The Prisoner's Dilemma as a Two-Level Game: An Experimental Investigation

Stephan Kroll*
St. Lawrence University

John A. List**
University of Maryland

Charles F. Mason***
University of Wyoming

November 1, 2001

Abstract

This paper extends the literature on individual versus group decision-making by combining a strategic inter-group game with an intra-group voting mechanism. The inter-group game is a simple Prisoners' Dilemma, while we consider two versions of an intra-group voting mechanism—a form of representative democracy and a majority vote. Our results suggest that players tend to be more cooperative in the representative democracy treatment than in the individual Prisoner's Dilemma. There are also important behavioral differences between the individual Prisoner's Dilemma and the majority vote treatments—individuals embedded in groups are more inclined to cooperate than subjects making independent decisions.

Classification Code: C92
Keywords: Experiments, Prisoner's Dilemma, two-level games

* Stephan Kroll, Department of Economics, St. Lawrence University, Canton, NY 13617. Phone: (315) 229-5428; fax: (315) 229-5819; e-mail: skroll@stlawu.edu
** John A. List, Department of Agricultural and Resource Economics, University of Maryland. Phone: (301) 405-1288; fax: (301) 314-9091; e-mail: jlist@arec.umd.edu
*** Charles F. Mason, Department of Economics and Finance, University of Wyoming, Laramie, WY 82071-3985. Phone: (307) 766-5336; fax: (307) 766-5090; e-mail: bambuzlr@uwyo.edu (Corresponding author)

Tom Crocker, Eric Nye, Owen Phillips, Jason Shogren, and Tommy Stamland provided valuable comments. Earlier versions were presented at the University of Rhode Island and at the World Congress of Environmental and Resource Economists in Venice, Italy. Funding for the project was received through the University of Central Florida and through the Strook Distinguished Professorship for Natural Resource Conservation and Management at the University of Wyoming.

1. Introduction

Theories in economics and political science typically assume that decisions are made at the individual level. Many times, however, the decision-making unit is a group—a family, legislature, board of directors or even a country—rather than an individual. Such a group can use one of several mechanisms to make its decision. For example, the members can discuss and then vote on what to do; they can vote anonymously and without prior communication, particularly in large groups; or they might appoint a representative to make the decision on the group's behalf (and who will be held accountable by the group afterwards, usually through elections). In each case, the decision-maker is not the unitary rational actor (URA) that most basic theories assume, and there are explicit and implicit social interactions between the group members that affect the group decision.

Despite the pervasiveness of groups as decision-making units in the real world, research by economists and political scientists on whether group decisions differ systematically from decisions made by individuals has been sparse. There are several reasons why such differences could occur: i) group members might have preferences that differ from those of the decision-maker, (Putnam, 1988; Milner, 1998); ii) social psychologists have found that groups often cooperate less and compete more with each other than do individuals; iii) group decisions made by majority rule are more likely to be "correct" than decisions made by an arbitrary single member of the group;[1] and iv) collective choice with voting tends to improve efficiency of group decisions compared to aggregated individual decisions (Walker et al., 2000).

Our paper is the first to combine a strategic inter-group game with an intra-group voting mechanism in an experiment. The inter-group game is a simple Prisoners' Dilemma, while we consider two versions of an intra-group voting mechanism—a form of representative democracy and a majority vote. Our goal in this paper is to compare behavior across two quite disparate settings—one that is consistent with the URA assumption and one that is inconsistent with the URA assumption. To this end, we report results from laboratory experiments that systematically vary the institutional setting of interactions. The first game is a finitely-repeated traditional two-person prisoner's dilemma game, which we refer to as the "prisoner's dilemma", or PD, treatment in what follows. The second game is a "Two-Level Game."[2] In the latter, there are two groups of seven players. In each of the finitely many periods, five voters in each group elect a representative from a set of two candidates. Each representative is involved in games on two levels. On the inter-group level, the representatives elected by the two groups interact in a one-shot prisoner's dilemma game. Representatives also make decisions that are meant to mimic an intra-group or domestic issue in each period. We refer to this treatment as "representative democracy", or RD, in the pursuant discussion. If the URA assumption is valid and groups and individuals behave similarly, one would expect similar behavior in the PD and RD treatments. Persistent and significant differences in behavior, however, would cast doubt on the validity of the URA assumption.

Differences between behavior in the PD and RD treatments could be based on strategic motives or they may arise if individuals behave differently when they are part of a group. To distinguish between these two explanations, we conduct a third treatment wherein two groups of five players interact without a representative making their decision. In this treatment, which we term "direct democracy" (DD) below, each agent votes for one of two actions, and the group's action is determined by majority rule. If individuals behave differently when they are part of a group—as social psychologists have claimed for some time—we would expect to see significant differences between the PD and DD treatments, but not between the DD and RD treatments. If the strategic explanation is accurate, there should be differences between the RD and DD treatments as well as between the PD and RD treatments.

Our results suggest that players tend to be more cooperative in the RD treatment than in the PD treatment. There are also important behavioral differences between the PD and DD treatments—individuals embedded in groups are more inclined to cooperate than subjects making independent decisions. Yet, individuals in the DD treatment become less inclined to vote for cooperation in the later stages of the repeated game. By contrast, we observe a substantial proportion of decision-makers in the representative democracy cooperating even in the final periods of the experiment. This latter observation contrasts with results from individually-played Prisoner's Dilemma games (see, e.g., Andreoni and Miller, 1993). Overall, these findings imply that groups behave more cooperatively and adhere to less individual rationality.

Interesting insights are gained by comparing our results to those in the literature that compares group and individual behavior in experimental games. These earlier papers consider a variety of games, including voting or jury games (Blinder and Morgan, 2000; Ladha et al., 1996)[3], ultimatum games (Bornstein and Yaniv, 1998), dictator games (Cason and Mui, 1997), beauty contest games[4] (Kocher and Sutter, 2000) and auctions (Cox and Hayne, 1998). Most of these designs allow pre-play communication amongst group members (of the papers listed above, only Ladha et al. do not allow pre-play communication). Empirical results from these studies are mixed; but many find that groups make "superior" decisions to individuals. For example, the five-person juries in Blinder and Morgan (2000) and Ladha et al. (1996) tend, on average, to make better decisions than individuals. Similarly, in Bornstein and Yaniv's ultimatum game experiment, proposing groups make on average lower offers to the responding group than do individuals. Since the rejection rate is equally low in group and individual treatments, the authors conclude that groups were also willing to accept less than individuals, which is closer to the outcome associated with the subgame perfect equilibrium. Kocher and Sutter find that teams learn faster than individuals and apply deeper depths of reasoning in later rounds than individuals. On the other hand, Cason and Mui find that two-person teams in the dictator game offer more to their counterparts than individuals (though the difference is modest), while Cox and Hayne find that teams bidding in a common-value auction deviate from rational behavior more than do individuals.

