

## Preplay contracting in the Prisoners' Dilemma

JAMES ANDREONI\*<sup>†</sup> AND HAL VARIAN<sup>‡</sup>

\*Department of Economics, University of Wisconsin, Madison, WI 53706; and <sup>†</sup>School of Information Management and Systems, University of California, Berkeley, CA 94720

Edited by Richard D. McKelvey, California Institute of Technology, Pasadena, CA, and approved July 13, 1999 (received for review May 13, 1999)

**ABSTRACT** We consider a modified Prisoners' Dilemma game in which each agent can offer to pay the other agent to cooperate. The subgame perfect equilibrium of this two-stage game is Pareto efficient. We examine experimentally whether subjects actually manage to achieve this efficient outcome. We find an encouraging level of support for the mechanism, but also find some evidence that subjects' tastes for cooperation and equity may have significant interactions with the incentives provided by the mechanism.

It is easy to write down simple games that have inefficient equilibria. The classic Prisoners' Dilemma is perhaps the simplest game with a dominant strategy equilibrium that is Pareto inefficient. Literally thousands of experiments on the Prisoners' Dilemma have been conducted across the social sciences. [See Rapaport and Chammah (1) and Dawes (2) for reviews of these experiments in sociology and psychology. Roth (3) surveys some of the studies by economists.]

Although researchers often find some fraction of cooperation, the general result in these experiments is that the incentives to defect can be very powerful. When subjects are faced with a single shot of a Prisoners' Dilemma they seldom reach mutually cooperative outcomes. It is interesting to note, however, that subjects do respond to incentives to increase cooperation, despite the existence of a dominant strategy. Roth and Murnighan (4), for instance, found that in a game with an uncertain end point the subjects were more cooperative the longer the anticipated horizon. Selten and Stoecker (5) and Andreoni and Miller (6) look at finitely repeated games and find that subjects will be more cooperative when they can build reputations.

There is a growing number of studies that examine mechanisms for implementing Pareto efficient outcomes in games with inefficient Nash equilibria. Some of the first studies were conducted by Smith (7, 8) on the Smith auction, a mechanism that requires unanimity to implement bids from all players. Although the theoretical properties of this mechanism are not well established, Smith nonetheless finds that his mechanism generates significant amounts of cooperation. [Banks *et al.* (9) generalized Smith's mechanism and find similar results.] Other research into more well-understood mechanisms has generated less optimism. Attiyeh *et al.* (10) examined the dominant strategy Groves-Clarke mechanism and were disappointed to find that fewer than 10% of subjects reveal preferences truthfully. Others have considered Nash equilibrium mechanisms. Bagnoli and McKee (11) examined a provision point game in which private pledges are collected only if a threshold of total pledges is met. They find Pareto efficient outcomes are provided only 54% of the time. The results were even less encouraging when Bagnoli *et al.* (12) turned to multiunit provision point games. Harstad and Marrese (13, 14) examined the Groves-Ledyard mechanism. They found the Pareto efficient outcome was implemented less than 20% of the time. Chen and Plott (15) and Chen and Tang (16) also considered Groves-Ledyard mechanisms and found similarly disappoint-

ing results with a low punishment parameter in the mechanism. However, when they assigned a high punishment parameter the choices converged to the Pareto efficient equilibrium in virtually every session. Chen and Tang (16) also consider the Walker mechanism, again with disappointing results. Chen (personal communication) makes the conjecture that convergence under the high punishment parameter is caused by the fact that only in this case is the mechanism supermodular. Supermodularity provides very robust stability properties that are consistent with learning dynamics. Cheng (18) explicitly considers the dynamic properties of the compensation mechanism and finds that, with a slight modification, it, too, satisfies supermodularity.

All of these prior experiments on mechanism design looked at dominant strategy or Nash equilibrium mechanisms (in ref. 19 Chen provides an excellent review of this literature). In this paper we take another approach and study a mechanism that implements the efficient outcome in a subgame perfect Nash equilibrium. Such subgame perfect mechanisms have gathered interest in the theoretical literature; see Moore and Repullo (20) for an extensive theoretical examination of implementability in subgame perfect equilibria.

