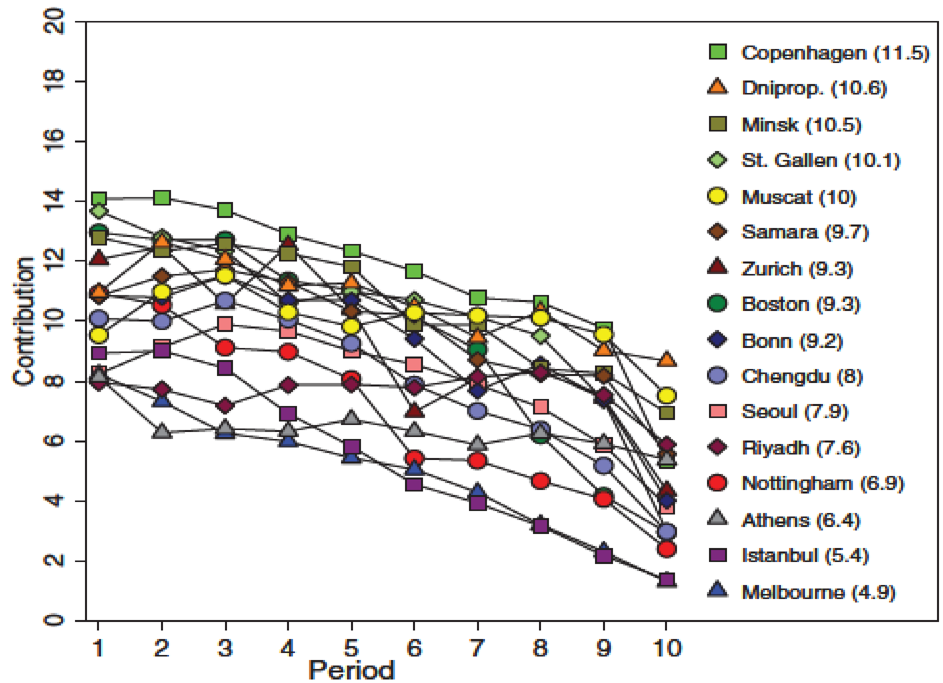

3. Mean contributions to the public good over the 10 periods of the N experiment. Each line ;ponds to the average contribution of a particular participant pool. The numbers in theses indicate the mean contribution (out of 20) in a particular participant pool.

More important for us is their result that cooperation (i.e. contribution levels) were similar in the U.S. and China for the multi-player public goods game (PGG).

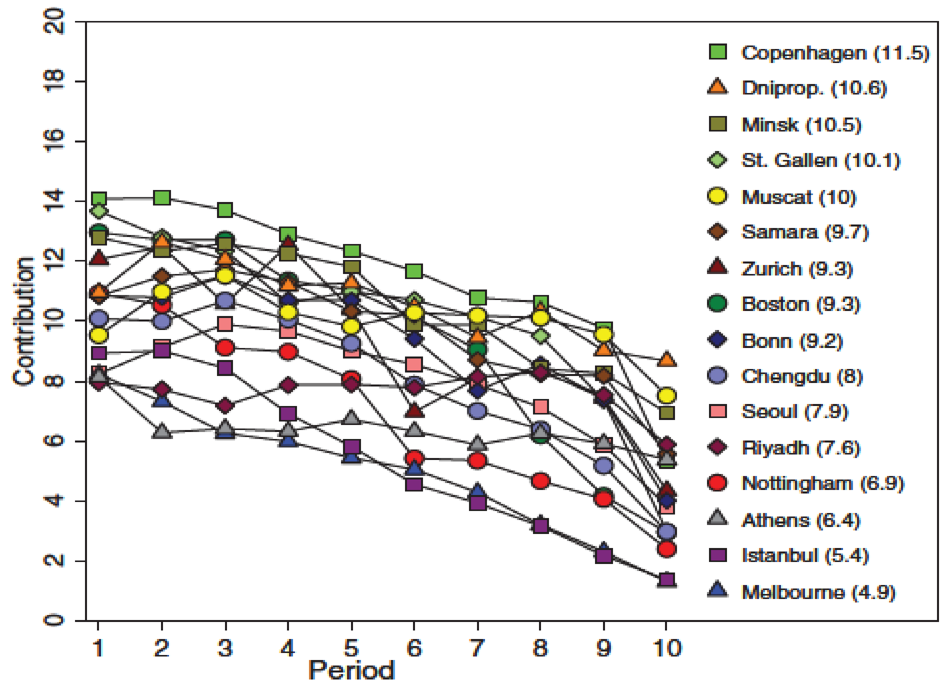

3. Mean contributions to the public good over the 10 periods of the N experiment. Each line ponds to the average contribution of a particular participant pool. The numbers in theses indicate the mean contribution (out of 20) in a particular participant pool.

We'll return to this later after I point out another major difference between the PD game experiment results in Beijing compared to Boston.

Distribution of first use of P

Consider those players who use P sometime during the repeated PD game with costly punishment. This is actually a small proportion of participants since the average use of P in any round is between 6 and 7 percent.

The following graph shows the distribution of which round these players first used P.

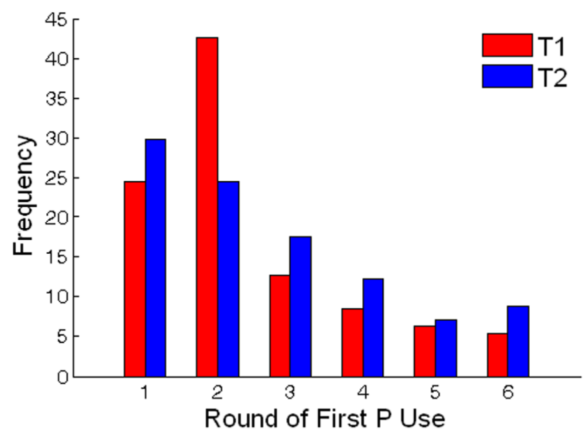

That is, 25 to 30 percent of P users do so in the first round. These players use P indiscriminately without knowing anything about their opponent's behavior. This could be an attempt to set oneself up as a dominant authority figure who is prepared to use P if the opponent does not behave as hoped (i.e. does not play C).