While economists and political scientists have not paid a great deal of attention to possible differences between group and individual decision-making, the research in social psychology on this distinction has been extensive. One of the main findings has been the "discontinuity effect"—a tendency for inter-group interactions to be more competitive and less cooperative than inter-individual interactions.[5]

Discontinuity between inter-group and inter-individual behavior has also been found when there was a dilemma situation within a group itself (in addition to the dilemma between the groups). Even in these situations, in which individual rationality differs not only from collective rationality but also from group rationality, groups are in general more competitive than individuals in one-shot games (Bornstein, 1992; Bornstein and Ben-Yossef, 1994); such behavioral differences tend to diminish over time in repeated games (Bornstein, Winter and Goren, 1996; Insko et al., 2001).[6]

Our experimental results are not consonant with the discontinuity effect, rather we find that groups tended to cooperate more than individuals and were less competitive. It is worthwhile to stress the differences between our experiments and those in social psychology. While Insko, Schopler and their co-authors allow subjects to communicate, there is no communication in our experiments: voting takes place anonymously. The major difference to Bornstein's and his co-authors' set-up is that voters from one group receive the same payoff independently of who, or what they vote for, so that there is no intra-group dilemma.

2. Experimental Design

We start with a simple two-person Prisoner's Dilemma, in which the players simultaneously choose between two actions, "cooperate" and "not cooperate." Table 1 presents the Prisoner's Dilemma game that forms the basis for our experiment. Numbers in each cell represent payoffs in a fictitious currency called tokens, where 40 tokens equaled $1. In each cell, the top right number is the payoff earned by the row player, while the bottom left number is the payoff earned by the column player.

As is well known, the one-shot Nash equilibrium is for each player to select "not cooperate," which results in a payoff of 15 tokens for each. This particular experimental setup is similar to past studies that analyze behavioral nuances of individual players in a Prisoner's Dilemma framework. Empirical results from these experiments overwhelmingly suggest that most subjects fail to play the dominant strategy of not cooperating. This willingness to play cooperatively declines in later periods, especially in experiments where the end period is common knowledge (Andreoni and Miller, 1993).

We conduct three treatments. Our baseline treatment is a finitely repeated game that uses the normal form represented in Table 1 as the stage game. In this game, two players simultaneously choose to either "cooperate" or "not cooperate." We refer to this treatment as the PD treatment below.[7] (A copy of the experimental instructions for this treatment, as well as the two alternative treatments, is contained in the Appendix.) Each repeated game had 25 rounds, a number that was common knowledge to the subjects.[8]

The second treatment, which we refer to as the RD treatment below, uses the same basic framework as the first treatment, but involves two groups of seven players. Two of the seven players from each group are designated as candidates. In each of the 25 periods, one of these two players has been selected by her group to be the representative. The selected representative then plays a prisoner's dilemma game with the other group's representative, with the two players simultaneously choosing to either cooperate or not cooperate. Each representative also makes an intra-group decision. Since the focus of the experiment is on behavior in the inter-group game, to avoid experimental complications we model the domestic issue as a pure public good. To this end, the elected representative in each period receives an endowment of ten tokens, some of which she may allocate to the group. Any tokens allocated to the group are shared in a non-rivalistic manner, while any tokens not allocated to the group are retained by the representative (the opposition candidate does not receive any tokens). One can think of a trade-off between supplying the public good and acquiring rents from an interest group that opposes the supply of the public good. At the same time that the elected representative makes a choice in the (inter-group) prisoner's dilemma game and a token distribution choice (both of which are binding), the subject who was not selected as the representative for the particular round makes non-binding choices for the two dimensions. Every voter in the group is rewarded the tokens from the intergroup prisoner's dilemma game and the portion of the ten tokens distributed by the representative. At the end of the period, each subject in the group learns the choices of both potential representatives and their total payoff for the period, which is the sum of their payment from the PD between the groups and the tokens transferred to them from their leader. Subjects then vote on whether to maintain the incumbent as the leader or replace him/her with the opposition candidate for the next period.[9,10]

In the third treatment, which we refer to as the DD treatment below, two groups comprised of five subjects each play the PD with one another for 25 periods. Each member of each group simultaneously and anonymously chooses between the two actions. As with the RD treatment, there is no preplay communication. Each group's decision is determined via simple majority rule and every member in a group receives the same payoff, independent of his or her vote. For example, if one group voted for cooperation and the other one did not, then all members of the first group would get 5 tokens, and all members of the second group would get 35 tokens.

The data in our experiments are gathered from subjects in two distinct locations: the University of Central Florida (UCF) and the University of Wyoming (UW). Five sessions were conducted at each location, with 58 undergraduate students participating at UCF and 62 undergraduate students participating at UW (see Table 2). In all sessions, subjects did not know who was in their group or against whom they played. Each of the five sessions lasted a little more than one hour, and in each location, subjects earned between $10 and $25 (the RD session lasted on average a few minutes longer, and the subjects in these sessions earned more tokens than subjects in the other treatments due to the domestic game). None of the subjects had previously taken part in a laboratory Prisoners' Dilemma experiment, and only a few had been exposed to any lectures on game theory. In each location we ran one session with the PD treatment and two sessions with each of the RD and DD treatments.

In the UCF sessions, each subject had to fill out a strategy sheet and physically hand it to the monitor in every period. The monitor then calculated the payoffs and gave the strategy sheets back to the subjects. In the UW sessions, subjects made their decisions on linked computers.

3. Results

We present three sets of results. The first set, which we term "unconditional", compares the pattern of behavior across treatment type without controlling for factors that might affect the level of cooperation. The second set, which we term "conditional", extends the unconditional analysis by accounting for factors that could affect cooperation rates. Here, we model individuals' actions using a bivariate panel data model, where we allow for learning and unobservable subject-specific effects. The third set, which we term "trends in play," provides evidence on the ultimate pattern of play. In this analysis we use a Markov chain model, where we focus on the intertemporal transition of the frequency distribution over actions. We then use the information on transitions to determine the steady state frequency distribution, which may be interpreted as the long run pattern of play.

Our general approach to examining the experimental results will be to analyze whether behavioral differences are observed across the three institutions. When behavioral differences are observed, we examine the role that the institution plays in the observed differences.

Unconditional Results

Our first step is to examine whether cooperation rates are similar across institutional settings. Figures 1-4 present the percentage of cooperative play for all three treatments. Figure 1 reveals the typical finding in an individual Prisoner's Dilemma experiment—subjects tend to cooperate more in early stages compared to later trials, particularly when the end period is common knowledge (see, e.g., Andreoni and Miller, 1993). Cooperation rates are above 30% in early rounds, but then decrease until the last few periods, in which subjects completely refrain from cooperative behavior. Alternatively, Figure 2 shows that this dissipation of cooperation rates is not important in the DD treatment—subjects vote for cooperation even in later stages of the treatment. Only in the last four periods does this strategy start to dissipate, as the endpoint phenomenon appears to affect behavior. Even here, willingness to cooperate does not disappear altogether.