The game we look at is modeled after work by Varian (21) regarding preplay contracting. More precisely, we add another stage to the Prisoners' Dilemma game in which the players can make binding commitments to pay the other player some amount if he chooses to cooperate. Of course, it has long been known that the ability to make binding commitments eliminates the dilemma in the Prisoners' Dilemma. However, there has been surprisingly little written about the exact form that such binding commitments might take. A notable exception would be experimental studies of bargaining games. In these games subjects must agree on how to divide a fixed pie, and failure to agree results in the loss of the pie. When subjects alternate making binding offers, as suggested by Rubinstein (22), the subgame perfect equilibrium involves the first player making an offer that is (barely) acceptable to the opponent. Hence, subgame perfection implies efficiency. As discussed by Roth (23), however, such subgame perfect behavior seldom is observed in these games.

In the experimental game we study here, we model the commitment stage explicitly in a two-stage game whose subgame-perfect equilibria are Pareto efficient. These equilibria imply a certain pattern of transfers between the agents that is reminiscent of a competitive equilibrium. In particular, to induce cooperation, each agent must be paid an amount at least as large as the amount he would receive if he were to defect. Roughly speaking, each agent is receiving his opportunity cost for cooperation.

We set a higher bar for ourselves by first training subjects on a basic Prisoners' Dilemma game, without a mechanism. This induces the normal amount of defection and frustration among subjects. Only then do we introduce the mechanism. We find this mechanism to be quite successful. By the end of the

experiment about two-thirds of all subjects chose cooperation, which compares favorably with the experimental results for the other efficiency-inducing mechanisms described above. We see some systematic deviations from our predicted equilibrium, but this behavior is consistent with subjects having tastes for equity similar to those observed in other experiments. We conclude that the mechanism was successful at increasing efficiency, and that the subjects played reasonably close to the subgame perfect equilibrium.

**Pay for Play**

We begin by describing the exact form of our two-stage game, which we call Pay for Play. Imagine a standard Prisoners' Dilemma game with the strategy set (cooperate, defect). We introduce a prior stage to this game where each agent can announce a non-negative number, indicating the amount that he will pay the other agent if he chooses the cooperate strategy in the second stage. This announcement is binding; once the agent offers a contract, he is obligated to carry it out.

Let us calculate the subgame perfect equilibrium of such a game. Consider the following example of an asymmetric Prisoners' Dilemma:

Now we add an announcement stage to this game where each agent simultaneously and independently announces how much

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	6; 7	0; 11
	Defect	9; 0	3; 4

he will pay the other agent if he cooperates. Player 1 announces a side payment of  $s_1$ , and player 2 announces a side payment of  $s_2$ . Given a pair of announcements the game becomes: If player 1 cooperates, player 2 receives 11 by defecting and 7 by cooperating. Therefore the minimal payment that would

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	$6 - s_1 + s_2; 7 + s_1 - s_2$	$0 + s_2; 11 - s_2$
	Defect	$9 - s_1; 0 + s_1$	3; 4

induce him to cooperate is 4. Similarly, the minimum payment to induce player 1 to cooperate when player 2 cooperates is 3. If these payments are announced, that is  $(s_1, s_2) = (4, 3)$ , then the second-stage game is transformed to:

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	5; 8	3; 8
	Defect	5; 4	3; 4

It can be shown  $(s_1, s_2) = (4, 3)$ , followed by mutual cooperation, is the unique subgame perfect equilibrium when the side payments can be any real number. See Varian (21) for a demonstration. Note that Ziss (24) showed that mutual cooperation is not an equilibrium for all constellations of payoffs in the Prisoner's Dilemma game. For the parameters used in this example, which are also those used in the experiment, the

Pareto efficient allocation is the unique subgame perfect equilibrium.