We take this as further evidence that developing a good reputation is not as important for Chinese subjects in two-player repeated PD games since guanxi networks cannot be established. We can again compare our results to those in Boston.

That is, 25 to 30 percent of P users do so in the first round. These players use P indiscriminately without knowing anything about their opponent's behavior. This could be an attempt to set oneself up as a dominant authority figure who is prepared to use P if the opponent does not behave as hoped (i.e. does not play C).

We take this as further evidence that developing a good reputation is not as important for Chinese subjects in two-player repeated PD games since guanxi networks cannot be established. We can again compare our results to those in Boston.

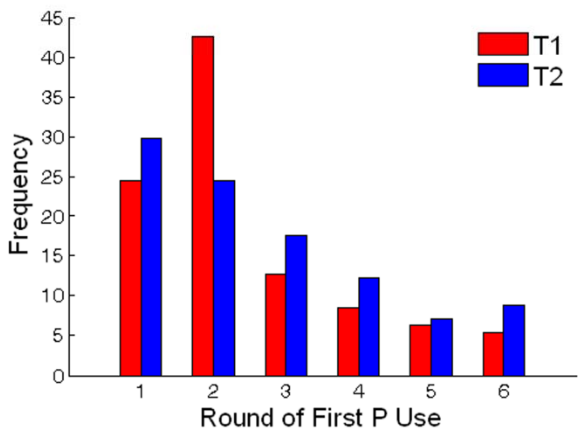

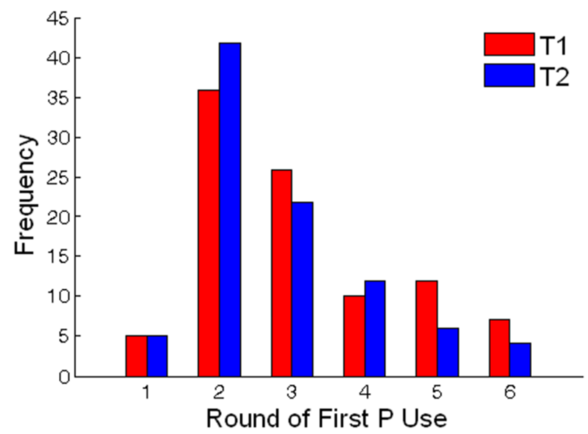

OUTLINE

Game Experiments on Cooperation

Repeated Prisoner's Dilemma (PD) Game

Cooperation with(out) costly punishment?
Experimental results in Beijing
Comparison to other experimental results
Reputation: dyadic versus network
Reputation and cultural differences
First use of costly punishment

Repeated Public Goods Game (PGG)

Cooperation without institutional incentives
Comparison to other experimental results
Institutional reward and/or punishment
Experimental results in Beijing
Comparisons to peer incentives
Effects on group incentives
Individual relative incentives
Evolution of cooperation

Repeated public goods game (PGG)

Our experimental setup:
Fifty rounds with four participants in a group.
In each round, each particpant is given 20 “units”.
Participants decide how much to put in public pool and how much to keep for themselves.
Contributions in public pool increased by 60% and distributed evenly to all group members.

Repeated public goods game (PGG)

Our experimental setup:
Fifty rounds with four participants in a group.
In each round, each particpant is given 20 “units”.
Participants decide how much to put in public pool and how much to keep for themselves.
Contributions in public pool increased by 60% and distributed evenly to all group members.

Let $c_{j,k}$ be player i's contribution in round k.
Then the payoff $\pi_{j,k}$ of player i in round k is

$$\begin{aligned}\pi_{j,k} &= 20 - c_{j,k} + \frac{1.6}{4}\sum_{i=1}^{4} c_{i,k} \\ &= 20 - 0.6c_{j,k} + 0.4\sum_{i=1,i\neq j}^{4} c_{i,k}.\end{aligned}$$

The only Nash equilibrium is $c_{j,k} = 0$ for $j = 1, ..., 4$ and $1 \leq k \leq 50$ with $\pi_{j,k} = 20$.
On the other hand, $\pi_{j,k} = 32$ if all participants contribute 20.