Figure 3 shows that the same is true in the DD treatment when we make the group, instead of the individual voter, the unit of observation. Groups continue to vote for cooperation throughout the entire session. We observe similar group behavior in the RD treatment shown in Figure 4. Cooperation rates in these treatments remain relatively stable throughout the 25 periods. And, in some RD sessions, we find cooperative outcomes in unlikely places. For example, in a session at UCF, we observe the leaders voting for cooperation in the very last period.

Although casual inspection of the data is insightful, a more rigorous statistical analysis is necessary to make meaningful inferences. We summarize our data in Table 3, which includes contingency tables that show the frequencies of cooperative outcomes, one-shot Nash outcomes, and other outcomes for each treatment in both locations.[11] To test for behavioral disparities between the three treatments, we use two methods. The first method takes a conservative approach; here we assume that each 25 round "treatment-tuple" yields only one usable observation. Under this approach, we therefore have 12, 4, and 4 PD, DD, and RD observations, respectively. To compute the mean cooperative play for each "observation", we sum the number of cooperative decisions for each treatment and divide by the number of decisions. In a 2-person 25-round PD game, this amounts to summing the 50 decisions (where non-cooperative play = 0 and cooperative play = 1) and dividing by 50. We then compare these computed means using nonparametric methods. We find that mean cooperation rates for the three institutions are 14.2, 27.0, and 28.5 percent for the PD, DD, and RD. We test the null hypothesis of no treatment effect (i.e., that the samples come from identical populations) using a Mann-Whitney rank-sum test, a standard nonparametric test where under the null hypothesis the test statistic is approximately distributed as a standard normal variate. We find that treatment effects are significant at the $p < .05$ level when comparing data across PD and DD ($z = 2.18$), and significant at slightly less than the $p < .05$ level when comparing data across PD and RD (z = 1.94). Treatment effects are insignificant when comparing RD and DD (z = 0.29).

Our second unconditional test assumes observations are generated by three multinomial distributions; using this assumption, we test for equality of the distributions. Our null hypothesis for this test is that the samples are drawn from the same underlying population distribution:

Hypothesis 1: The frequency distribution of outcomes is identical across treatments.

This hypothesis may be examined using a chi-squared test. Under the null hypothesis that all three treatments have the same distribution over outcomes, the test statistic is distributed as a chi-squared variate with four degrees of freedom. Given that our sample size is now considerable larger compared to the nonparametric tests above, we obtain test statistics for three subject cohorts: all subjects, subjects at UW, and subjects at UCF. These test statistics are 28.42 (all subjects), 22.25 (UW subjects), and 10.18 (UCF subjects). All of these test statistics exceed the critical value for a chi-squared variate with 4 degrees of freedom (9.49), and therefore we reject the null hypothesis for each cohort.

Rejection of an independence hypothesis in a three-way classification is difficult to interpret since it is unclear which two distributions are not independent—pairwise tests are needed. Moreover, we are most interested in comparing individual behavior when individuals act in isolation, as compared to when they are part of a group. We therefore consider two additional hypotheses:

Hypothesis 2: The frequency distribution of outcomes in treatment PD is identical to the frequency distribution of outcomes in treatment DD.

Hypothesis 3: The frequency distribution of outcomes in treatment PD is identical to the frequency distribution of outcomes in treatment RD.

The test statistic for both hypotheses is distributed as a chi-squared variate with two degrees of freedom under the null hypothesis. Again, we report test statistics for each of the three cohorts. For the hypothesis that the distribution of actions is identical across the PD and DD treatments, the test statistics are 13.04 (all subjects), 13.29 (UW subjects), and 3.24 (UCF subjects). The first two are statistically significant, though the third is not. For the hypothesis that the distribution of actions is identical across the PD and RD treatments, the test statistics are 22.21 (all subjects), 15.52 (UW subjects), and 9.59 (UCF subjects). All three of these are statistically significant at conventional levels. We conclude that there is evidence that subjects behave differently when placed in groups.[12]

Conditional Results

The above results suggest that individuals behave differently when they are part of a group. While interesting, these results should be considered preliminary since there was no attempt to control for other factors that might affect the decision process. Accounting for such factors would yield a more powerful test of the main hypotheses.

The nature of our data calls for an estimation procedure much more involved than the standard dichotomous regression model. In our case, we gathered data from subjects that provided responses across 25 trials at two different locations. Accordingly, it is important to control for factors that may affect cooperation rates, such as learning and other subject-specific effects that are both observed (location) and unobserved (individual propensity to cooperate). In our econometric approach, we make use of Butler and Moffit's (1982) random effects probit model to control for these various factors while testing our hypothesis of behavior across different voting institutions. The regression model is:

$$nc^*_{it} = \beta`X_{it} + e_{it}, \qquad e_{it} \sim N[0,1], \qquad (1)$$

$$nc_{it} = 1 \text{ if } nc^*_{it} > 0;\ nc_{it} = 0 \text{ if } nc^*_{it} \leq 0 \qquad i = 1,2,\ldots n;\ t = 1,2,\ldots 25;$$

where nc_{it} equals 1 if individual i choose "non-cooperate" in period t, or 0 otherwise. We specify $e_{it} = u_{it} + \alpha_i$, where the two components are independent and both are normally distributed with mean zero. It follows that the variance of the disturbance term e_{it} is $Var(e_{it}) = \sigma_u^2 + \sigma_\alpha^2$. By construction, the individual random effects α_i will capture important heterogeneity across subjects that would be left uncontrolled in a standard cross-sectional model.

Exogenous factors in X_{it} include dichotomous variables that indicate the particular voting regime (DD or RD) of the subject. The response coefficients of these variables provide insights into the differences of cooperation rates across regimes. We also include interactions of these factors with the location at which the subject participated in the experiments. The coefficient estimates of these terms provide insights into whether the differences between cooperation rates (PD versus DD; DD versus RD, PD versus RD) are similar across locations, controlling for individual random effects.

In addition to these primary regressors, we also include other control variables in X_{it}. In the spirit of reinforcement learning models (e.g., Roth and Erev, 1995, 1998, or Sarin and Vahid, 1999), we follow Mason and Nowell (1998) and include a regressor defined as *Diff* to capture the effect of the success rate of previous strategies on current decisions. *Diff* equals the difference of the average payoff from playing cooperation versus the average payoff of playing non-cooperation in all previous periods t-1, t-2, etc. Thus, its response coefficient measures the effect of the success rate of the previous strategies.[13,14] To provide another control for learning, we include a linear time trend to capture behavior that is monotonically trending over trials.

Two aspects of equation (1) warrant further comment. First, we estimate equation (1) using both individual and "group" decision data. Because individuals in the RD treatments do not vote on cooperation, we use only PD and DD data in the individual regression framework. In the "group" regression, we use the DD and RD final group decision in conjunction with the PD decisions to test for differences in "group" decision making across the three voting regimes. Second, we use a random effects model because it treats unmeasured characteristics as error components, economizes on degrees of freedom, and yields coefficients that are not conditioned on unmeasured person and time effects.