In the experimental game that we consider, the side payments are restricted to be integers. This restriction adds a new subgame perfect equilibrium to the game, namely one where each agent pays 1 unit more than the break-even announcement. In the example we are considering here we would have  $s_1 = 5$  and  $s_2 = 4$  so the game becomes

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	5; 8	4; 7
	Defect	4; 5	3; 4

equilibrium is supported by the pessimistic expectations that if the other player is indifferent between cooperating and defecting, he will choose to defect. The other equilibrium is supported by the optimistic expectations that if the other player is indifferent between his two strategies he will choose to cooperate. Of course, it follows trivially that  $(s_1, s_2) = (4, 4)$  or  $(s_1, s_2) = (5, 3)$ , along with mutual cooperation, are also equilibria.

The idea of adding a contracting stage to the Prisoners' Dilemma is a variation on Varian's (21) compensation mechanism. The idea is that each player offers to compensate the other for the costs that he incurs by making the efficient choice. Varian (21) shows that this sort of compensation mechanism is very powerful. It can be used to internalize nearly any sort of externality, resolve public goods problems, regulate bilateral monopolies, etc. In addition to being robust, the mechanism yields outcomes that are competitive equilibria, with externalities priced at the appropriate efficiency price.

**The Experiment**

Our experiment, which was run on a computer network, is designed and presented to look like a card game. An example of the computer screen that the subjects saw is illustrated in Fig. 1. Instructions used in the experiment are presented in the Appendix.

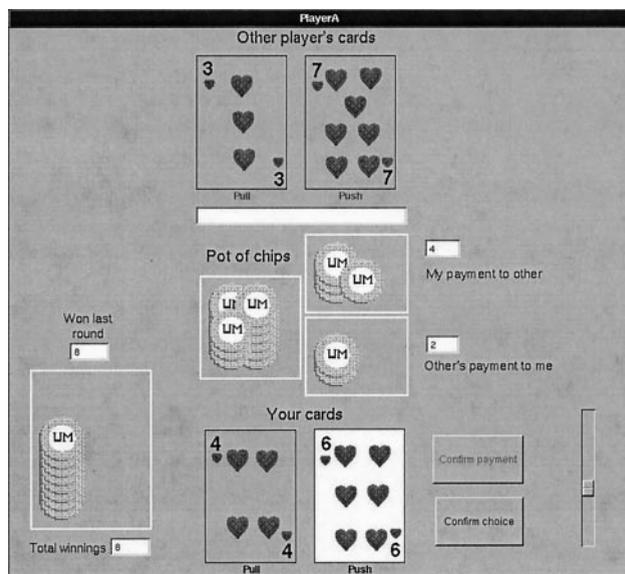


FIG. 1. Sample of the computer screen seen by subjects.

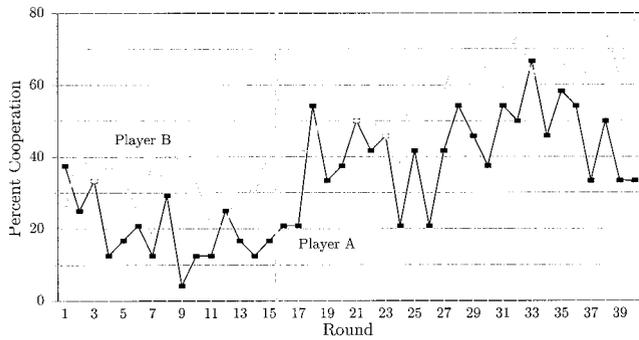


FIG. 2. Cooperation by rounds. Rounds 1–15, Prisoners’ Dilemma; rounds 16–40, with mechanism.

Each game has two players with a pot of chips between them, as shown on the screen in Fig. 1. Each chip is worth 5 cents. Player A has two playing cards, a push card with a 6 on it, and a pull card with a 4 on it. If player A plays his push card, then six chips are pushed from the pot to the other player. If player A plays his pull card, then he pulls three chips from the pot to himself. Player B has a push card of 7 and a pull card of 3 and has the same options. Players must move simultaneously. Note that this game has the same payoff structure as the Prisoners’ Dilemma game illustrated above; the dominant strategy is for each player to choose the pull card, while the Pareto efficient strategy is for each player to choose the push card.

To give the mechanism the greatest challenge, we first ran 15 rounds of this push-pull game with each subject randomly reassigned a different opponent each time. As is commonly observed, the players started out cooperating but soon switched to defecting. By the last round, most players were playing the defect strategy of pull.

