Equilibrium Cooperation in Two-Stage Games: Experimental Evidence

Douglas D. Davis and Charles A. Holt[*]

July, 1995
A revision was publised in the *International Journal of Game Theory*)

Abstract

This paper reports results of an experiment designed to investigate the nature of cooperation and punishment. Subjects are matched in a series of two-person, two-stage games with a sequential equilibrium that supports first-stage cooperation with a credible threat of subsequent punishment. Participants sometimes used a consistent punish/reward strategy, and when they did, cooperation rates increased dramatically. The results thus contradict "payoff relevance": second-stage behavior can be influenced by first-stage outcomes that have no effect on the payoff structure. Nevertheless, high cooperation rates were often not observed, even with a Pareto undominated "punishment" equilibrium in the second stage.

1. Introduction

Recent theoretical work shows that noncooperative behavior can generate "cooperative" outcomes in the early stages of multi-stage games. This paper reports the results of an experiment designed to assess factors that facilitate such cooperation. The experiment involved a two-stage game. A cooperative outcome in the first stage is an outcome yielding payoffs that Pareto dominate those determined by any noncooperative equilibrium *for the first-stage game*. By definition, cooperation would not be an equilibrium in the absence of the second stage, but such cooperation can occur if (noncooperative) equilibrium strategies incorporate appropriate threats and rewards.

Threat strategies specify that a departure from cooperative behavior in the first stage will

---

[*] Virginia Commonwealth University and University of Virginia, respectively. We received helpful suggestions from Ronald Harstad, Robert Reilly, Roger Sherman, Barry Sopher, from seminar participants at the Universities of Arizona, Pittsburgh and Iowa, and from an anonymous referee. This research was supported by grants from the National Science Foundation (SBR 93-19842 and SBR 93-20044). The data and experiment instructions are available at FTP address fido.econlab.arizona.edu.

be followed by play that alters the outcome in the second stage in a way that punishes the "defector."[1] In theory, the credibility of threats depends on the structure of the final stage. It is well known (e.g., van Damme, 1983) that sequential Nash equilibria can support cooperative outcomes when there are multiple Nash equilibria in the second stage game; the threat alters the outcome following a defection from one second-stage equilibrium to another that is worse for the player who defected.[2] Such threats are credible in the sense that the second-stage punishment outcomes are themselves Nash equilibria for the stage game. Threats that are not credible in this sense, i.e. "non-Nash" threats, are generally ruled out in theoretical work (Selten, 1965).

Other than avoiding the use of non-Nash threats, there seems to be no clear consensus among theorists concerning the game structure necessary to generate cooperative outcomes in initial stages of noncooperative games. Bernheim and Ray (1989) and Farrell and Maskin (1989) argue that threats are only effective if the outcome that the strategies specify to follow a first-stage defection is not Pareto dominated by any equilibrium outcome for the subgame following defection. Otherwise, players would have collective incentives to reconsider the planned behavior after the defection occurs. This requirement that the threat outcome not be Pareto dominated was termed "collective dynamic consistency" by Bernheim and Ray and "renegotiation proofness" by Farrell and Maskin.

Although there is considerable interest in the theoretical analysis of cooperation supported by punishments and rewards, an argument to the contrary can be based on a "payoff-relevance" assumption. This assumption precludes threats and rewards that are contingent on earlier decisions if those decisions have no effect on subsequent payoffs. For example, suppose that no decision in the first stage of a game has any effect on the structure of any subsequent subgame. Under the payoff-relevance assumption, even threats that are renegotiation proof would not support cooperation in the first stage of this game. Payoff relevance is used in any analysis in which threat strategies are possible but are ruled out by assumption, e.g., a static Cournot analysis

---

[1] Here, "defection" indicates a departure from cooperative behavior in the first stage. In contrast, a "deviation" will indicate a departure from an equilibrium strategy. If the equilibrium generates a cooperative outcome in the first stage, then a defection is also a deviation.

