

The Carrot or the Stick: Rewards, Punishments, and Cooperation

By JAMES ANDREONI, WILLIAM HARBAUGH, AND LISE VESTERLUND*

Many institutional arrangements suggest that punishments and rewards each play a separate role in providing incentives. In New York City's recent negotiations with its teacher's union, for instance, the city sought a contract that would strengthen school principals' ability to assign teachers to be cafeteria monitors, in part because this allows them to punish underperforming teachers. The contract also includes rewards for teachers with good performance. Another example is that some universities now use a combination of raises and differential teaching loads to encourage good performance. Similarly, procurement and production contracts and government regulations in areas ranging from meat inspection to sulphur dioxide pollution often include both bonuses for good performance and various sorts of clawbacks for bad.

Punishments and rewards also play an important role in informal relationships. We seem to care about the unwritten social contract that people should be cooperative and kind to others, and we are often willing to pay to enforce these ideals. For instance, every time we take a taxi or sit in a restaurant we entrust our happiness to another person. If that trust is protected then we may reward it with a generous tip, but if not we may leave less than usual, or nothing at all. Likewise we may shun unfriendly colleagues but invite the friendly ones to our homes, and

secretaries may perform more promptly for those who are polite or bring gifts and less promptly or poorly for those who are rude or unfriendly. Absent complete contracts, voluntary punishments and rewards are the mechanisms we use to sustain cooperation. Understanding what triggers the demands for each and their effectiveness may help us understand informal relations and to design institutions that can harness cooperation and improve social welfare.

The objective of this paper is to begin a systematic look at both punishments and rewards in economic laboratory experiments.¹ Experimental studies have demonstrated substantial demands for both punishments and rewards.² In one-shot proposer-responder games people are willing to sacrifice personal payoffs to punish those who are unkind, and reward those who are kind. These studies typically examine either rewards or punishments, but not both.³ Thus they make it difficult to determine what motivates these demands, how they interact, or how they affect cooperation.

¹ To our knowledge we are the first to study these questions in a simple proposer-responder setting. David L. Dickinson (2001) examines a team production problem where subjects may be exogenously rewarded or punished. He finds that this mechanism increases efficiency. Martin Sefton et al. (2002) examine repeated linear public goods games and find that combined use of rewards and punishments result in the most generous public good contributions. In sequential public goods games Andreoni et al. (2002) find that subjects punish small contributions but seldom reward generous contributions.

² To keep the exposition simple we will refer in two-stage games to a punishment as a decrease in payoff that a second-stage player imposes on a first-stage player. Similarly we refer to a reward as an increase in payoffs at the second stage. As examples of punishments see Werner Güth et al. (1982), Robert Forsyth et al. (1994), Gary E Bolton and Rami Zwick (1995), Ernst Fehr and Simon Gächter (2000), and Andreoni et al. (2003). Examples of a demand for rewards are Richard D. McKelvey and Thomas R. Palfrey (1992), Fehr et al. (1993), Joyce Berg et al. (1995), Gary Charness (1996), Charness et al. (2000), and Charness and Ernan Haruvy (2002).

³ Exceptions are Klaus Abbink et al. (2000) and Theo Offerman (2002).

* Andreoni: Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706 (e-mail: andreoni@wisc.edu); Harbaugh: Department of Economics, 1285, University of Oregon, Eugene, OR 97403 (e-mail: bill@harbaugh.org); Vesterlund, Department of Economics, 4T18 Posvar Hall, University of Pittsburgh, Pittsburgh, PA 15260 (e-mail: vester@pitt.edu). This research was in part conducted while Harbaugh and Vesterlund were visiting the Center on Philanthropy and the Department of Economics at IUPUI, and while Vesterlund later was visiting Harvard Business School. We are grateful for their hospitality and financial support. We also acknowledge support from the National Science Foundation. Finally, Vesterlund acknowledges financial support from Iowa State University.

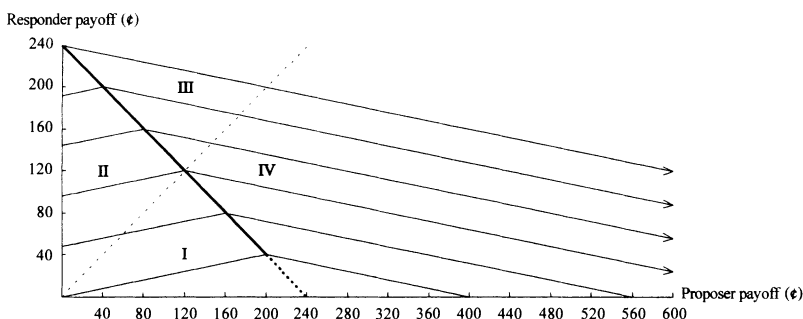


FIGURE 1. POSSIBLE PAYOFFS IN THE CARROT-STICK

We examine punishments and rewards both separately and jointly. We study a series of two-person proposer-responder games with costly punishments and rewards. Proposers choose how much to share of a fixed pie. Responders are in one of four conditions: punish or reward, reward only, punish only, or neither. By considering two-person games we avoid free-riding on punishments by others, and by randomly changing partners we avoid any repeated game effects. By looking at rewards and punishments both separately and jointly we can identify any interaction or complementarity.