We then switched to a new game called Pay for Push. In this game, each player independently names an amount of chips that he will pay the other player if the other player chooses the push card. Each player manipulates a slider on a computer screen to indicate how many chips he would transfer to the other player if the other player chooses to push. Once both players had committed to their payments, the amounts were revealed to the other player and we moved on to the push-pull game described above. We repeated this for 25 rounds. Hence, subjects played a total of 40 rounds, 15 of push-pull and 25 of pay for push.

We conducted six sessions in all, with eight players per session, using a total of 48 subjects. Subjects 1–4 played against subjects 5–8. After each round, players 1–4 were reassigned to a different player in the 5–8 group. Each player had cards (4 and 6 or 3 and 7) and switched cards after each play of the game. Each session of the experiment was complete within an hour, and the subjects earned about \$9–15 in the experiment, plus \$3 for showing up. We recruited some subjects from a subject pool that had been used in a very different bargaining experiment several months earlier, but we do not know if any of our subjects had participated in an experiment before. Other subjects were recruited from economics principles courses. We were careful to recruit subjects before any mention had been made of Prisoners’ Dilemma in the classroom.

**Results**

Fig. 2 shows the fraction of subjects choosing to cooperate over all 40 iterations of the game. Look first at rounds 1–15, which are just the standard Prisoners’ Dilemma game. Over these rounds, 25.8% of moves are cooperative, which declines to 22.9% over rounds 10–15. This finding is consistent with results from other recent Prisoners’ Dilemma games. For example, Roth and Murnighan (4) find 10.1% cooperative

moves, Cooper *et al.* (25) find 20%, and Andreoni and Miller (6) find 18%. A surprising result from this part of the study, however, is that the propensity to cooperate differs significantly across the two types of players. Players in the role of A cooperate 19.2% of the time, while those in the B player role cooperate 32.5% of the time. This difference is statistically significant ( $t = 4.13$ ). The gap remains in the final five rounds of Prisoners’ Dilemma, with As at 16.7% and Bs at 29.2% cooperation ( $t = 2.32$ ). What could explain this asymmetry? A players have a push card of 6 while Bs have a push card of 7. One possibility is that the warm glow from giving may be higher for Bs because the gift to the other player is higher. At the same time, As have a pull card of 4 while Bs have a pull card of 3, so the cost of cooperating is also lower for the Bs. Hence, if there is an independent utility from the act of cooperating, as many have suggested, then we may expect more cooperation from Bs.

Given this training at defecting, how do subjects respond to the introduction of the pay-for-push mechanism? Fig. 2 shows that the response is tentative at first, but after three rounds there is clear movement toward more cooperation. During the entire pay-for-push segment, 50.5% of all moves are cooperative, and 54.5% are cooperative over the last five rounds of play. Again, however, there is a significant difference between As and Bs, with cooperation by As at 41.8% and by Bs at 59.2% ( $t = 6.09$ ).

Clearly the mechanism is having some effect; cooperation has doubled. But the mechanism is also far from 100% successful. There are two reasons that this could be so. First, subjects could be failing to make the subgame perfect side payments in first phase of mechanism, and second, when they receive the appropriate side payment they fail to respond optimally. We examine both of these possibilities next.

**Side Payments: Phase 1 of the Mechanism.** Fig. 3 shows the average side payment made during the experiment for both As and Bs. The subgame perfect prediction is that player A should be offering 3 or 4, and player B should be offering 4 or 5. We see that actual behavior deviates some from this prediction, especially for Bs. Over the 25 rounds, A players’ average offer is 3.39, which rises to 3.57 for the final five rounds, which is consistent with the prediction. For Bs, however, the overall average side payment offered is 3.30 on average, rising to 3.52 over the last five rounds, but is well short of the predicted 4–5.