[2] Benoit and Krishna (1985) provide a useful characterization of the set of outcomes that can be generated in this manner in finitely repeated games.

of a repeated, quantity-choice game.[3]

The experiment reported below involves a first stage that is similar to a prisoners' dilemma, i.e., with a unique Nash equilibrium that is Pareto dominated by a cooperative outcome. Noncooperative behavior is predominant in such one-stage games with unique equilibria after players gain experience from being paired with a sequence of different partners (Holt, 1985; Cooper, DeJong, Forsythe and Ross, 1991). The first stage is followed by a game with the same players. Second-stage treatments include a non-Nash punishment, a Nash punishment that is Pareto dominated by another equilibrium, and a Nash punishment that is not Pareto dominated by the other second-stage Nash outcomes. The paper is organized as follows. Section 2 presents the basic two-stage game and the configuration of equilibria with cooperative outcomes. Experiment procedures, treatments, and results are discussed in sections 3, 4, and 5 respectively. The final section contains a conclusion.

2. Equilibrium Cooperation in Two-Stage Games

---

[3] Both Maskin and Tirole (1988) and Eaton and Engers (1990) are able to generate a type of cooperation in alternating choice models without the use of threat strategies.

First stage: Player R chooses R1 or R2, Player C chooses C1 or C2.
Second stage: Player R chooses R3 or R4, Player C chooses C3 or C4.

C's Actions

R's Actions	C1	C2	C3	C4
R1	60 60	70 10	|	|
R2	40 0	50 20	|	|
R3	-- --	-- --	40 30	50 50
R4	-- --	-- --	10 60	0 20

Figure 1. A Two-Stage Game (the U10 Design).

The experiment consisted of a series of sessions conducted under variations of the simple two-stage game shown in Figure 1. The payoff for the row player is listed in the lower left corner of each box. The row player chooses between R1 and R2 in the first stage and between R3 and R4 in the second; the column player chooses between C1 and C2 in the first stage and between C3 and C4 in the second. Participants are presented with asymmetric incentives in each stage; this was done to minimize coordination problems. The asymmetries provide clear incentives for one player (column) to be identified as a defector, and for the other (row) to have the opportunity to punish defection.[4]

The R2 and C2 decisions constitute a Nash equilibrium for the first stage in Figure 1. Row, however, has little incentive to choose R2 unless it is very likely that column will choose C2. But C2 is the best response to both R1 and R2, so column is presented with an asymmetric

---

[4] Results of a pilot session indicated that a symmetric game could cause sizable coordination problems. The design in the pilot was composed of a symmetric 2x2 prisoners' dilemma followed by a 3x3 matrix with three Nash equilibria. Of the three Nash outcomes, one had equal payoffs, and the other two had unequal payoffs, which allowed one to player punish the other. The symmetry permits either player to defect in the first-stage and to punish the other's defection. The session involved 10 participants in a series of 20 two-stage matchings against anonymous opponents. A cooperative two-stage outcome was never observed. The subjects chose their first-stage cooperative decisions in only 6% of the matchings (and they never chose them at the same time).

incentive to defect from R1/C1. In the second stage, row has an asymmetric capacity to punish a first-stage defection because column’s payoff is more sensitive to row’s decisions. The second-stage game is a battle-of-the-sexes game for which there are two pure-strategy equilibria, R4/C3 and R3/C4, with payoffs of (60,10) and (50,50) respectively. The unique mixed equilibria in the second stage involves probabilities of .5 for each player and expected payoffs of (40,25), assuming risk neutrality.

As is commonly the case for games with multiple equilibria in a stage game, there are many equilibrium outcomes for the game as a whole. Equilibria of interest are enumerated in the leftmost four columns of Table 1. Consider first the equilibria E4, E5 and E6, listed in the bottom half the table. These equilibria arise because any combination of stage-game equilibria will be an equilibrium for the two-stage game. Equilibria E4, E5, and E6 each involve defection, R2/C2, followed by the use of the Nash strategies for one of the three second-stage equilibria in all contingencies. Therefore, the entries in the "reward" and "punish" columns for equilibria E4-E6 are the same.