We find substantial demands for both rewards and punishments. As expected, an increase in the offer by proposers, on average, decreases the punishment and increases the reward. Interestingly, while the average demand for punishment appears to be independent of the reward option, we find that the demand for rewards is significantly larger when the responder doesn't have the option of punishing. This suggests that the ideals subjects try to enforce may be affected by the tools they have available. We also find that, on average, the proposed offer is largest when a combination of rewards and punishments are available, and smallest when neither option is available. Although the average offer of the rewards-only treatment exceeds that of punishments-only, we find that rewards are much less effective in moving the proposers away from the minimum possible offer. Thus one might expect less cooperation in societies where good behavior is rewarded than in those where poor behavior is punished. Designing an institution around rewards only and omitting an option for punishments may be a mistake, even if in the end punishments are rarely used.

I. Experimental Design

We consider four variations of a proposer-responder game. In the first stage of each variation a proposer decides what portion of \$2.40 he wants to transfer to the responder. The only difference between the four treatments is in the responder's ability to punish or reward at the second stage. We refer to the treatments as: *Dictator*, *Carrot-Stick*, *Carrot*, and *Stick*. In the Dictator treatment the responder can neither punish nor reward. In the Carrot-Stick treatment the responder can, at a cost of 1 cent, increase or decrease the proposer's earnings by 5 cents.⁴ In the Stick treatment the responder can, at a cost of 1 cent, decrease the proposer's earnings by 5 cents. Finally, in the Carrot treatment the responder can, at a cost of 1 cent, increase the proposer's earnings by 5 cents.

The set of payoff combinations that are available in the Carrot-Stick game is illustrated in Figure 1. The proposer's payoff is measured on the horizontal axis and the responder's payoff

⁴ For identification purposes we wanted the cost of punishments and rewards to be independent of the offer. This is not the case in the ultimatum game nor in Fehr and Gächter (2000) where a larger offer increases both the perceived generosity and the cost of punishing. While there are cases where the cost of say a \$1 reward is \$1, there are many others where the cost versus consequence of rewards and punishments is not one-for-one. For example, the cost of complaining over bad service in a restaurant is likely to be much smaller than the consequences imposed on the waiter; similarly the cost of praise may be smaller than the benefit. As the objective of this study is to examine the interaction between rewards and punishments we chose a cost-consequence ratio where we expected to observe both.

on the vertical axis. The proposer chooses an offer along the bold solid line, and conditional on that offer the responder has the option of choosing any point on the reward and punishment lines originating at the proposer's offer. These choices are indicated by the lighter solid lines. To secure the responder the opportunity of decreasing the proposer's payoff to zero, we do not allow the proposer to make offers below 40 cents. The dashed portion of the bold line shows the proposals excluded by this rule.

The 45° line, shown with light dashes, indicates the possible payoff combinations that result in equal payoffs to the responder and proposer. While the responder decreases her payoff by choosing any outcome off of the bold line, note that moves into areas I and III result in a more equal distribution of payoffs than under the original proposal. In contrast, subjects who pick outcomes in areas II and IV could have chosen an outcome which both resulted in a larger personal payoff as well as a more equal distribution of payoffs.

The feasible payoff combinations in the Carrot treatment include all points on or to the right of the bold line, the Stick treatment include points on or to the left of the bold line, while the Dictator treatment only includes points on the bold line. If it is common knowledge that all individuals seek to maximize their personal payoff, then the subgame-perfect equilibrium outcome is the same for all treatments. The responder should neither reward nor punish, and given this response the proposer should choose the minimum required transfer of 40 cents.

We conducted three sessions of each treatment, each with ten proposers and ten responders, for a total of 30 subjects in each role in each treatment. Subjects were undergraduate business students and were randomly assigned to a treatment. Upon arriving to the experiment they were randomly assigned to a computer terminal and were given a set of written instructions. The experimenter read the instructions aloud, after which the subjects were asked to calculate the payoffs in a specific example of the game. The answers to the quiz were collected, and the example was reviewed verbally by the experimenter. Half of the subjects were then randomly assigned to be proposers and half to be responders. They remained in that role throughout the experiment.