We estimate equation (1) using the maximum likelihood approach derived in Butler and Moffit (1982).[15] Empirical results are presented in Table 4. Column 1 in Table 4 contains estimates from the individual-based regression model, whereas column 2 in presents estimates from the "group" decision model. When considering our results, it is important to note that likelihood-ratio tests of the significance of our regression models suggest that both model types are significant at the $p < .01$ level.[16] This outcome implies that our error-components model explains a significant portion of the variation in cooperation rates.

In each model type displayed in Table 4, we include dummy variables for the RD and DD treatments with a constant term. The baseline therefore represents the case where both RD and DD are zero, i.e. the PD treatment; coefficients on the RD and DD dummies reflect differences between the associated treatment and the PD treatment. For example, the coefficient on the DD dummy in the individual decision regression is statistically significantly at the $p < .01$ level, which suggests that subjects were considerably more likely to select the cooperative action in the DD treatment than in the PD treatment. In general, coefficient estimates in both model types corroborate (and strengthen) the unconditional results presented above.

Other results from the individual model suggest that we cannot reject the learning model with a one-sided alternative at the $p < .01$ level. Accordingly, predictions of the learning model—that subjects compare their average past payoffs from playing a certain action with the average payoff of the entire population, or at least with their own average payoffs over all actions—are consonant with our findings. Coefficients of the time trend indicate that players tend to cooperate less in later periods, a result consistent with the intuition of Kreps et al. (1982).

These results are broadly supported in the "group" decision regression model in column 2 of Table 4. Again, the baseline category is the PD treatment, and the coefficients of DD and RD represent deviations from this baseline. Cooperation rates for groups in the DD treatment are significantly greater than cooperation rates in the PD treatments, at the $p < .10$ (respectively, $p < .05$) level for a two-sided (respectively, one-sided) alternative. Likewise, we find that group decisions in the RD treatment tend to be more cooperative than choices in the PD treatment. As in the individual decision regression, we fail to reject the hypothesis that UCF subjects behaved the same as subjects at UW. Overall, these results present evidence that subjects tend to be more cooperatively when they are part of a group decision-making effort as compared to when they act in isolation.

Other control variables in the group model are consistent with results from the individual model. Namely, the variation of learning we incorporate into the regression model performs well, and subjects tend to lower their cooperation rates later in the experiments. Finally, there is no indication of systematic differences in subject behavior between the two locations.

One catalyst for these observed deviations in behavior between behavior in the RD treatment and the other two treatments is the basic premise underlying the theory of two-level games: because governments are obliged to consider group interests, they cannot behave analogously to unitary actors. Electoral concerns have a strong impact on the decisions of the leaders. Figure 4 indicates that cooperation in the RD treatment can be more easily sustained than cooperation in the PD or DD treatments—leaders tend to stick to what has been successful in previous periods. Moreover, inter-group issues are only one arena in which governments make important decisions. By making choices on several issues simultaneously, politicians have the opportunity to compensate potential losers on one issue with better performance on another. In our simple experiments, the second issue is intra-group or domestic in nature—determining the optimal token transfer to the electorate.

To obtain insights into the token transfer issue, we analyzed voting patterns by each electorate. Our focus is on explaining the collective decision to replace an incumbent. To this end, we ran a logit regression, where the dependent variable is coded "1" if the current incumbent is re-elected and "0" if the current challenger is elected. We attempt to explain voting patterns with three explanatory variables: differences between token offers by incumbent and challenger, differences between earnings in the PD game received on the basis of the incumbent's choice and the challenger's announced choice, and differences in the historical tendencies to cooperate (or announce a cooperative action) between incumbent and challenger.[17]

Table 5 reports the results from the logit regression. We observe that payoff differences (from token gifts or PD earnings) are strongly correlated with the decision to retain or replace an incumbent, and have the anticipated sign. Differences in past historic cooperative tendencies also are statistically and economically important, though evidently it becomes more likely for incumbents to be replaced as their tendency to cooperate in the PD game increases. This is somewhat surprising, since we observed a stronger tendency towards cooperation in the RD games than in the PD game. In addition, we note that token gifts apparently play a significantly larger role in influencing voters' decisions: the point estimate on differences in token distribution is on the order of four times the size of the coefficient on differences in PD earnings; this difference is statistically significant. One possible explanation for this divergence is that PD earnings also depend on the behavior of the other team's representative, while the group's representative uniquely determines token distributions. That is, there is some degree of extra risk associated with the PD earnings, and so it is conceivable that voters focus primarily on the token issue.

Trends in play

While our trend variable in the probit random effects models provides an idea concerning subject behavior over time, there is another, more precise technique to address this issue. Given that understanding sustained cooperation is important in any temporal bargaining process, we follow the approach in Friedman (1967) and Mason and Phillips (2000), who treat each vector of frequencies with which an option is chosen as a Markov chain to yield an asymptotic distribution of choice probabilities.

Formally, a transition matrix Π consists of elements π_{jk}, j, k $\in$ {1, 2}, which are probabilities that option j will be selected in round t+1 given that option k was selected in round t. The vector of frequencies at time t+1, x_{t+1}, is related to the vector of frequencies at time t, x_t, in the following manner:

$$x_{t+1} = \Pi x_t, \qquad (2)$$

Asymptotically, play tends to the distribution x*, which satisfies $x^* = \Pi x^*$. This distribution is properly interpreted as the steady state, or long run equilibrium of the system.

To estimate the elements in the transition matrix, we define $p_{jk} = n_{jk}/n_k$, where n_k is the frequency that subjects played option k in period t and n_{jk} is the frequency that subjects who had played option k in period t now play option j in period t+1. With $x^* = (x_1^*, x_2^*)^T$ we can write

$$\begin{pmatrix} x_1 * \\ x_2 * \end{pmatrix} = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} \begin{pmatrix} x_1 * \\ x_2 * \end{pmatrix}. \quad (3)$$

Since $x_1* + x_2* = 1$ and also $p_{1k} + p_{2k} = 1$, we calculate the asymptotic distribution $x^*=(x_1^*, 1- x_1^*)^T$, where $x_1* = p_{12}/(1-p_{11}-p_{12})$.

Because the results reported above indicate the lack of systematic differences in behavior across location, we pool all observations within each of the three treatments. We then consider two samples to produce the set of estimated frequencies: i) all periods pooled and ii) periods 11 through 25 pooled. Our elimination of the first ten periods in the second case is based on the notion that subjects (or subject groups) may experiment in the beginning periods. To the extent this is true, the asymptotic distributions estimated on the basis of observations after period 10 would more accurately reflect the ultimate pattern of play. Our major findings are not significantly changed if we select other periods as the starting point.

The upper panel in Table 6 presents the transition matrix and the asymptotic distributions for all periods of the PD treatment, and the lower panel of Table 6 shows the estimates for the last 15 periods of the same treatment. Note that the limiting probability for playing non-cooperation estimated on the basis of the last 15 periods' observations is larger than the corresponding estimate based on all periods. The same is true in the DD treatment with respect to voting for non-cooperation, individually (Table 7) or at the group level (Table 8). These results corroborate our earlier findings that suggested cooperation rates decrease over time.