Equilibria E1, E2, and E3 in Table 1 involve contingent strategies. Equilibrium E1 is composed of the cooperative outcome R1/C1 in the first stage, a "reward" R3/C4 in the second stage, and a "punishment" R4/C3 to any first-stage deviation. This R4/C3 punishment supports the [R1/C1, R3/C4] equilibrium path because it makes deviation from C1 unprofitable for column. That is, in the event column tries to gain 10 by playing C2 in the first stage, R4/C3 reduces column’s second-stage payoff by 40 (from 50 to 10). Equilibrium E2 also generates the [R1/C1, R3/C4] equilibrium path, but supports it with a threat of mixing with a .5 probability in the event of a deviation. The [R1/C1, R3/C4] outcome generated by both E1 and E2 (with payoffs of 110 each) is the Pareto dominant "cooperative" outcome that is of primary interest. A third equilibrium, E3, also supports the R1/C1 first-stage cooperative play. This cooperative equilibrium, however, does not Pareto dominate each of the noncooperative equilibria E4-E6. The equilibrium path for E3 involves mixing in the second stage, and is supported by a

Table 1. Equilibria, Payoffs, and Associated Punish/Reward Propensity Predictions for the U Designs.

Equilibria				Payoffs (U10 design)		Second-Stage Propensities			
	1st Stage	2nd Stage Reward	2nd Stage Punishment	R	C	R_p (R4|C2)	R_r (R3|C1)	C_{ap} (C3|C2)	C_{ar} (C4|C1)
Equilibria with Threat/Reward Strategies									
E1	R1C1	R3C4	R4C3	110	110	1	1	1	1
E2	R1C1	R3C4	mix	110	110	0.5	1	0.5	1
E3	R1C1	mix	R4C3	100	85	1	0.5	1	0.5
Other Equilibria									
E4	R2C2	R3C4	R3C4	70	100	0	1	0	1
E5	R2C2	R4C3	R4C3	80	60	1	0	1	0
E6	R2C2	mix	mix	60	75	0.5	0.5	0.5	0.5

punishment outcome of R4/C3 in the event of any deviation from R1/C1.[5]

The four columns on the right side of Table 1 present four independent measures of contingent behavior in the second stage, given that row makes the cooperative R1 choice in the first stage. These propensities are of the form (X|Y), i.e., the proportion of times that the second-stage decision X followed the first-stage decision Y:

R_p -- row's propensity to punish, defined as (R4|C2),

R_r -- row's propensity to reward, defined as (R3|C1),

C_{ap} -- column's propensity to accept punishment, defined as (C3|C2), and

---

[5] There are at least two other sequential equilibria for the game in Figure 1. First, the decision pair R1/C2 followed by R4/C3 is a weak equilibrium outcome supported by R3/C4 as a punishment for any first-period deviation. Second, the R1/C2 followed by R4/C3 on the equilibrium path can also supported by mixing with a probability of .5 in the second stage.

C_{ar} -- column's propensity to accept a reward, defined as (C4|C1).[6]
In E1, for example, the proportion of times that row follows the deviation C2 with the punishment R4 is 1, so R_p = 1 in the E1 row of the table.

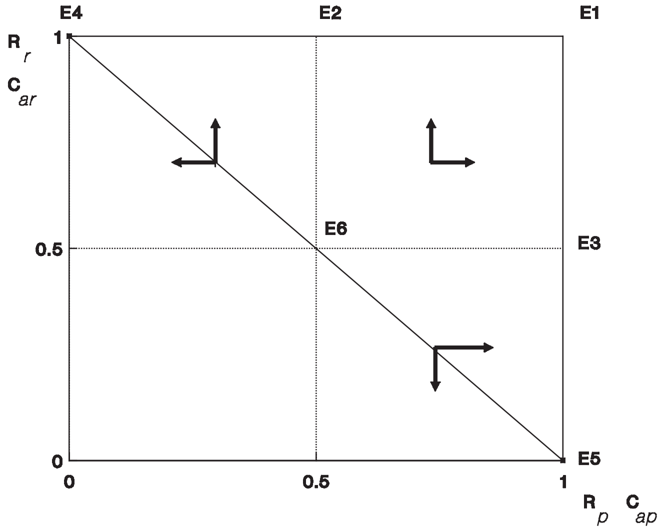

Figure 2. Equilibrium Punish/Reward Propensities for U Design Sessions.