Subjects played ten iterations of the game. In each iteration subjects were randomly and anonymously paired, with the stipulation that no two subjects met more than once. Subjects' identities were never revealed to one another. Subjects' earnings for all ten rounds were tallied and added to a \$5 show-up payment. While waiting for their payment, subjects answered a questionnaire. They were paid anonymously with cash in envelopes which were handed out by subject number. The experiment typically lasted less than an hour, and including the show-up fee the average earnings were \$17.41 (standard deviation of \$4.80, maximum of \$49.35, and minimum of \$6.70).⁵ The instructions were kept as neutral as possible by referring to the punishments and rewards simply as changes to the proposer's payoff. A copy of the instructions is available from the authors.

II. Results

We first examine whether the observed behavior is consistent with the standard subgame-perfect equilibrium prediction. Next we examine how the responder's options affect the cooperative behavior of the proposer, and discuss the demands for rewards and punishments.

A. Equilibrium Predictions

Figure 2 shows the average offers in each of the ten rounds for the four treatments. In each treatment and in each of the ten rounds we reject the hypothesis that the average offer equals 40 cents.⁶ Furthermore offers are not independent of the treatment. The average offer is largest in the Carrot-Stick treatment and smallest in the Dictator treatment.⁷ As expected the offers do

⁵ The subject with the highest earnings was a proposer in the Carrot treatment. That subject made generous offers and received total rewards of \$36. The subject with the lowest earnings was a proposer in the Stick treatment. This subject never made an offer below 200, and indicated in the questionnaire that her main objective was to avoid punishments.

⁶ Treating each proposer as an observation we test if the average proposal equals 40. Across the ten rounds the *t*-statistic is 12.6 for Carrot, 10.7 for Carrot-Stick, 9.6 for Stick, and 8.0 for Dictator. The results for the last round and the last five rounds are similar.

⁷ Treating each proposer as an observation, we use a *t*-test to determine if average offers differ across treatments.

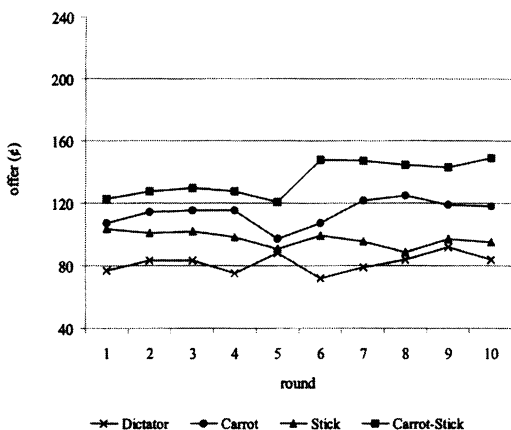


FIGURE 2. AVERAGE OFFERS

not decrease towards the end of the experiment, indicating that the proposers correctly understood the one-shot nature of the interactions.

Next we turn to the responders' decisions. Here we find statistically significant evidence of both punishments and rewards.⁸ In fact the willingness to reward and punish is substantial. Across the three treatments responders changed the proposer's payoff 43 percent of the time, and over the ten rounds 80 percent of the responders chose to make at least one change of the proposer's payoff. During the last five rounds, 75 percent of the responders changed the proposer's payoff.

Figure 3 shows the proportion of offers that were either increased or decreased by the responder in a given round.⁹ Once again we observe differences across treatments. During the last five rounds the average responder in the

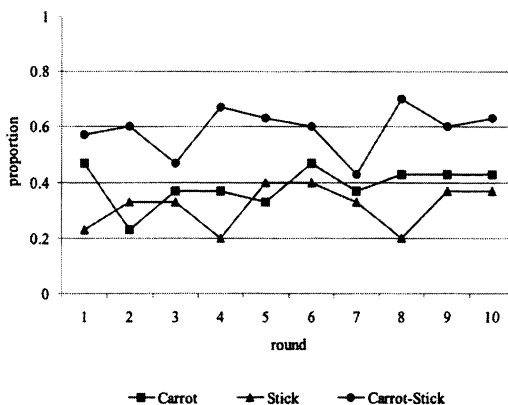


FIGURE 3. AVERAGE RESPONSES

Carrot treatment spends 20.4 cents changing the proposer's payoff. In comparison, the number in the Carrot-Stick is 14.7, and in the Stick treatment it is only 5.9.¹⁰

B. Cooperation Production Function

We saw in Figure 2 that the proposers' offers differ across treatments, with Carrot-Stick offers the highest (136 on average), followed by Carrot (114), then Stick (97), and Dictator (82).¹¹ This seems to indicate that punishments alone are the least effective in moving the proposer away from the minimum possible offer. However, when examining the distribution of offers we find a much more subtle and interesting relationship between rewards, punishments, and cooperation.