We do note an important behavioral difference between group and individual behavior in the DD treatments. The transition probability associated with voting for cooperation in period t+1, given cooperation in t, is larger than 0.5 for individuals in both samples. This finding indicates that subjects who voted for cooperation in the prior round are more likely to cooperate in the current round. By contrast, the corresponding transition probability for groups is much less than 0.5 for both samples. For both individuals and groups, the other diagonal element (the transition probability associated with staying at non-cooperate) is also well above 0.5. This type of inertia is even more evident in the RD treatments (Table 9), where the diagonal elements of the transition matrix are much larger than the off-diagonal elements. Table 9 confirms our previous findings—cooperation in RD is difficult to achieve. For example, the probability of moving from non-cooperate to cooperate is 0.1338 for all periods and 0.1548 for periods 11-25. But, once cooperation is established, the inertia is strong enough to sustain high rates of cooperation: the estimated probability of sustaining cooperation is 0.64 based on data from all periods and 0.6944 based on data from the last 15 periods. A comparison between the DD and RD asymptotic distributions is also interesting. While the asymptotic distribution based on data from all periods is almost identical between DD and RD, there is a gap between these distributions when we consider only the last 15 periods.

These estimated transition matrices and asymptotic distributions indicate three main findings: i) there is less cooperation in the PD treatment compared to the group treatments; ii) in the limit, there is less cooperation in the DD treatment compared to the RD treatment, even though individuals tend to vote at a surprisingly high rate for cooperation in the DD treatment; and iii) cooperation is easier to sustain in the RD treatment compared to the PD and DD treatments.

4. Concluding remarks

Despite the merits of models that use groups such as legislatures or countries as unitary actors, it is important to recognize that they do not address all of the other interactions inherent in representative decision-making. In this experimental study, we used three different experimental treatments—a standard two-player Prisoner's Dilemma, a representative democracy, and a direct democracy—to compare individual and collective behavior when subjects face variations in the voting institution.

We find that the propensity to cooperate increases when subjects move from individual games to group games. We also observe an interesting result in the representative democracy regime that is usually not observed in other Prisoner's Dilemma experiments, namely that cooperation is robust in later periods in the experiment. Several potential conjectures can be made to explain these results: candidates may be able to use a second political issue (transfer payments) to compensate the electorate for potential losses in the Prisoner's Dilemma; the lack of within-group communication may inhibit a strong group identification; or the majority vote might induce group members to choose riskier (more cooperative) behavior since one vote makes up only at most 1/3 of a majority. This last conjecture is consistent with results reported by Walker et al. (2000) who find that subjects facing a commons problem in the lab tend to be more cooperative when they had to vote on a scheme how to use the resource compared to when they made individual decisions.

Overall, our findings suggest that there is a behavioral component inherent in games that group subjects, a component traditional game-theoretic models do not incorporate. More work needs to be done in this area to determine what conditions might induce groups to behave differently than individuals in economic settings. Some possible extensions of our basic approach are to allow more than two options, different forms of communication across and within groups, heterogeneous voters, or incomplete information about voter preferences. Researchers in economics and political science have recognized the necessity to join social psychologists and go beyond the URA assumption when they examine group interactions. Nevertheless, theoretical and empirical studies concerning these issues are still in their infancy. We hope findings from our study will further this research agenda.

Table 1: Payoff Table

		Column player: not cooperate	Column player: cooperate
Row player	not cooperate	15 (row), 15 (column)	35 (row), 5 (column)
	cooperate	5 (row), 35 (column)	25 (row), 25 (column)

Table 2: Experimental Design

Session	Treatment	Number of Subjects	Location
1	PD	10 in 5 dyads	UCF
2	PD	14 in 7 dyads	UW
3	DD	10 in 2 groups	UCF
4	DD	10 in 2 groups	UCF
5	DD	10 in 2 groups	UW
6	DD	10 in 2 groups	UW
7	RD	14 in 2 groups	UCF
8	RD	14 in 2 groups	UCF
9	RD	14 in 2 groups	UW
10	RD	14 in 2 groups	UW