It is not obvious why side payments would differ in this way. One possibility is that this could be another manifestation of pure tastes for cooperation we saw in the regular Prisoners’ Dilemma game. Alternatively, some subjects could have been using the side payments to try to even out earnings in the mutually cooperative outcome, which could be accomplished if player A’s side payment is one greater than player B’s. In the subgame perfect equilibrium, however, B’s side payment is one greater than A’s. It could be that some A players could have

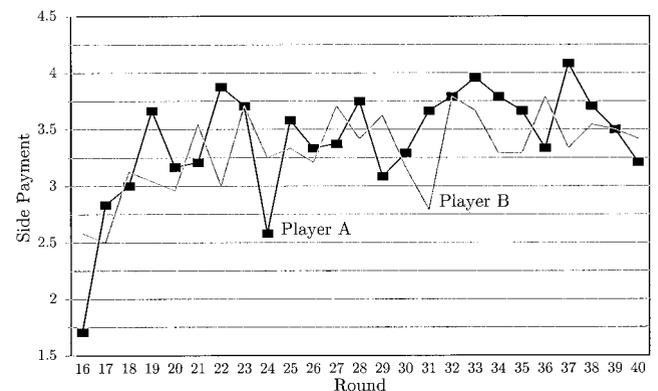


FIG. 3. Average side payments by round.

tried to increase their side payment, and some Bs decrease theirs, in the hopes that tastes for cooperation or equity would be enough to enforce the mutually cooperative outcome, leading to the choices observed.

Let's examine this hypothesis by looking more carefully at the offers made. First, ask how often an offer at or above the equilibrium offer was made. Overall, 63.5% of all side payments were at or above the predicted amounts, rising to 84.9% for the last five rounds. This fraction, however, is higher for the A players than for B players: 78.7% for As and 48.3% for Bs, which is a significant difference ( $t = 8.91$ ). Over the last five rounds, however, the fraction is similar and the difference is insignificant: 85.8% for As and 84.2% for Bs ( $t = 0.36$ ). How often did a player make an offer strictly above the predicted amount, that is, above 3 for As and above 4 for Bs? Overall this happened 39.9% of the time, but again it is far more likely to happen among the As. The difference is 57.5% for As and 22.3% for Bs, which is significant ( $t = 13.32$ ), and it grows to 65.8% versus 12.5% for the last five rounds ( $t = 10.10$ ). Hence, by the end of the experiment players in both roles are equally likely to make an offer that should generate a cooperative reply, but the surplus in the offer is significantly higher in offers made by A players. This finding indeed indicates a bias in play toward more fair allocations.

**Cooperation: Phase 2 of the Mechanism.** Next we ask how subjects respond to the side payments offered in phase 1. Fig. 4 shows how often subjects cooperated when they were offered a side payment at least as high as the subgame perfect Nash equilibrium prediction. The average over the 25 rounds of the mechanism is 67.9%. Hence, when it is optimal to cooperate, two-thirds of subjects do. What about when the offer makes cooperation a dominant strategy? Now the fraction who cooperate when given such an offer rises to 77.4%.

Is there a difference between the A and B players in this regard? Disaggregating this we find that As respond optimally to good offers 63.1% of the time, while Bs do 71.0%. For offers generating a dominant strategy, As cooperate 71.6% of the time and Bs 79.7%. Again, there is a difference in the actions of players in the two roles, albeit a smaller difference than in the offers.

What about the propensity to defect when the side payments offered are so low that defection is the best reply to either of the other player's choices? Given such a bad offer, 79.9% of people will defect. Note this rate of defection is only slightly higher than that in the ordinary Prisoners' Dilemma game, which as we reported earlier was 77.1%.

Again, let's separate the behavior of the A and B players. When As got a bad offer, that is of 3 or less, they defected 78.1% of the time. When Bs get a bad offer, that is of 2 or less, they defected 84.4% of the time. In contrast to the earlier results, these differences are not significant ( $t = 1.59$ ). Hence, any significant difference between As and Bs tends to vanish when they fail to get a good offer in the mechanism.

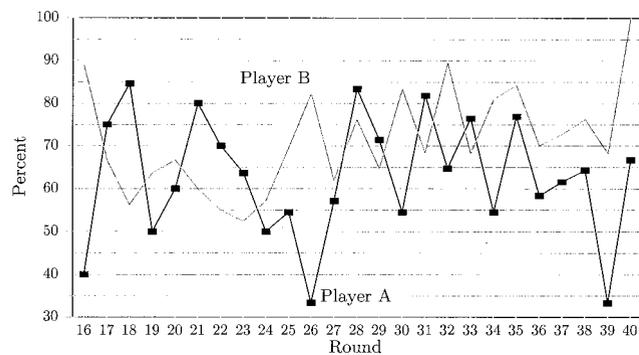


FIG. 4. Probability of cooperation when receiving a good offer.