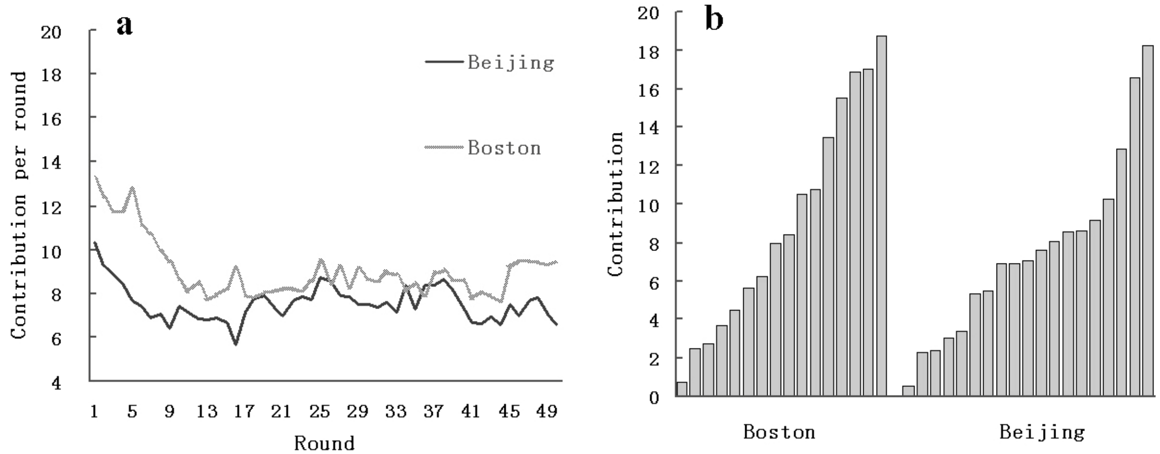

a
b
Beijing
Boston
Contribution per round
Round
Contribution
Boston
Beijing

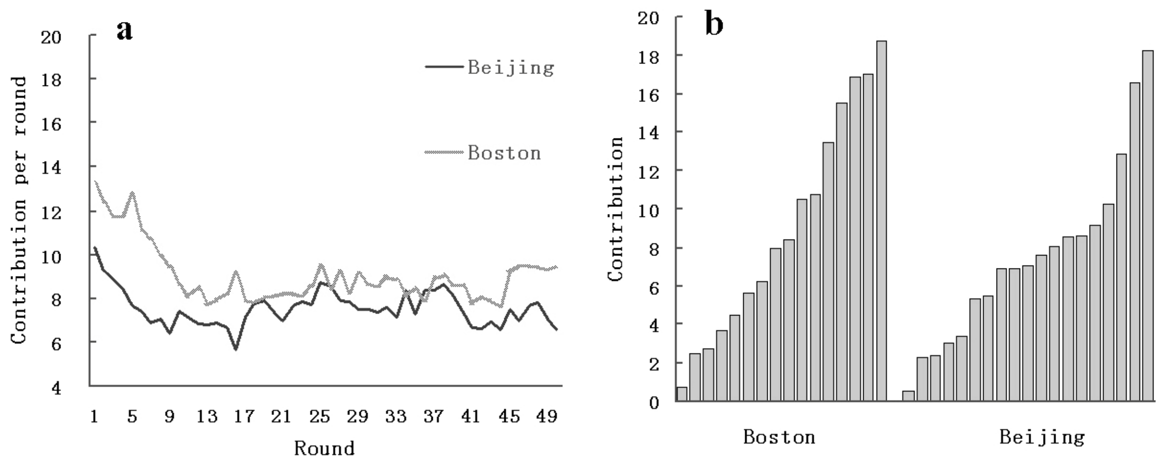
a
Beijing
Boston
Contribution per round
Round
b
Contribution
Boston
Beijing

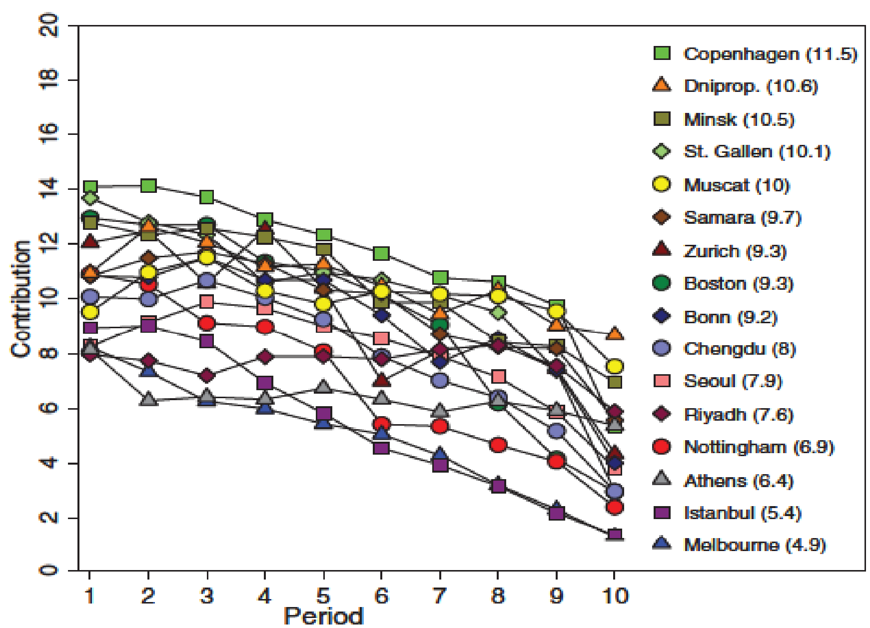
Copenhagen (11.5)
Dniprop. (10.6)
Minsk (10.5)
St. Gallen (10.1)
Muscat (10)
Samara (9.7)
Zurich (9.3)
Boston (9.3)
Bonn (9.2)
Chengdu (8)
Seoul (7.9)
Riyadh (7.6)
Nottingham (6.9)
Athens (6.4)
Istanbul (5.4)
Melbourne (4.9)
Contribution
Period

Sefton, Schupp, Walker (2007)

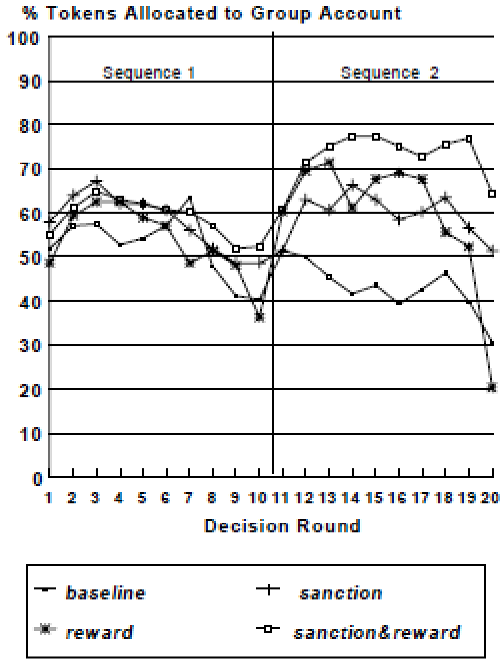

Andreoni, Harbaugh, Vesterlund (2003)