Figure 1: Cooperative Play in PD Treatment

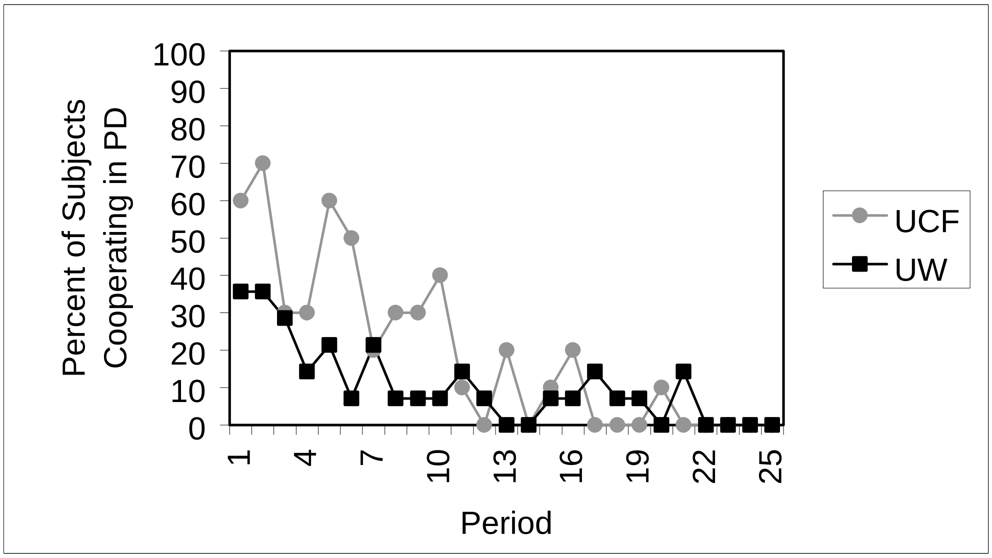

Figure 2: Cooperative Play in DD Treatment

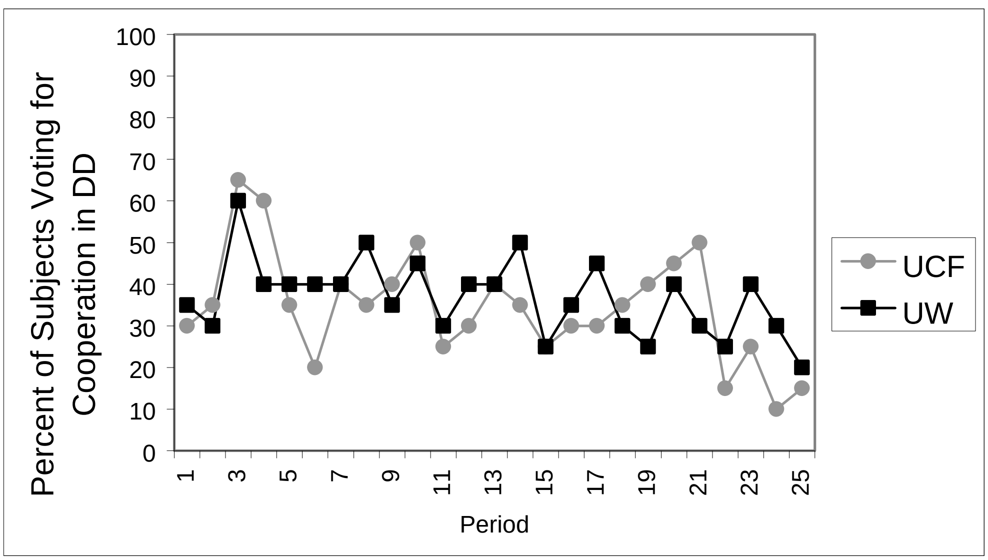

Figure 3: Cooperative Group Play in DD Treatment

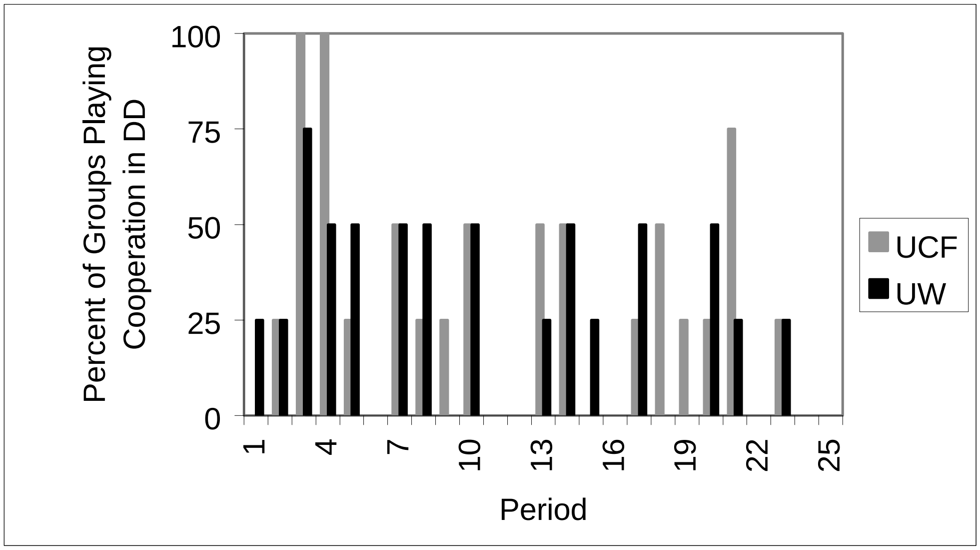

Figure 4: Cooperative Group Play in RD Treatment

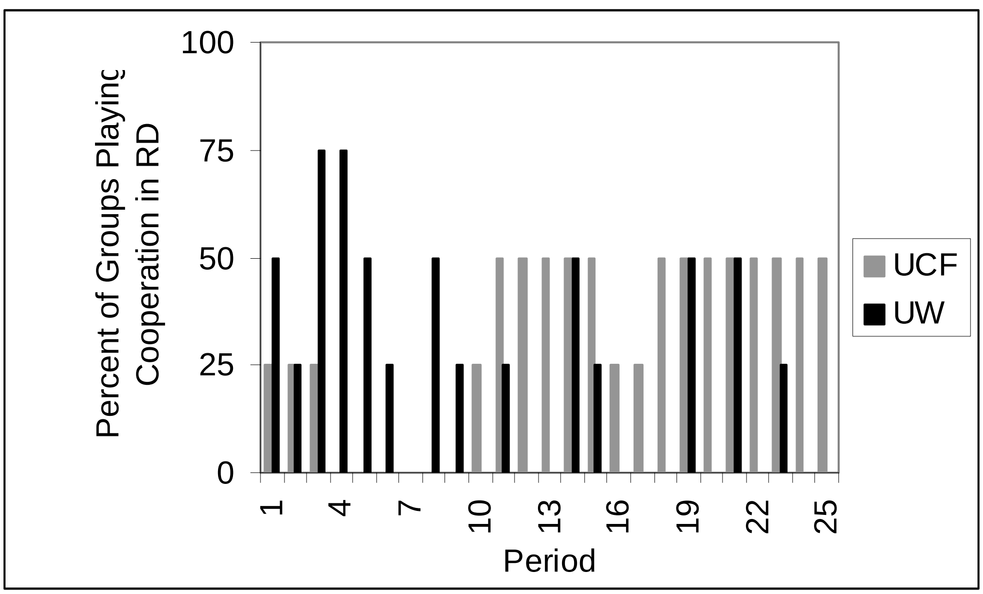

Table 3: Count of Outcomes in Different Treatments

All Sessions	PD	DD	RD
Cooperative outcomes	13	12	19
One-shot Nash outcomes	228	58	62
Neither	59	30	19
Total	300	100	100

UCF Sessions	PD	DD	RD
Cooperative outcomes	11	9	13
One-shot Nash outcomes	88	30	31
Neither	26	11	6
Total	125	50	50

UW Sessions	PD	DD	RD
Cooperative outcomes	2	3	6
One-shot Nash outcomes	140	28	31
Neither	33	19	13
Total	175	50	50

Table 4. Probit random effects estimates

	Model	
Variable	**Individual Decision**	**Group Decision**
Constant	0.57* (2.3)	0.12 (0.80)
Direct Democracy	-0.90** (-3.4)	-0.68* (-1.7)
Direct Democracy UCF*	0.22 (1.0)	0.57 (0.84)
Representative Democracy	–––	-0.43* (-1.9)
*Representative Democracy*UCF*	–––	0.28 (1.2)
Diff.	0.02 (1.6)	0.07** (5.2)
*Diff.*UCF*	-0.007 (-0.46)	-0.02 (-1.2)
Trial	0.04** (9.5)	0.04** (7.7)
χ^2 (d.f.)	280.0 (5)	111.0 (7)
LL	-734.06	-387.76
N	1464	864

Notes:
[a]Dependent variable equals 1 for non-cooperative behavior, 0 for cooperative.
[b]Group regression model uses the group decisions in the direct and representative democracies, and the individual decisions in the prisoners' dilemma.
[c]t-ratios are beneath coefficient estimates.
[d]Regime type prisoners' dilemma is excluded category and represents the baseline group.
[e]χ^2 is a test of model significance.
[f]Each model contains individual specific random effects.
** Significant at the .05 level.
Significant at the .10 level

Table 5. Logit estimates of voting patterns in RD game

Variable		estimate
Constant		1.615***
		(0.2451)
Token distribution difference		0.4602***
		(0.1058)
PD earning difference		0.1087***
		(0.0283)
Cooperative tendency difference		-1.888***
		(0.4595)
χ^2	47.58	
(d.f.)	(3)	
Log-Likelihood	-84.18	
Percent Correctly Predicted:	78.6458	
Madalla's pseudo R-square:	0.2195	
N	192	

Notes:
[a] Dependent variable equals 1 if incumbent re-elected, 0 if challenger elected.
[b] standard errors are beneath coefficient estimates.
[c] χ^2 is a test of model significance.
*** Significant at the .01 level.