Figure 4 shows the distribution of offers over the last five rounds. Notice that the modes of the distributions differ substantially across treatments. The most frequent offer in the Dictator is 40, the minimum possible. Adding just rewards leaves the modal offer at 40, but greatly in-

Differences are significant at the 5-percent level for the ten rounds, at the 10-percent level for the last five rounds, and with exception of the Dictator-Stick comparison at the 10-percent level for the last round.

⁸ For all three treatments we reject the hypothesis that rewards and/or punishments are zero. Treating each responder as an observation we test if responders on average chose not to change the proposer's payoff. Over the ten rounds the *t*-statistic is 4.6 for Carrot, 7.5 for Carrot-Stick, and 5.9 for Stick. The results for the last round and the last five rounds are similar.

⁹ Neither the likelihood nor the expenditure on changes decreases over time. This suggests that both proposers and responders behaved in a manner consistent with the one-shot interaction. This contrasts with the repeated game result of Sefton et al. (2002).

¹⁰ Over all ten rounds the average cost of changing payoffs are 14.7 in Carrot, 13.3 in Carrot-Stick, and 5.7 in Stick. Treating each responder as an observation we test for treatment differences. The *t*-statistic is 0.35 for the Carrot and Carrot-Stick comparison, 2.68 for the Carrot and Stick comparison, and 3.75 for the Carrot-Stick and Stick comparison. The results for the last round and the last five rounds are similar.

¹¹ The data for the last five rounds is similar, with 146 for Carrot-Stick, 118 for Carrot, 95 for Stick, and 82 for Dictators.

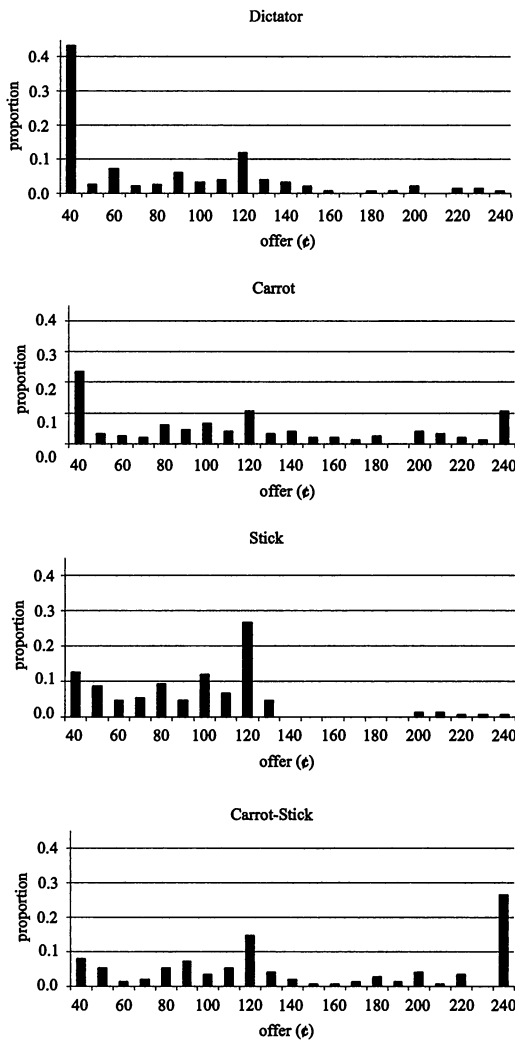


FIGURE 4. DISTRIBUTION OF OFFERS FOR LAST FIVE ROUNDS

increases the variance in offers. A sizable minority of people, 10 percent, actually share everything by offering 240 (and hence raise the average in this treatment). Except for this small group of generous individuals, most of the proposers actually behave as in the Dictator game. This indicates that for most subjects, rewards alone are not effective. The Stick treatment, however, moves the behavior dramatically. The modal offer is now 120, with virtually no offers above 120. Hence, punishments can move people from the selfish to the equitable offer. Fi-

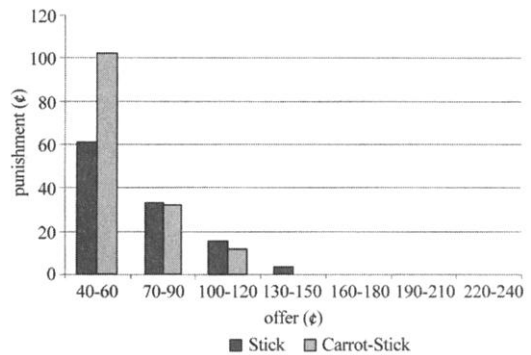


FIGURE 5. AVERAGE PUNISHMENTS FOR LAST FIVE ROUNDS

nally, the Carrot-Stick has a substantial effect, further increasing the mode all the way to 240, the maximum possible offer. Overall, the distribution of offers in the Carrot-Stick treatment first-order stochastically dominates that of the three other treatments. Thus, while adding rewards only had little effect, adding rewards to punishments has a profound effect. In other words, rewards and punishments seem to act as complements in encouraging proposers to increase their offers.