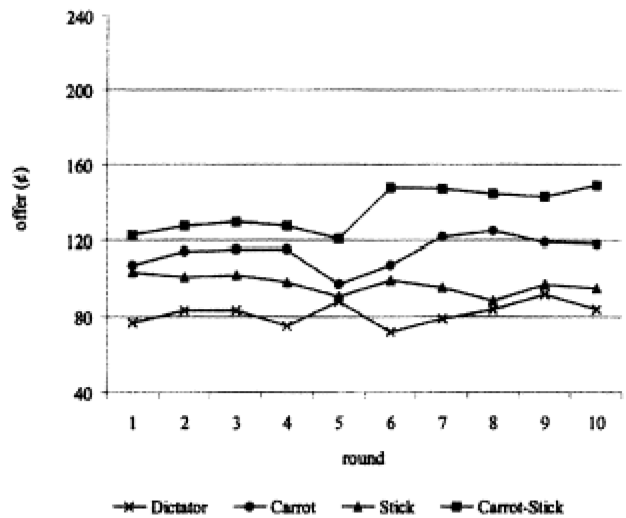

FIGURE 2. AVERAGE OFFERS

Institutional Incentives

Institutional Reward: At the end of each round, one group member is chosen with probability increasing as the amount this person contributes to the public pool increases.

This person is given a reward equal to 20% of the total group units.

Institutional Punishment: At the end of each round, one group member is chosen with probability decreasing as the amount this person contributes to the public pool increases. This person is punished by an amount equal to 20% of the total group units.

Institutional Reward and Punishment: At the end of each round, one group member is chosen to be given a 20% reward with above probability. Independently, one group member is chosen to be punished at 20% with above probability.

Institutional Reward and Punishment: At the end of each round, one group member is chosen to be given a 20% reward with above probability. Independently, one group member is chosen to be punished at 20% with above probability.

Forced Ranking

Typical system: At regular intervals (e.g. yearly) employees are ranked as to their performance to the institution.

Those ranked in the bottom 10% are fired.

Those ranked in the top 20% are given monetary bonuses or other rewards.

Institutional Reward and Punishment: At the end of each round, one group member is chosen to be given a 20% reward with above probability. Independently, one group member is chosen to be punished at 20% with above probability.

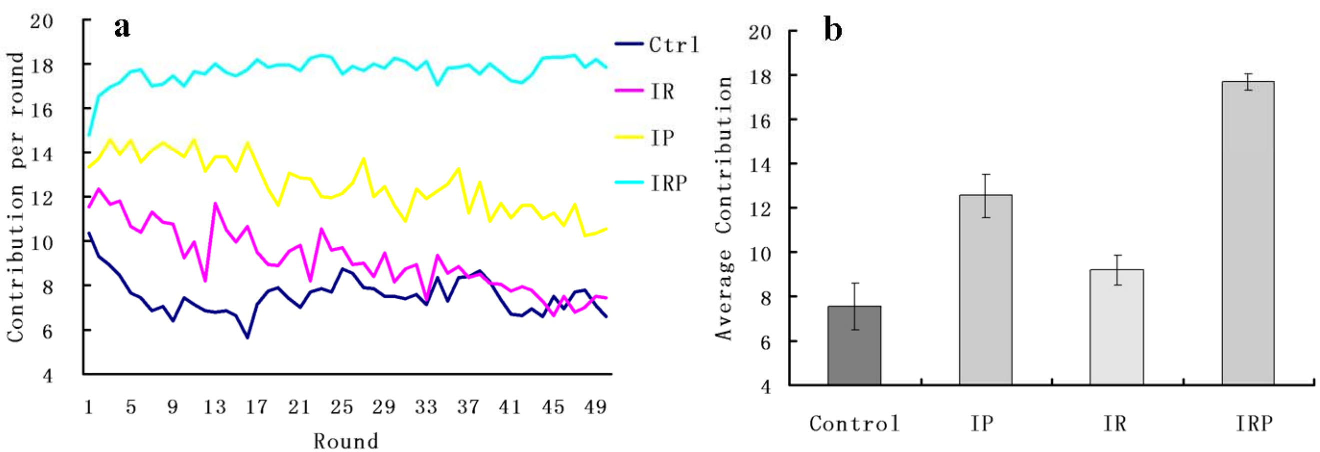

Institutional Incentives (Beijing)

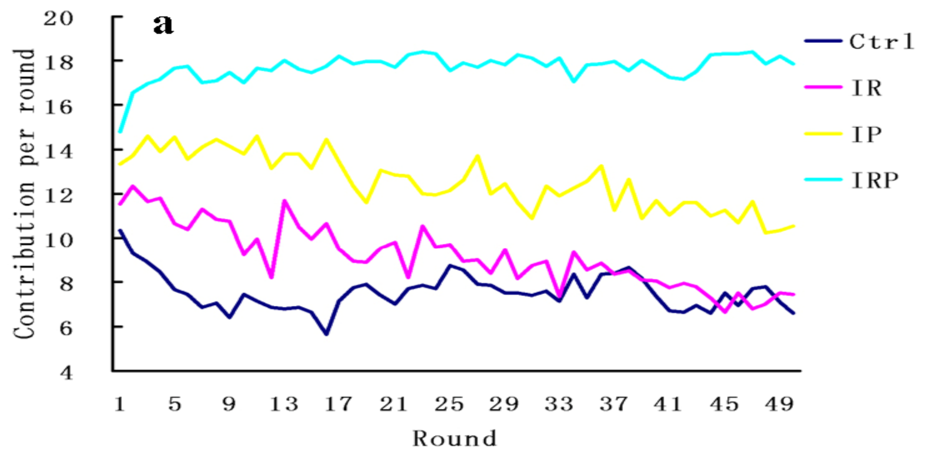

Peer Incentives (Boston)