Table 6: Transition matrices and asymptotic distribution in the PD Treatment

PD all periods	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.91	0.64	0.88
Cooperation in period t	0.09	0.36	0.12

PD Per. 11-25	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.97	0.70	0.96
Cooperation in period t	0.03	0.30	0.04

Table 7: Transition matrices and asymptotic distribution for individuals in the DD Treatment

DD indivi. all periods	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.77	0.41	0.65
Cooperation in period t	0.23	0.59	0.35

DD indivi. per. 11-25	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.82	0.41	0.70
Cooperation in period t	0.18	0.59	0.30

Table 8: Transition matrices and asymptotic distribution for groups in the DD Treatment

DD group all periods	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.72	0.72	0.72
Cooperation in period t	0.28	0.28	0.28

DD group per. 11-25	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.80	0.82	0.80
Cooperation in period t	0.20	0.18	0.20

Table 9: Transition matrices and asymptotic distribution in the RD-treatment

RD all periods	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.87	0.36	0.73
Cooperation in period t	0.13	0.64	0.27

RD per. 11-25	Non-cooperation in t-1	Cooperation in t-1	Asymptotic distribution
Non-coop. in period t	0.85	0.31	0.66
Cooperation in period t	0.15	0.69	0.34

References

Andreoni, James and John H. Miller, 1993. Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. Economic Journal, 103, 570-585.

Austen-Smith, David and Jeffrey S. Banks, 1996. Information aggregation, rationality and the Condorcet jury theorem. American Political Science Review, 90, 34-45.

Berndt, Ernst, Bronwyn Hall, Robert Hall and Jerry Hausman, 1974. Estimation and inference in nonlinear structural models. Annals of Economic and Social Measurement, 3, 653-65.

Blinder, Alan S. and John Morgan, 2000. Are two heads better than one? An experimental analysis of group versus individual decision-making. NBER Working Paper 7909.

Bornstein, Gary, 1992. The free-rider problem in intergroup conflicts over step-level and continuous public goods. Journal of Personality and Social Psychology, 62, 597-606.

Bornstein, Gary and Meyrav Ben-Yossef, 1994. Cooperation in intergroup and single-group social dilemmas. Journal of Experimental Social Psychology, 30, 52-67.

Bornstein, Gary, Eyal Winter, and Harel Goren, 1996. Experimental study of repeated team-games. European Journal of Political Economy, 12, 629-39.

Bornstein, Gary and Ilan Yaniv, 1998. Individual and group behavior in the ultimatum game: Are groups more 'rational' players? Experimental Economics, 1, 101-108.

Butler, J. S. and Robert Moffitt, 1982. A computationally efficient quadrature procedure for the one factor multinomial probit model. Econometrica, 50, 761-764.

Camerer, Colin and Teck-Hua Ho, 1999. Experience-weighted attraction learning in normal form games. Econometrica, 67, 827-874.

Cason, Timothy N. and Vai-Lam Mui, 1997. A laboratory study of group polarisation in the team dictator game. Economic Journal, 107, 1465-1483.

Cox, James C. and Stephen C. Hayne, 1998. Group vs. individual decision-making in strategic market games. Mimeo.

El-Gamal, Mahmoud A. and David M. Grether, 1995. Are people Bayesian? Uncovering behavioral strategies. Journal of the American Statistical Association, 90, 1137-1145.

Erev, Ido and Alvin E. Roth, 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. American Economic Review, 88, 848-881.

Friedman, James W., 1967. An experimental study of cooperative duopoly. Econometrica, 35, 379-397.

Guarnaschelli, Serena, Richard D. McKelvey, 2000. An experimental study of jury decision rules. American Political Science Review, 94, 407-423.

Iida, Keisuke, 1993. When and how do domestic constraints matter? Two-level games with uncertainty. Journal of Conflict Resolution, 37, 403-426.

———, 1996. Involuntary defection in two-level games. Public Choice, 89, 283-303.

Insko, Chester A., John Schopler, Lowell Gaertner, Tim Wildschut, Robert Kozar, Brad Pinter, Eli J. Finkel, Donna M. Brazil, Candy L. Cecil and Matthew R. Montoya, 2001. Interinvidual-intergroup discontinuity reduction through the anticipation of future interaction. Journal of Personality and Social Psychology, 80, 95-111.

Kocher, Martin G. and Matthias Sutter, 2000. When the 'decision maker' matters: Individual versus team behavior in experimental 'beauty-contest' games. University of Innsbruck, Institute of Public Economics Discussion Paper 2000/4.

Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson, 1982. Rational cooperation in the finitely repeated prisoners' dilemma. Journal of Economic Theory, 27, 245-252.

Ladha, Krishna, Gary Miller and Joe Oppenheimer, 1996. Information aggregation majority rule: Theory and experiments. Mimeo.

Mason, Charles F. and Cliff Nowell, 1998. An experimental analysis of subgame perfect play: The entry deterrence game. Journal of Economic Behavior and Organization, 37, 443-462.

Mason, Charles F. and Owen R. Phillips, 2000. An experimental evaluation of the effects of vertical integration. International Journal of Industrial Organization, 18, 471-496.

Milner, Helen V., 1997. Interests, institutions, and information: Domestic politics and international relations (Princeton University Press, Princeton).

———, 1998. Rationalizing politics: The emerging synthesis of international, American, and comparative politics. International Organization, 52, 759-786.

Milner, Helen V. and B. Peter Rosendorff, 1996. Trade negotiation, information and domestic politics. Economics and Politics, 8, 145-189.

———, 1997. Democratic politics and international trade negotiations: Elections and divided government as constraints on trade liberalization. Journal of Conflict Resolution, 41, 117-146.

Mo, Jongryn, 1994. The logic of two-level games with endogenous domestic coalitions. Journal of Conflict Resolution, 38, 402-422.

———, 1995. Domestic institutions and international bargaining: The role of agent veto in two-level games. American Political Science Review, 89, 914-924.

Nagel, Rosemarie, 1999. A survey of experimental beauty contest games: Bounded rationality and learning. In: David Budescu, Ido Erev and Rami Zwick (eds.), Games and human behavior: Essays in honor of Amnon Rapoport. Lawrence Erlbaum Assoc., New Jersey, pp. 105-142.

Pahre, Robert and Paul A. Papayoanou, 1997. Using game theory to link domestic and international politics. Journal of Conflict Resolution, 41, 4-11.

Putnam, Robert D., 1988. Diplomacy and domestic politics: The logic of two-level games, International Organization, 42, 427-460.

Rabbie, Jacob M., 1998. Is there a discontinuity or a reciprocity effect in cooperation and competition between individuals and groups? European Journal of Social Psychology, 28, 483-507.

Rabbie, Jacob M. and Hein F. M. Lodewijkx, 1994. Conflict and aggression: An individual-group continuum. In: B. Markovsky, K. Heimer and J. O'Brien (eds.), Advances in group processes 11. JAI Press, Greenwich, pp. 139-174.

Roth, Alvin E. and Ido Erev, 1995, Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term, Games and Economic Behavior, 8, 164-212.