Figure 2 shows the relationship between the second-stage propensities and the equilibria E1-E6. The propensities to punish (R_p), or accept punishment (C_{ap}) are measured on the horizontal axis, while propensities to reward (R_r) or accept reward (C_{ar}) are measured on the vertical axis. The equilibria E1-E6 are plotted as ordered pairs of punish/reward propensities for row (from the R_p and R_r columns of Table 1), and accept-punishment/accept-reward propensities for column (from the C_{ap} and C_{ar} columns of Table 1). Equilibrium E1 appears in the upper right hand corner of Figure 2, because both row and column always engage in pure-strategy punish/reward behavior, so all propensities are equal to 1. Similarly, E2 involves an R3/C4 reward outcome, supported by the mixed equilibrium as a punishment. Consequently, E2 appears

---

[6] It is difficult to interpret actions by Row in terms of punishments and rewards when Row plays R2 in the first stage. For this reason, only outcomes on the E1 equilibrium path are included in these calculations. Thus, R_r is actually calculated as R3|(R1/C1) and R_p is calculated as R4|(R1/C2). Similarly C_{ap} is actually C3|(R1/C2) and C_{ar} is C4|(R1/C1). The simpler notation in the text is used for expositional ease. Second-stage responses to an R2 play are, however, relevant to evaluation of noncontingent equilibria E4 - E6. With the exception of session U10x-3, (discussed below in note 14) the incorporation of Row and Column responses to plays off the equilibrium path has very little effect on results.

at the vertical maximum (indicating that a pure strategy reward always follows a C1 play), and at the midpoint of the horizontal axis (indicating a mixed punishment, with .5 probabilities). The symmetric positioning of E3 at the midpoint of the vertical axis indicates that in this equilibrium the reward is a mix, but the penalty is always R3/C4. Points E4, E5 and E6 represent the three equilibria predicted by the "payoff relevance" assumption. In equilibrium E4, the strategies specify R3 and C4 regardless of first-stage behavior, while in E5 the strategies specify R4 and C3. In E6, each player mixes with probability .5 in the second stage.[7]

The upper right quadrant of Figure 2 has a special significance. On the right half of the figure, $R_p > .5$, and the best response for column is to accept punishment when it is more likely to be given, i.e. to follow C2 with C3. Similarly $C_{ap} > .5$ on the right side of the figure, and the best response of row is to punish, i.e. follow C2 with R4. These two tendencies reinforce each other, so we have drawn arrows to the right in the right half of the figure. An analogous consideration of decisions following C1 indicates that the two propensities on the vertical axis reinforce each other at levels above .5, so upward arrows are drawn in the top half of the figure. The pattern of directional arrows suggests a certain stability for the equilibria E1, E4, and E5, i.e. the equilibria that do not involve any mixing.

3. Experimental Procedures

Data were collected in a series of 23 laboratory sessions conducted at Virginia Commonwealth University and at the University of Virginia. In each session, a group of 10 participants worked through self-paced, interactive instructions at visually isolated personal computers. Following the instructions, participants made decisions anonymously in a series of two-stage games. Decisions were transmitted and recorded via a PC network.

Participants were predominantly undergraduate student volunteers recruited from principles and intermediate economics courses.[8] In total, the decisions of 140 different participants were

---

[7] The equilibria described in note 5, in which [R1/C2, R4/C3] is supported by either a pure R3/C4 punishment, or by mixing in the second stage, cannot be illustrated in Figure 2 without changing the axes. In these equilibria the punishments are contingent on row's rather than on column's first-period action.

[8] Students from upper level undergraduate economics courses participated in some experiments. Also, graduate students very occasionally participated. In one instance two MBA students at VCU participated in a cohort along with

collected in the 23 sessions. The number of participants is less than ten times the number of sessions, because we were interested in assessing the effects of prior experience with the game on performance. Nine of the sessions used "experienced" subjects who were randomly drawn from the pool of participants who had previous experience in a session with two-stage games. Although the experience level was always common across sessions, the experience treatment was not common knowledge. No general announcement regarding experience levels was made during the sessions.