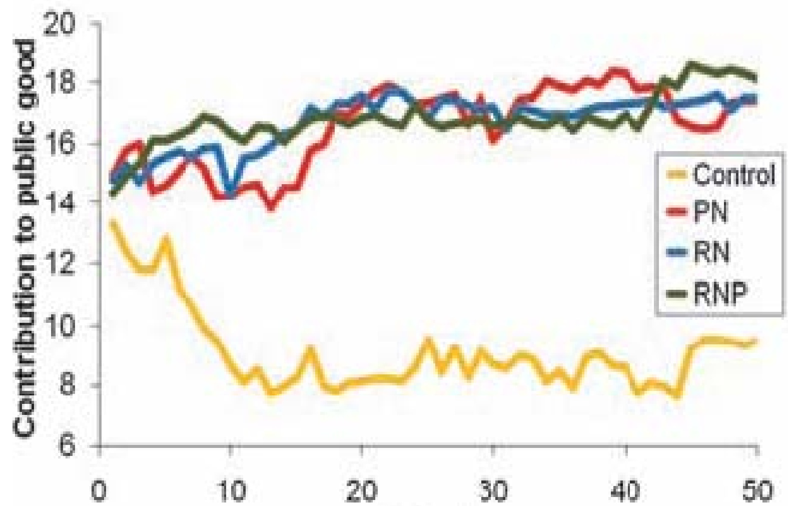

There have been many studies (both in theory and by experiments) of PGG with peer incentives. To explain the increased level of cooperation under reward and/or punishment, the effects of these incentives on group payoffs are often examined.

This approach does not work for our institutional incentives model. Specifically, group payoff is an elementary linear function of total group contribution.

Let $\overline{c}_k$ be the average group contribution in round k of our experiments.

Let π_k^C, π_k^{IR}, π_k^{IP}, π_k^{IRP} be the average group payoff in round k for C, IR, IP, IRP respectively. Then

$$
\begin{aligned}
\pi_k^C &= 20 + 0.6\overline{c}_k \\
\pi_k^{IR} &= 24 + 0.72\overline{c}_k \\
\pi_k^{IP} &= 16 + 0.48\overline{c}_k \\
\pi_k^{IRP} &= 20 + 0.6\overline{c}_k.
\end{aligned}
$$

Thus, the maximum group payoff occurs when all four members contribute 20 in all cases. These maximum payoffs are $32, 38.4, 25.6, 32$ for C, IR, IP, IRP respectively. A more accurate comparison of group contribution on group payoff is to examine the percentage of maximum payoff (called "group efficiency" in the literature).

$$\frac{\textbf{group payoff}}{\textbf{maximum payoff}}$$

In all four protocols, this is simply $62.5 + 1.875\bar{c}_k$.

That is, from the group perspective, there is the same payoff incentive in all four protocols to contribute a given amount to the public pool.

Thus, the maximum group payoff occurs when all four members contribute 20 in all cases. These maximum payoffs are 32, 38.4, 25.6, 32 for C, IR, IP, IRP respectively. A more accurate comparison of group contribution on group payoff is to examine the percentage of maximum payoff (called "group efficiency" in the literature).

$$\frac{\textbf{group payoff}}{\textbf{maximum payoff}}$$

In all four protocols, this is simply $62.5 + 1.875\bar{c}_k$.

That is, from the group perspective, there is the same payoff incentive in all four protocols to contribute a given amount to the public pool.

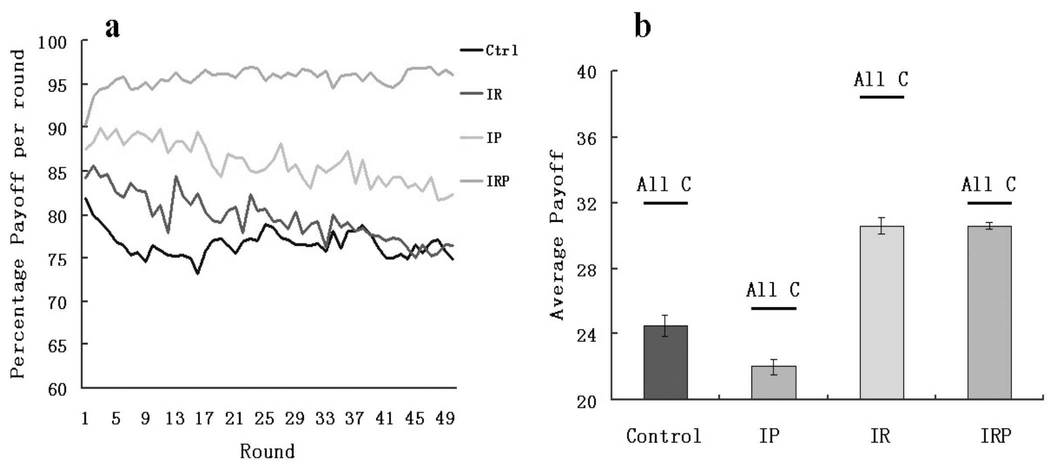

Relative Contributions and Payoffs

Instead of group contributions effects on payoffs, our experimental results can be understood through comparisons of individuals to their group.

Relative contribution of individual i in round k:

$$x_{j,k} \equiv c_{j,k} - \overline{c}_k \; where \; \overline{c}_k \equiv \frac{1}{4} \sum_{i=1}^{4} c_{i,k}.$$

Relative payoff:

$$y_{j,k} \equiv \pi_{j,k} - \overline{\pi}_k \; where \; \overline{\pi}_k \equiv \frac{1}{4} \sum_{i=1}^{4} \pi_{i,k}.$$

By definition, in PGG without incentives (i.e. in the control experiment C)

$$y_{j,k} = -x_{j,k}.$$

That is, the NE of PGG is for everyone to free ride (i.e. contribute nothing).