**Summary: A Closer Look at Equity.** We have seen that, while the mechanism works reasonably well, there is significant asymmetry between the play of As and Bs. In Prisoners' Dilemma without the mechanism, Bs are significantly more cooperative than As. With the mechanism, As make side payments that have a significantly higher surplus. How do we interpret the asymmetry?

To understand the asymmetry of behavior, we should begin by examining the asymmetry in the payoff matrix in the standard Prisoners' Dilemma game. When each player is choosing between his push card and pull card, he is weighing whether to be nice or be selfish. A players can give 6 to the other player at a cost of 4 to themselves. This implies a price of cooperation of  $2/3$ . B players face a similar decision, but they can give 7 at a cost of 3, implying a price of  $3/7$ . Because  $3/7 < 2/3$ , we would predict Bs to be more cooperative, which is what we see.

How would this effect translate to the pay-for-push game? Implicit in the last paragraph is that individuals have a natural desire to be nice in these situations, a hypothesis that has received a fair degree of support in prior experiments. [See refs. 1 and 2 in the broad social science literature. Within economics see, for example, Palfrey and Rosenthal (26) for a discussion of social dilemmas, Andreoni and Miller (6) on Prisoners' Dilemma, Andreoni (27) and Palfrey and Prisbrey (28) on public goods, and Andreoni and Miller (29) for general tastes for giving.]

If there is a utility component in cooperation above and beyond the monetary gain, then it may be that one need not pay the full opportunity cost of cooperation to elicit cooperation in the mechanism. Because Bs face a lower price than As, they already are more inclined to cooperate; hence a side payment below the subgame perfect payments is more likely to induce cooperation for the Bs than for the As. This means that, relative to the subgame perfect side payment, As should get higher payments than Bs. Unfortunately, this is exactly the opposite of what happened.

This means that simple demands for being nice will not be enough to explain this asymmetry. A more complete explanation can be found by looking at final payoffs. If the subgame perfect Nash equilibrium is reached, the payoffs will be fairly unequal; As will earn 8 and Bs 5. Another finding from experiments on games with unequal equilibrium payoffs is that subjects tend to dislike unequal payoffs, especially when on the losing end (see refs. 17 and 30). If subjects do indeed dislike inequality, then this should put pressure on side payments to equal out earnings, not simply to pay the opportunity cost of cooperation. This in turn should suppress side payments by B and increase side payments by A, which is exactly what was observed in the experiment.

It is fair to conclude therefore that the kinds of taste for equity seen in other experiments may be at work here and, as in other settings, is interacting with the incentives of the mechanism. The asymmetry of the payoffs seems particularly powerful in affecting the choices of our subjects.

## Conclusion

We considered a subgame perfect Nash equilibrium mechanism that was applied to the Prisoners' Dilemma. This is perhaps the simplest form of this mechanism that could be implemented in an experiment. We saw that overall the mechanism was largely successful at implementing Pareto efficient allocations. We found that players made offers of side payments that should induce cooperation about 63.5% of the time. When such offers were received, subjects responded with cooperation nearly 70% of the time. Accounting for the usual levels of noise and confusion in experiments, we find these results encouraging for this mechanism.

The presence of concerns for equity also were found in our data, and they appeared to interact with the incentives presented in the mechanism. Especially important was the apparent effort of subjects to use the side payments to undo the inequality implied by the asymmetric payoff matrix, as well as to cover the opportunity costs of cooperation. A symmetric game, while less revealing about the ability of subjects to learn the mechanism, may eliminate some of the concerns with inequality that entered our experiment. Alternatively, one could follow the experiences of Chen and Plott (15) and use a steeper payoff space with a greater gain from cooperation in hopes that this could swamp the concerns for equity and get an even stronger result for the mechanism. These could be interesting topics for future research.