Participants were paid $3.00 for keeping their appointment, plus their earnings in the course of the session. Earnings varied from $8.00 to slightly less than $25.00 for sessions that lasted approximately one hour. Earnings levels were high on an hourly basis, and in general, subjects appeared to be intensely interested in the information provided on their screen displays.

A session consisted of 20 two-stage *games*. In each game, every participant was anonymously matched with another participant. We divided sessions into four five-game *sequences*. At the start of each sequence, participants were given a two-letter "ID" composed of a capital letter indicating their player type (R or C) and a lower-case letter identifier a,b,c,d or e. IDs remained fixed for all participants within each sequence, and it was emphasized to participants in the instructions that they would be paired with a different partner in every game of a sequence. Subjects could verify this rematching in the sense that they were given the list of the IDs for the five subjects with whom they would be matched in a sequence. A participant with an Ra identifier, for example, would see a list of matchings with participants Ca, Cb, Cc, Cd and Ce.

To promote anonymity, lower-case identifier letters and player types were shuffled between sequences. After the first sequence, all R and C assignments were reversed. After the second sequence, only one half of the participants switched roles. This limited reassignment of player-types increased the number of different participant pairings in the session. For example, an R-type player that did not switch player-types after sequence 2 would be paired with some C-type players in the third sequence that had been R-type players in the first sequence. All R

---

eight VCU undergraduate students. In another instance an economics graduate student at UVA participated in a cohort along with nine UVA undergraduate students.

and C roles were again reversed after the third sequence.[9] In total, each participant was a R-type player for 10 games and a C-type player for 10 games.[10]

A record sheet was displayed after each game. The record sheet listed a player's own previous decisions in the sequence, the player's earnings in each game of the sequence, and the decisions of the subjects with whom the player had been paired in each game of the sequence. Participants were not able, however, to see other players' histories or to retrieve the history of their own actions and payoffs in previous sequences. As emphasized in the instructions, all participants of each type faced the same payoff parameters, which did not change during the session, although subjects had no way to know this for more than one sequence at a time. It was announced in advance that payoffs were denominated in pennies, and that participants would be paid in private after the session.

4. Treatments

Our original intention was to focus on the design illustrated in Figure 1, along with some immediate variations. An initial session with the Figure 1 payoffs, however, generated remarkably low rates of cooperation. Column played C1 in only three of the 50 matchings in the final two sequences of this session (denoted U10-1), and the predominant outcome (42 percent of outcomes in the final two sequences) was a Nash outcome in both stages (R2/C2 followed by R4/C3). The low rates of first-period cooperation in this session led us to consider an additional design that we expected to support higher rates of cooperation, while retaining the

---

[9] Thus, participant pairings were anonymous but not random. A random pairing process may rematch the same pair of participants in successive matchings. We used a deterministic matching schedule in an effort to attenuate the possibility that participants might try to build reputations across two-period matchings. Of course, the 10 participants in each 20-matching experiment were occasionally rematched with previous partners. We thought, however, that ID changes across sequences and uncertainty about length of the game would minimize efforts to develop reputations across matchings. The effects of alternative matching schemes are discussed in Cooper, et. al (1991) in the context of a prisoners' dilemma experiment. Participant choices were invariant to a change in their matching procedure that eliminated the possibility of being matched with a previous partner, or with someone who had been matched with a previous partner.

[10] Computer network crashes in sessions U5-2 and U10x-2, along with subsequent reinitializations caused some deviation from this rotation schedule. Although both of these sessions consisted of four 5-game sequences, not all participant's assignments were evenly divided between row and column decision roles.

configuration of equilibria in Figure 1.[11]

Figure 3. The "Nickel" Design with Treatments (U5, D5, N5).