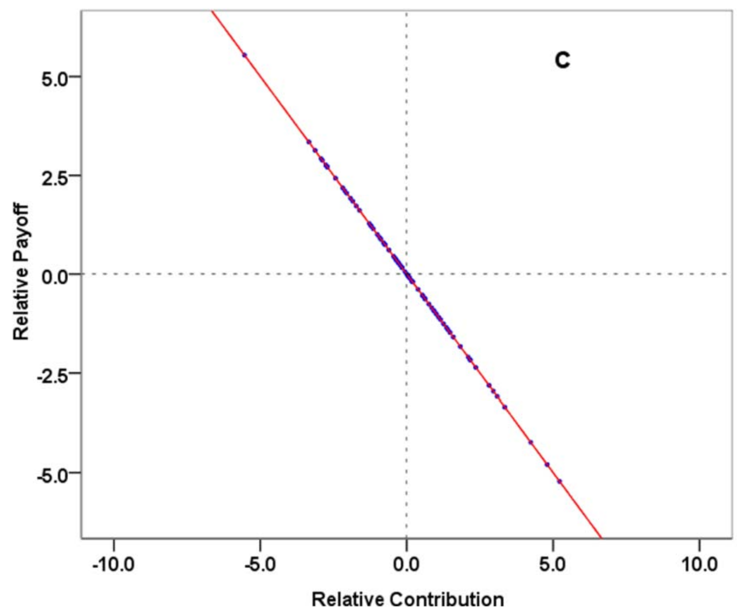

Due to the stochastic nature of rewards and/or punishments in the three treatments, there is no exact equation for $y_{j,k}$ in terms of $x_{j,k}$.

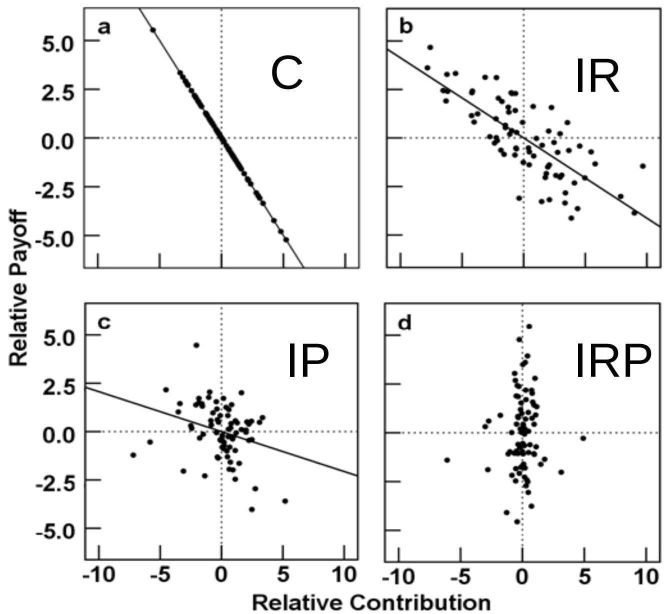

Due to the stochastic nature of rewards and/or punishments in the three treatments, there is no exact equation for $y_{j,k}$ in terms of $x_{j,k}$.

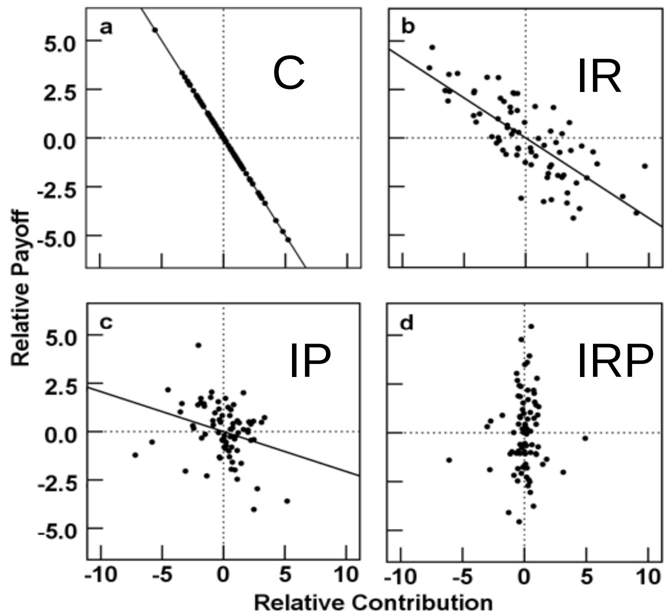

We take the negative slope of the lines of regression as a measure of the advantage to an individual of free riding. This advantage is strongest in C (-1), then in IR (-0.413) and finally in IP (-0.206). In fact, by this measure, there is no free-riding advantage in IRP since there is no correlation there.

Number of incentives versus relative payoff

One would expect that an individual's relative payoff should increase with the more rewards he/she receives and decrease with the number of punishments.

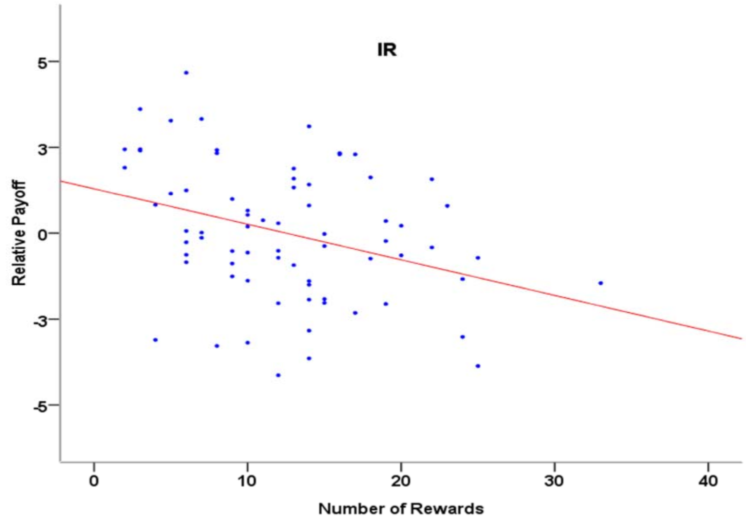

Number of incentives versus relative payoff

One would expect that an individual's relative payoff should increase with the more rewards he/she receives and decrease with the number of punishments.