### Appendix: Subjects' Instructions

**Instructions for Zenda.** Zenda is a simple card game that you play with one other person. Each game of Zenda will consist of 15–25 rounds of play. Players are matched up randomly each round, so that you will play a different person each round. All of your choices and earnings in the experiment will be confidential.

In Zenda, you will be playing for chips. The value of the chips corresponds to cash earnings for you. In particular, each chip you earn is worth 5 cents. So if you earn five chips in a round, you earn 25 cents in the round. If you earn 10 chips, you earn 50 cents in that round. The earnings that you make each round will be totaled by the experimenter and paid to you privately and in cash at the end of the experiment. No other subject will know your earnings.

After each round is finished, a dialog box will be displayed informing you of this fact and asking you to wait until the other players are finished. Please click on the OK button as soon as you have read and understood the material because the system will wait until everyone has clicked OK before it proceeds. Dialog boxes will be displayed at other times during the game; after you have read and understood a dialog box, click OK.

Please do not talk to any other player or look at any other player's screen. If you have a problem, please raise your hand and someone will come to help you. We expect the experiment to last about 50 min.

**Information About You.** The first thing that you see will be a panel that asks for information about you. The university requires that we collect this information because we are going to give you money. This information is not recorded as part of the experimental data; it is only there to satisfy university rules about dispersing money. In the experiment you are identified only by a player number, and we maintain strict anonymity. No one in the experiment will ever know your name, your choices, or your earnings.

**Phase 1 of Zenda.** There are two phases to Zenda. Phase 1 is called push-pull. When you play push-pull, you will see two cards in front of you, two cards in front of the other player, and a pile of chips between the two of you. The pile of chips between you is the pot; it is the source of the payments.

Your cards are labeled push and pull. You can choose to play a card by clicking on it with the mouse. When you choose a card it will be highlighted but your choice will not be final until you click the confirm choice button. If you choose the pull card, then you will pull the number of chips on that card from the pot to your pile of chips. If you choose the push card, then you will push the number of chips on that card from the pot to the other player.

Note that the chips that you push or pull come from of the pot in the middle of the table, not from either player's pile of chips. When both players have made their decision to push or to pull, you will see a message appear telling you what has happened and your earnings will be displayed. When all players have made their choices, you will see a panel and a beep

will announce the start of a new round. Click on the OK button on the panel to start playing the new round.

**An Example.** Suppose that your push card is a 6 and your pull card is a 4. Then if you choose to push, the other player will get six chips from the pot in the middle of the table. If you choose to pull, then you will get four chips from the pot.

Suppose that the other person has a push card of 7 and a pull card of 3. Then if he pushes, you will get seven chips from the pot. If he pulls, you will get no chips from the pot.

The total number of chips that you end up with depends on the choices made by you and the choices made by the other player.

**Summary Phase 1 of Zenda.** Choose to push or pull and click confirm choice. If you push, the other player gets the number of chips on your push card; if you pull, you get the number of chips on your pull card. When you see a dialog box that tells you the round has ended, click on OK so that the play can proceed. You will play 15 rounds of phase 1.

**Phase 2 of Zenda.** In this phase you have a new option that we call Pay for Push. Everything else about Zenda will be the same as in phase 1. In particular, during each round of phase 2 you will be randomly paired with another subject each time you play Zenda.

On your screen you will see a new button labeled confirm payment and a slider. You can use the slider to offer a payment to the other player to encourage them to choose to play the push card.

You set this payment by moving the slider up or down. As you do this you will move chips from your pile of chips to a pile in front of the other player. When you are satisfied with your decision about how much you are willing to pay the other player, you click on your confirm payment button. Neither player will see how much the other player is willing to pay until both players have clicked on their confirm payment button. Once you've seen how much the other player has offered to pay you to push, you can decide whether to push or to pull. Your payoff will be as before but now the payment you make will be subtracted from your pile of chips if the other player chooses to push. If the other player chooses to pull, then you will get your payment back. Likewise, if you choose to push, the payment offered to you by the other player will be added to your earnings and subtracted from the other player's earnings.