Sally, David, 1995. Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992, Rationality and Society, 7, 58-92.

Sarin, Rajiv and Farshid Vahid, 1999. Payoff assessments without probabilities: A simple dynamic model of choice. Games and Economic Behavior, 28, 294-309.

Schopler, John and Chester A. Insko, 1992. The discontinuity effect in interpersonal and intergroup relations: Generality and mediation. European Review of Social Psychology, 3, 121-151.

Selten, Reinhard, 1978. The chain-store paradox. Theory and Decision, 9, 127-159.

Smith, Alastair and David R. Hayes, 1997. The shadow of the polls: Electoral effects on international agreements. International Interactions, 23, 79-108.

Tarar, Ahmer, 2001. International bargaining with two-sided domestic constraints. Journal of Conflict Resolution, 45, 320-340.

Walker, James M., Roy Gardner, Andrew Herr and Elinor Ostrom, 2000. Collective choice in the commons: Experimental results on proposed allocation rules and votes. Economic Journal, 110, 212-234.

Wildschut, Tim, Hein F.M. Lodewijkx, and Chester A. Insko, 2001. Toward a reconciliation of diverging perspectives on interindividual-intergroup discontinuity: The role of procedural interdependence. Journal of Experimental Social Psychology, 37, 273-285.

Williams, Kenneth C., 1994. Sequential elections and retrospective voting: Some laboratory experiments. Journal of Theoretical Politics, 6, 239-255.

Wit, Jörgen, 1998. Rational choice and the Condorcet jury theorem. Games and Economic Behavior, 22, 364-376.

ENDNOTES

1. This is the "Condorcet Jury Theorem." There has been some theoretical controversy about the behavioral assumptions behind this theorem (see, e.g., Austen-Smith and Banks, 1996; and Wit, 1998).

2. The term "two-level game" was coined by Putnam (1988) in an international context, in which the international interactions take place on level 1 and the domestic interactions on level 2. For theoretical research on two-level games see Iida (1993; 1996), Milner (1997), Milner and Rosendorff (1996; 1997), Mo (1994; 1995), Pahre and Papayoanou (1997), Smith and Hayes (1997), and Tarar (2001).

3. Guaranschelli et al. (2000) also conduct an experimental study of jury decision rules, but they focus on whether jury members vote strategically (many do), and whether unanimity rules lead to more incorrect decisions than majority rules (they do not).

4. For a survey of individuals in laboratory beauty-contests see Nagel (1999).

5. In an extensive series of Prisoner's Dilemma experiments, Insko and Schopler, together with several co-authors, have demonstrated the robustness of this effect (see the overview in Schopler and Insko, 1992). The explanations for the discontinuity effect are given by the "identifiability hypothesis" (group membership provides more anonymity), social-support-for-shared-self-interest hypothesis (group members give each other social support for being greedy), and schema-based distrust hypothesis (based on the tendency to mistrust groups more than individuals).

6. Apart from this one-shot versus repeated game inconsistency, some researchers have started to challenge the generality and robustness of the discontinuity effect because of procedural matters (Rabbie, 1998; Rabbie and Lodewijkx, 1994). Even though Wildschut et al. (2001) has tried to reconcile the diverging perspectives on interindividual-intergroup discontinuity, the final word on the discontinuity effect and under what circumstances it appears has not been spoken.

7. An alternative baseline treatment would have two subjects making decisions for two groups with real people. Since in that case any subjects not playing the role of representative would be idle for some time, we opted for the simpler version.

8. At least since Selten's (1978) chain-store paradox it has been well-known that the optimal strategy for rational actors is to not cooperate in each round of a finitely repeated Prisoners' Dilemma when they assume that the other actors are also rational. Kreps et al. (1982) show that the rationality assumption is crucial, and that cooperating can be part of an optimal strategy for a rational subject if he/she thinks others might be irrational. Because we are interested in comparing behavior across treatments rather than within treatments, the matter of a finitely versus an infinitely repeated game is less important.

Moreover, Sally (1995) finds in his meta-analysis of social dilemmas that knowledge of the final period does not have a significant impact on the rate of cooperation while the number of iterations does—cooperation rates and the number of periods are inversely related.

9. This voting game was employed at the end of periods 1, 2,…, 24. There was no voting game at the end of period 25, as the repeated game ended at that point. The decision as to which potential representative to select for period 1 was slightly different. Prior to the start of period 1, both candidates announced their intentions as if they were the representatives for period 1, and a representative was chosen based on these announcements. All announcements were non-binding for all subsequent periods.

10. The voting game we use in the Representative Democracy treatment follows that of Williams (1994), who examines whether voters prefer to use a traditional, purely retrospective voting rule or a retrospective-prospective rule. Using the retrospective voting rule, the electorate will presumably vote for the incumbent when their payoffs in the last period are higher than in the next-to-last period. Using the retrospective-prospective rule, voters compare the last payoff they received from the incumbent with the payoff they expect to receive in the next period from the challenger, where their expectations are shaped by the challenger's non-binding announcement. Williams (1994) finds that in times of stability voters tend to use the retrospective voting rule. Note the distinction between our experiments and Williams' (1994): while these announcements are costly to observe for the voters in Williams' set-up, they are free in ours.

11. "Other outcomes" refer to outcomes where the cooperation/non-cooperation tuple was realized.

12. We note that the hypothesis for the third possible pairwise comparison, that behavior is identically distributed across the DD and RD treatments, cannot be rejected at conventional significance levels. Here the test statistics are 4.18 (all subjects), 2.28 (UW subjects), and 2.21 (UCF subjects).

13. Alternatives to the Roth-Erev learning model are proposed in, for example, El-Gamal and Grether (1995) and Camerer and Ho (1999).

14. Three of the fourteen subjects never cooperated so we do not include them in the analysis. We exclude observations from round 1 since there is no previous round, and so no possibility of learning from past play.

15. The likelihood function can be succinctly written as:

$$L = \prod_i L_i = \int_{-\infty}^{\infty} (2\pi)^{-1/2} \prod_t \exp(-e_{it})^2 \, \phi(g_{it} q_{it}),$$

where $g_{it} = 2nc_{it} - 1$; and $q_{it} = \beta`X_{it} + [\text{corr}(e_{it}, e_{is}) / (1 - \text{corr}(e_{it}, e_{is}))]^{1/2} e_i$. Estimation of this particular model is quite complex, but is amenable to Hermite integration. To estimate the model, we use an eight-point quadrature and use the Berndt et al. (1974) estimator to compute the covariance matrix.

16. The test statistic is 280.0 (respectively, 111.0) for the individual (respectively, group) decision regression. As there are 5 (respectively, 7) degrees of freedom, the 1% critical value is 15.09 (respectively, 18.48).

17. In calculating these differences, we used all information available to voters–both observed actions (by incumbent) and announced choices (by challenger). We also considered a variant of the regression model that included differences in historic average total payoffs between incumbent and challenger, along with the contemporaneous differences between token gifts and PD earnings, as reported in the text. Inclusion of this fourth variable had very little effect, and so we focus on the version discussed in the text.