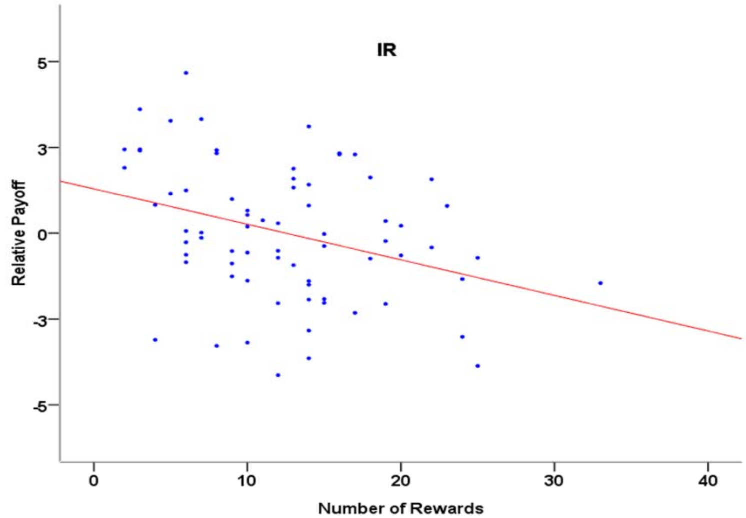

This seemingly counterintuitive result means that the expected payoff gain from more rewards (due to higher contribution level) in IR does not offset the relatively high free-riding benefit of contributing less.

Similar counterintuitive results do NOT occur in IP or IRP.

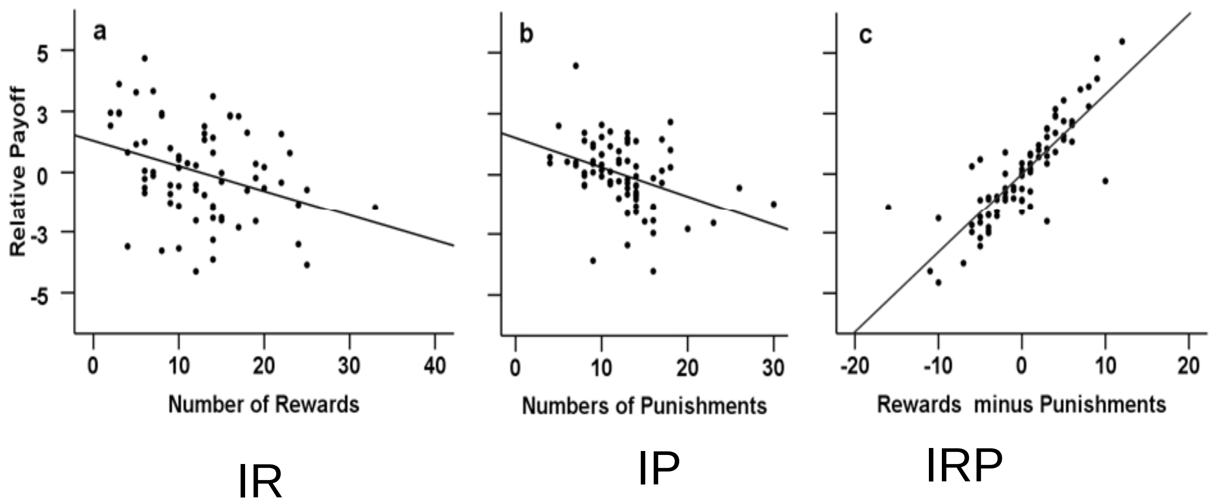

Thus, it is then not surprising that IR is the least effective incentive in increasing cooperation.

Number of incentives versus relative contribution:

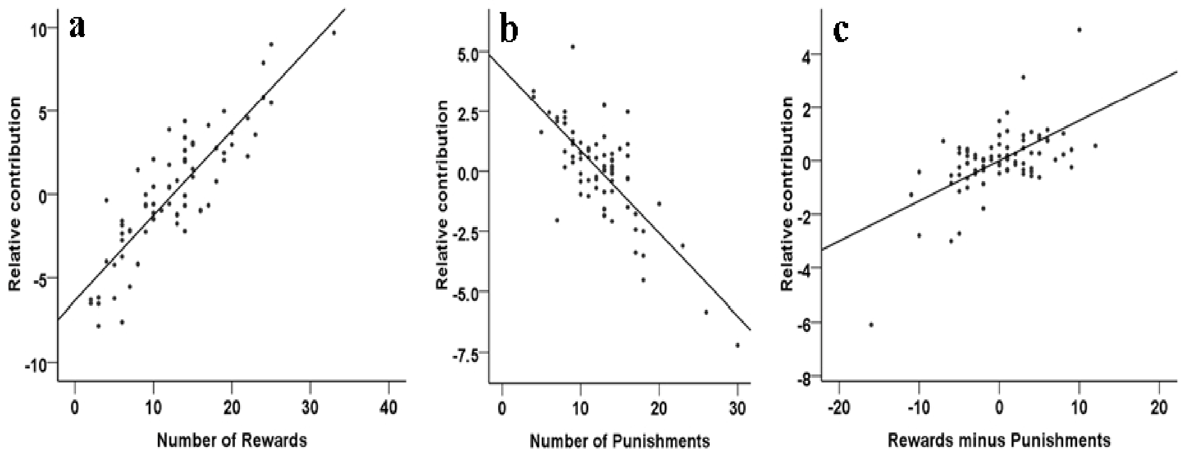

IR IP IRP

Notes

IRP is the only incentive scheme where cooperation rises in later rounds of the repeated game. This suggests an evolutionary advantage to systems that incorporate both mechanisms.