C. Demands for Punishments and Rewards

Punishments and rewards are common. More than half of the offers in excess of 120 result in an increase in the proposer's payoff, and more than 40 percent of the offers below 120 result in a decrease in the proposer's payoff.¹² In this subsection we examine how these punishments and rewards differ with proposals as well as with the treatment.

We start with punishments in the Carrot-Stick and Stick treatments. The demands for punishments during the last five rounds are shown in Figure 5. First note that the average punishment decreases with the size of the proposer's offer. This is similar to the ultimatum

¹² Across all ten periods the proportion of offers above 120 that are rewarded is 54 percent in the Carrot and 59 percent in the Carrot-Stick. Of offers below 120 the proportion rewarded is 31 percent in Carrot and 11 percent in Carrot-Stick. Of offers below 120 the proportion punished is 43 percent in the Stick and 51 percent in the Carrot-Stick, and of offers above 120 the proportion punished is 8 percent in the Stick and 2 percent in the Carrot-Stick.

game, where the cost of punishing is not independent of the offer. Second, the decrease in punishment is rather steep—after the equal split there are essentially no punishments. Of the offers of 120 approximately 15 percent result in some decrease in the proposers payoff.¹³ Finally, the responder's ability to reward has limited effect on the demand for punishments.¹⁴ The punishments of medium-sized offers are the same across treatment, and although punishments of low offers appear larger in the Carrot-Stick treatment, this difference is only significant at the end of the experiment.¹⁵ This finding is surprising. One would not expect rewards of small offers, yet it is for small offers that the availability of rewards appears to affect the demand for punishments.

Next we examine the demand for rewards in the Carrot-Stick and Carrot treatments. The demands for rewards in the last five rounds are shown in Figure 6. Similar to the evidence from the trust game, a higher offer on average leads to a larger reward. But the data reveal a result which one may not have expected based on the results from trust games: average rewards at or below equal-split offers are quite substantial.¹⁶ Another puzzling finding is the substantial difference in the demand for rewards between the

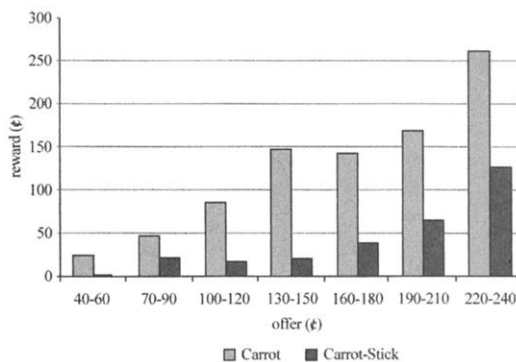


FIGURE 6. AVERAGE REWARDS FOR LAST FIVE ROUNDS

Carrot and Carrot-Stick treatments.¹⁷ Conditional on the offer, the average reward given in the Carrot is larger than in the Carrot-Stick.¹⁸ Even for very large offers, where we know there is no demand for punishment, we see that the absence of the ability to punish results in substantially larger average rewards.

While punishments are unaffected by the availability of rewards, we find that the rewards are larger when there is no option of punishing. It is the combination of these results that makes them puzzling. If one were presented simply with the reward results, then it is tempting to argue that the difference is caused by a simple substitution effect. If rewards and punishments are viewed as substitutes in enforcing a certain level of proposer kindness, then the demand for either one may be decreasing in own cost and

¹³ During the last five rounds the proportion of 120 offers punished is 17 percent in the Stick and 13 percent in the Carrot-Stick.

¹⁴ We use a conservative approach to labeling observations as censored when testing if the demand for punishment depend on the availability of rewards. We denote any observation that results in zero payoffs to the proposer, and any observation where the responder chose not to change the proposer's payoffs as censored. Using a random-effects model for censored data we regress the truncated demand for punishments on the offer, a dummy for the Stick treatment, and an interaction between the two, and account for censoring as described above. Looking across all ten rounds, or just the last five rounds, we cannot reject the hypothesis that the demands for punishments are the same across treatments. As expected an increase in the proposer's offer significantly decreases punishments.