**Summary of Phase 2 of Zenda.** In each round of play you will move your slider to determine how many chips you want to pay the other player to push. Once you have decided this, you click on your set payment button. When both players have clicked their confirm payment button, each will be able to see how much the other player has offered. At that point, each player can choose whether to push or pull as before. When both players have clicked their confirm payment button, each player sees the payoffs and a new round begins.

You will play 25 rounds of Pay for Push.

**Things to Remember.** You play against a new person every time. You should click OK as soon as you have read and understood a message. You are playing for real money; each chip is worth 5 cents.

We thank Yan Chen for helpful comments. We are grateful to the National Science Foundation for financial support.

1. Rapaport, A. & Chammah, A. M. (1965) *Prisoners' Dilemma* (Univ. of Michigan Press, Ann Arbor).
2. Dawes, R. M. (1980) *Annu. Rev. Psychol.* **31**, 169–193.
3. Roth, A. E. (1988) *Econ. J.* **98**, 974–1031.
4. Roth, A. E. & Murnighan, J. K. (1978) *J. Math. Psychol.* **17**, 189–198.
5. Selten, R. & Stoecker, R. (1986) *J. Econ. Behav. Organ.* **7**, 47–70.
6. Andreoni, J. & Miller, J. H. (1993) *Econ. J.* **104**, 570–585.
7. Smith, V. L. (1979) *Scand. J. Econ.* **81**, 198–215.

8. Smith, V. L. (1980) *Am. Econ. Rev.* **70**, 584–599.
9. Banks, J. S., Plott, C. R. & Porter, D. (1988) *Rev. Econ. Stud.* **60**, 301–322.
10. Attiyeh, G., Franciosi, R. & Isaac, R. M. (1999) *Public Choice*, in press.
11. Bagnoli, M. & McKee, M. (1991) *Econ. Inquiry* **29**, 351–366.
12. Bagnoli, M., Ben-David, S. & McKee, M. (1992) *J. Public Econ.* **47**, 85–106.
13. Harstad, R. M. & Marrese, M. (1981) *J. Econ. Behav. Organ.* **2**, 129–151.
14. Harstad, R. M. & Marrese, M. (1982) *J. Public Econ.* **19**, 367–383.
15. Chen, Y. & Plott, C. R. (1996) *J. Public Econ.* **59**, 335–364.
16. Chen, Y. & Tang, F. (1998) *J. Polit. Econ.* **106**, 633–662.
17. Andreoni, J., Brown, P. M. & Vesterlund, L. (1999) Social Systems Research Institute working paper #9904 (University of Wisconsin, Madison).
18. Cheng, J. Q. (1998) Ph.D dissertation (University of Michigan, Ann Arbor).
19. Chen, Y. (1999) in *Handbook of Experimental Economics Results*, eds. Plott, C. & Smith, V., in press.
20. Moore, J. & Repullo, R. (1988) *Econometrica* **47**, 1191–1220.
21. Varian, H. R. (1994) *Am. Econ. Rev.* **84**, 1278–1293.
22. Rubinstein, A. (1982) *Econometrica* **50**, 97–109.
23. Roth, A. E. (1995) in *The Handbook of Experimental Economics*, eds. Kagel, J. H. & Roth, A. E. (Princeton Univ. Press, Princeton, NJ), pp. 253–348.
24. Ziss, S. (1997) *Am. Econ. Rev.* **87**, 231–235.
25. Cooper, R., DeJong, D. V., Forsythe R. & Ross, T. W. (1996) *Games Econ. Behav.* **12**, 187–218.
26. Palfrey, T. R. & Rosenthal, H. (1988) *J. Public Econ.* **35**, 309–332.
27. Andreoni, J. (1995) *Am. Econ. Rev.* **85**, 891–904.
28. Palfrey, T. R. & Prisbrey, J. E. (1997) *Am. Econ. Rev.* **87**, 829–846.
29. Andreoni, J. & Miller, J. H. (1998) Social Systems Research Institute working paper #9902 (University of Wisconsin, Madison).
30. Prasnikar, V. & Roth, A. E. (1992) *Q. J. Econ.* **108**, 865–888.