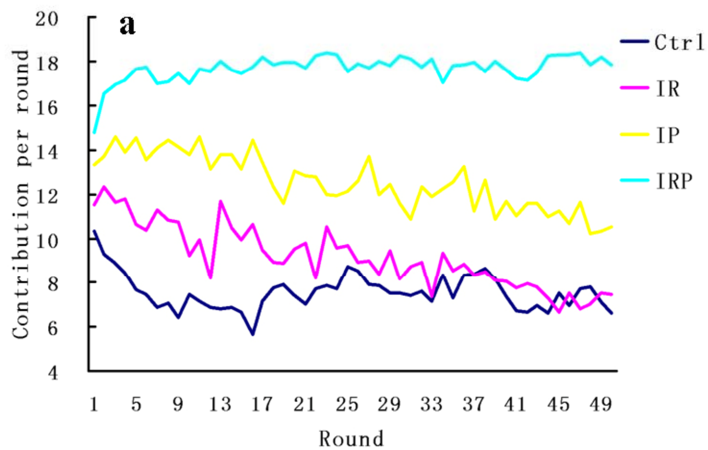

In IRP, all individuals in a group tend to make almost the same contribution. Should this incentive scheme be continued indefinitely?

In IRP, all individuals in a group tend to make almost the same contribution. Should this incentive scheme be continued indefinitely?

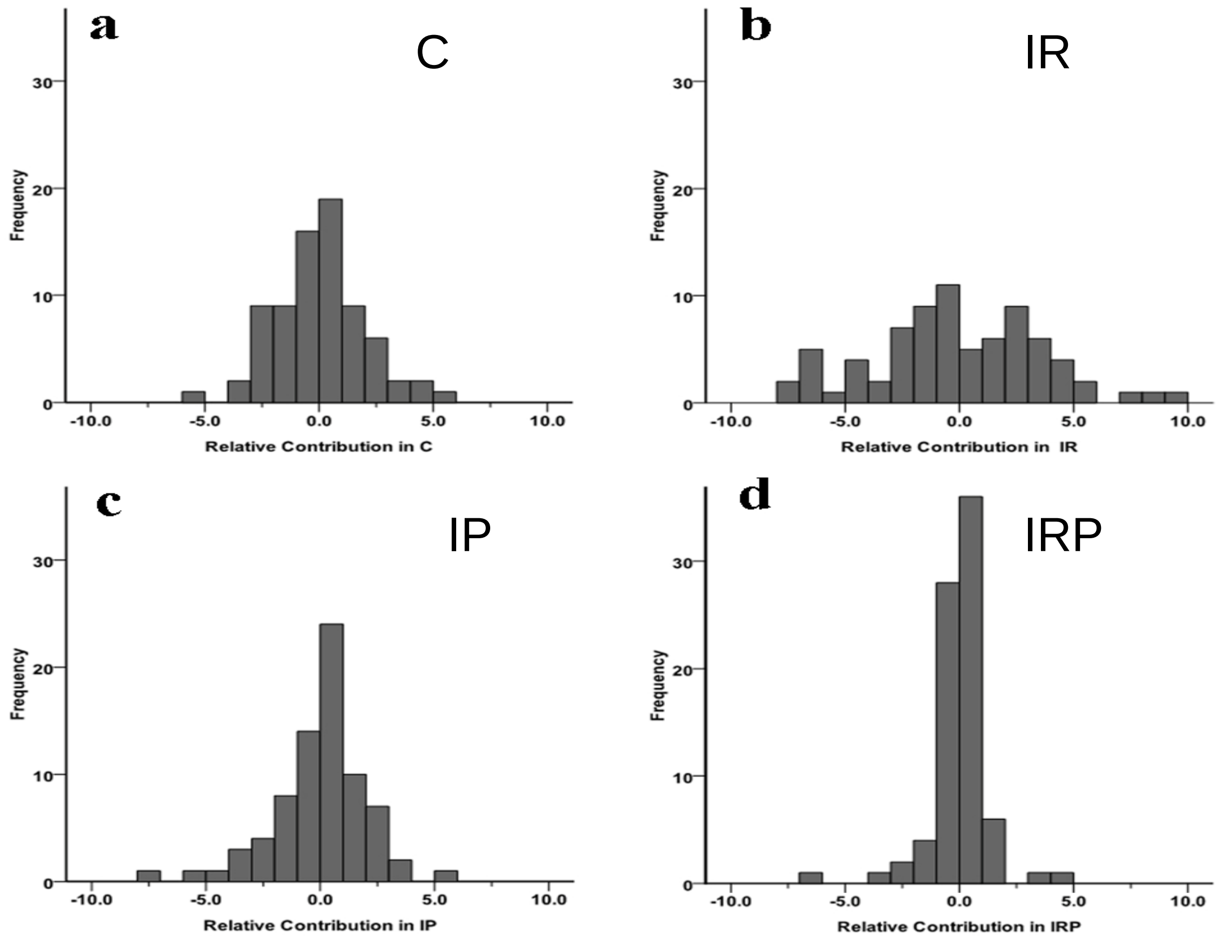

Future research involves designing experiments to further examine how cultural differences affect cooperation in repeated PD games, perhaps by including interactions between subjects of varying cultural backgrounds.

Thank You

14th International Symposium on Dynamic Games and Applications (ISDG)
June 20-23, 2010
The Banff Centre, Banff, Alberta, Canada