¹⁵ Counting each responder as an observation we test if the option of rewarding increases average punishments of low offers. Across the last five rounds the p -value of the test is 0.21, and for the last round the p -value is 0.10. Note that in each of the last five rounds the punishment of low offers exceeds that observed in the absence of rewards.

¹⁶ Such choices cannot be revealed in the trust game, where all transfers result in the responder receiving more than half the pie (Berg et al., 1995). Note that the cost-consequence ratio is one-for-three in the trust game.

¹⁷ To determine if the responders' demands for rewards change when punishments are available, we take a conservative approach to labeling observations as censored. We denote any observation that results in zero payoffs to the responder, and any observation where the responder chose not to change the proposer's payoffs as censored. Using a random-effects model we regress the truncated demand for rewards on the offer, a dummy for the Carrot treatment, and an interaction term between the two, and account for censoring as described above. The joint hypothesis that there is no effect from the punishment option is rejected with a p -value of 0.04 over the ten rounds, and 0.03 over the last five rounds. As expected an increase in the proposer's offer significantly increases rewards.

¹⁸ Treating each responder as an observation we test if the average reward is larger in the Carrot. The p -values for the last five rounds and for each of the seven offer ranges are from smallest offer to largest: 0.04, 0.30, 0.07, 0.03, 0.12, 0.17, and 0.05. The insignificant results are generally found for offers where there are fewer observations.

increasing in the cost of the other input. If the ideal is the same across treatments, then excluding one of these substitutable tools simply suggests that rewards should be larger in the Carrot than in the Carrot-Stick treatment. If subjects are limited to only using one of two tools, then they may use the available tool more. In contrast to our experimental evidence, this also implies that the demands for punishments will be larger in the Stick than in the Carrot-Stick treatment. Therefore a simple substitution argument cannot explain the combined results.

To better understand what motivates rewards and punishments we examine the distribution of final payoffs in the Carrot, Stick, and Carrot-Stick treatments. Figure 7 shows the payoffs during the last five rounds in each of the three treatments. The bold line illustrates the set of possible proposals, and the thin lines illustrate the possible changes that the responders can choose when given a particular proposal. The circles illustrate the observations, where the area of the circle shows the number of observations as indicated by the key. Circles on the right edge of the figure indicate observations where the proposer's payoff exceeds \$4.00.

In the Carrot-Stick treatment, 90 percent of responders chose to either punish or reward at least one time, and 60 percent of all choices involved a punishment or reward. A substantial number of responders choose outcomes that simultaneously decrease equality and decrease their personal payoff—areas II and IV of Figure 1. A total of 23 percent of the observations during the last five rounds are in this area, corresponding to more than half the subjects choosing at least one outcome which decreases both their absolute and relative payoffs.

The substantial treatment differences between the Carrot and Carrot-Stick treatments are apparent when comparing the distributions of final payoffs. Relative to the Carrot-Stick treatment, the Carrot treatment shows much larger rewards overall, as well as larger rewards of low offers. Here 35 percent of all responses make relative payoffs more unequal, which corresponds to 70 percent of the subjects making at least one choice which decreases both their absolute and relative payoffs.

Comparing the outcomes of the Stick and the Carrot-Stick treatments reveals that more punishments in the Stick treatment result in a final

outcome where the responder's payoff exceeds that of the proposer's. Thus, responses appear more extreme when only one tool is available.

Based on these payoff distributions what can we conclude about the motive for punishments and rewards? It is clear that responders do not simply aim at equalizing payoffs, otherwise they should not be choosing outcomes that decrease both their absolute and relative payoffs, nor should rewards be larger in the absence of punishments.¹⁹ Rather our results may suggest that the ideals enforced by responders depend on the tools they have available.²⁰ Specifically it appears that proposers in the Carrot-Stick treatment are held to a higher standard than those in the other treatments. In the Carrot-Stick treatment the offers first-order stochastically dominate those of the other treatments. Thus, for any given offer a larger fraction of offers exceed that offer in the Carrot-Stick game than in any of the other games. As a result, the same offer may be seen as less generous in the Carrot-Stick treatment than if it were made in the Carrot or Stick treatment. All else equal, this would result in smaller rewards and larger punishments in the Carrot-Stick treatment. Combined with the substitutability arguments, two opposing factors may be affecting the average demand for punishments. On one hand the absence of rewards may cause a substitution towards the punishment option, and on the other hand the lack of rewards may imply that a given offer is perceived as being more generous, suggesting less punishment. In contrast, both of these effects suggest larger rewards when there is no punishment option. Our findings are consistent with both of these effects.

III. Conclusion

We have examined demands for rewards and punishments and their effects on cooperation.

¹⁹ Examples of equity models are presented in George Loewenstein et al. (1989), Bolton (1991), Fehr and Klaus M. Schmidt (1999), and Bolton and Axel Ockenfels (2000).

²⁰ The reciprocity models by Matthew Rabin (1993) and Martin Dufwenberg and Kirchsteiger (1998) cannot account for these differences in demands. The reason is that they predict that the kindness associated with a particular offer is independent of the availability of rewards or punishments.

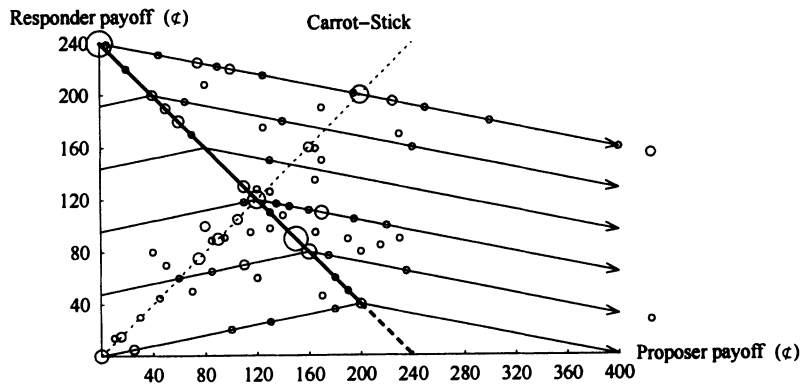
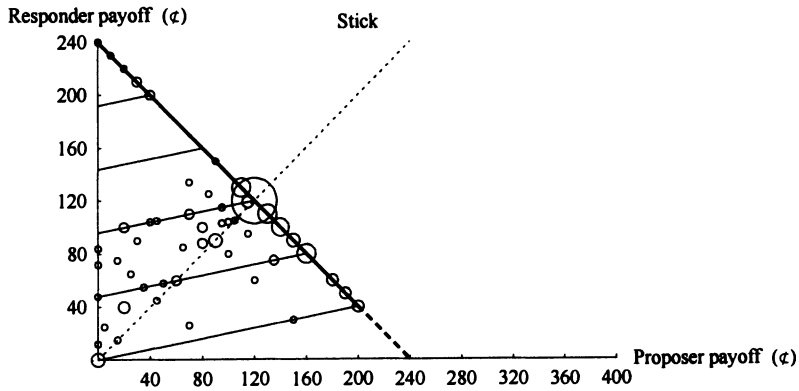
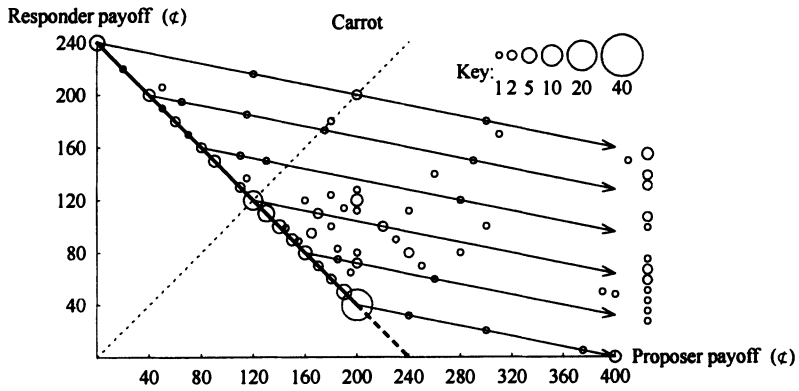


FIGURE 7. FINAL PAYOFF DISTRIBUTION FOR LAST FIVE ROUNDS

We considered a simple proposer-responder environment with randomly rematched partners. In this way, our experiment allowed us to concentrate on the pure demands for rewards and punishments. By considering four conditions—punish or reward, reward only, punish only, and neither—we can identify the effect each has separately and jointly.

We find, first, and perhaps surprisingly, that rewards alone are relatively ineffective in moving the modal offer away from the most selfish one possible. Second, punishments improved cooperation by eliminating extremely selfish offers, pushing proposers in the Stick treatment to modest degrees of cooperation. Combining rewards and punishments had a very strong effect. In the Carrot-Stick treatment the modal offer was the most generous one possible, often leading to rewards by responders. Even though generous offers were not punished, such generosity was only reached when the threat of punishments existed. This indicates that rewards and punishments act to complement one another.

In addition to this, we also found some surprising treatment effects in the responders' choices. While demands for punishments are unaffected by the availability of a reward option, we found that rewards are larger when there is no option of punishing. The combination of these two findings is quite puzzling. An explanation may require a definition of kindness that changes by treatment. Given the distribution of offers it may be that a particular offer is perceived as less kind in the Carrot-Stick treatment than in the other treatments. This suggests that the two tools are not merely substitutes in enforcing a fixed objective, but rather that their availability alters the ideals that they enforce. Thus, compared to punishment-only institutions we may observe harsher conditional punishments when rewards are also available.

Finally, what do our results suggest about how these voluntary demands shape economic institutions? While more work clearly needs to be done, cooperation in our experiment is most successfully enforced in an environment in which both punishments and rewards are available. The process suggested by our data is that the stick can help by getting people to move away from perfect selfishness and to test the waters of cooperation. The carrot can then take over by encouraging further cooperation, ren-

dering the stick a rarely used but necessary tool. Our results show that when devising incentive systems it is important to recognize that in some environments the absence of a reward is not equivalent to a punishment—it is important that both tools be present.

REFERENCES

- Abbink, Klaus; Irlenbusch, Bernd and Renner, Elke.** "The Moonlighting Game: An Experimental Study on Reciprocity and Retribution." *Journal of Economic Behavior and Organization*, June 2000, 2(42), pp. 265–77.
- Andreoni, James; Brown, Paul and Vesterlund, Lise.** "What Makes an Allocation Fair? Some Experimental Evidence." *Games and Economic Behavior*, July 2002, 40(1), pp. 1–24.
- Andreoni, James; Castillo, Marco and Petrie, Ragan.** "What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game." *American Economic Review*, June 2003, 93(3), pp. 672–85.
- Berg, Joyce; Dickhaut, John and McCabe, Kevin.** "Trust, Reciprocity, and Social Norms." *Games and Economic Behavior*, July 1995, 10(1), pp. 122–42.
- Bolton, Gary E.** "A Comparative Model of Bargaining: Theory and Evidence." *American Economic Review*, December 1991, 81(5), pp. 1096–136.
- Bolton, Gary E and Ockenfels, Axel.** "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review*, March 2000, 90(1), pp. 166–93.
- Bolton, Gary E and Zwick, Rami.** "Anonymity versus Punishment in Ultimatum Bargaining." *Games and Economic Behavior*, July 1995, 10(1), pp. 95–121.
- Charness, Gary.** "Attribution and Reciprocity in an Experimental Labor Market." Mimeo, Universitat Pompeu Fabra, 1996.
- Charness, Gary and Haruvy, Ernan.** "Altruism, Equity, and Reciprocity in a Gift-Exchange Experiment: An Encompassing Approach." *Games and Economic Behavior*, August 2002, 40(2), pp. 203–31.
- Charness, Gary; Haruvy, Ernan and Sonsino, Doron.** "Social Distance and Reciprocity: The Internet vs. the Laboratory." Mimeo, Technion, 2000.
- Dickinson, David L.** "The Carrot vs. the Stick in

- Work Team Motivation." *Experimental Economics*, 2001, 4(1), pp. 107–24.
- Dufwenberg, Martin and Kirchsteiger, Georg.** "A Theory of Sequential Reciprocity." CentER Discussion Paper No. 9837, March 1998.
- Fehr, Ernst and Gächter, Simon.** "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, September 2000, 90(4), pp. 980–94.
- Fehr, Ernst; Kirchsteiger, Georg and Riedl, Arno.** "Does Fairness Prevent Market Clearing?" *Quarterly Journal of Economics*, May 1993, 108(2), pp. 437–59.
- Fehr, Ernst and Schmidt, Klaus M.** "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics*, August 1999, 114(3), pp. 817–68.
- Forsythe, Robert; Horowitz, Joel; Savin, Nathan G. and Sefton, Martin.** "Fairness in Simple Bargaining Games." *Games and Economic Behavior*, May 1994, 6(3), pp. 347–69.
- Güth, Werner; Schmittberger, Rolf and Schwarze, Bernd.** "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization*, December 1982, 3(4), pp. 367–88.
- Loewenstein, George; Thompson, Leigh and Bazerman, Max.** "Social Utility and Decision Making in Interpersonal Contexts." *Journal of Personality and Social Psychology*, September 1989, 57(3), pp. 426–41.
- McKelvey, Richard D. and Palfrey, Thomas R.** "An Experimental Study of the Centipede Game." *Econometrica*, July 1992, 60(4), pp. 803–36.
- Offerman, Theo.** "Hurting Hurts More Than Helping Helps." *European Economic Review*, September 2002, 46(8), pp. 1423–37.
- Rabin, Matthew.** "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, December 1993, 83(5), pp. 1281–302.
- Sefton, Martin; Shupp, Robert and Walker, James.** "The Effect of Rewards and Sanctions in the Provision of Public Goods." Mimeo, University of Indiana, 2002.