C's Actions

		C1	C2	C3	C4
	R1	60 60	65 0	| |	| |
R's Actions	R2	50 0	55 5	| |	| |
	R3	-- --	-- --	55 55^{U} 35^{D} 55^{N}	60 60
	R4	-- --	-- --	10 65^{U} 45^{D} 45^{N}	0 50

^{U} Pareto Undominated Punishment Treatment
^{D} Pareto Dominated Punishment Treatment
^{N} Non-Nash Punishment Treatment

The alternative parameterization is illustrated in Figure 3. Values in Figure 3 are based on the following reasoning. First, the R1/C1 outcome would be more likely if column (C) gained very little from playing C2, if row (R) gained very little from playing R2, and if R could hurt C very little by playing R2. The first-stage payoffs were chosen to maximize these characteristics, subject to the constraints that the equilibrium structure was unchanged and that differences were no smaller than $.05. In particular, a defection by C from the R1/C1 outcome

---

[11] Low rates of cooperation were also observed in pilot sessions, but these led to alterations in the structure of the game. For example, we conducted two pilot sessions using variations of the U design which led us to provide R and C with asymmetric incentives. (One of these sessions is described in note 4.) The three other pilot sessions we conducted had a 1x2 first stage, (where row has only one first-stage choice, R1, and observes column's actions without recourse) and a 2x2 second stage similar to that observed in Figure 1. The simplicity of this alternative design is appealing. It is deficient, however, in that the two-stage cooperative equilibrium [R1/C1 and R3/C4] does not Pareto dominate a two-period equilibrium combination of R1/C2 in the first stage, and R3/C4 in the second stage. *Ex ante*, we did not expect this to be important, because we expected R to respond to defection with second-period punishment. We modified the design after observing nearly universal play at this "defect/reward" Nash equilibrium in one session.

would only gain a nickel and would cost R sixty cents, thereby making row more likely to punish in the subsequent stage. Note that R would be most likely to punish C (with an R4 play) if R could reduce the earnings of C significantly, at little risk. This objective was satisfied subject to the structure-preserving, nickel difference constraints used in the first stage.

The superscripts on row's second-stage payoffs in Figure 3 indicate different treatments. With the U superscripts, the R4/C3 outcome is a Pareto undominated equilibrium in the second stage, as required by renegotiation proofness. With the D superscripts, the R4/C3 outcome is Pareto dominated by R3/C4, and with the N superscripts there is a unique Nash equilibrium in the second stage at R3/C4, and in this sense R4/C3 is a non-Nash threat.[12] To calibrate results, we also conducted a series of sessions in an "F" treatment, consisting of the first stage structure shown in Figure 1 or Figure 3, and a "flat" second stage in which participants earned $.45 regardless of their choice. The payment level for the second stage was set at $.45 to keep overall earnings at about the level that had been observed in earlier sessions with a U treatment.

The matrix of treatments is summarized by the session identifiers listed in the leftmost column of Table 2. The first letter of each session identifier indicates the structure of the second-stage game: U (Pareto undominated punishment outcome), D (Pareto dominated punishment outcome), N (non-Nash punishment outcome), and F (flat second-stage earnings). Following the treatment letter, the "10" indicates the design in Figure 1 (with a $.10 gain to C for first-stage defection), and the "5" indicates the nickel design in Figure 3. The "x" in a session identifier indicates the use of experienced subjects. Finally, the session number in a treatment cell appears at the end of the session identifier. Thus, for example, U10x-2 refers to the second session conducted in the U10 design that used experienced participants. The 14 sessions with inexperienced participants are grouped at the top of the table, and the 9 sessions with experienced

---

[12] For every treatment with multiple Nash equilibria in the second stage game, there will also be a mixed strategy equilibrium. We were concerned that a mixed-strategy equilibrium close to either of the pure-strategy equilibria in the second stage could bias results in a particular direction. So, in most instances we chose second-stage row and column parameters so that the sums of row's payoffs (column's payoffs) were equal in the second stage. Consequently, the mixed-equilibrium strategies will involve probabilities of .5 for such games. (Obviously, this equal-sums condition is violated for the N treatment with the unique second stage Nash equilibrium.)

Table 2. First and Second-Stage Propensities in the Last Half of Each Session.

Session	First Stage		Second Stage			
	$\frac{R1}{R1+R2}$	$\frac{C1}{C1+C2}$	R_p (R4|C2)	R_r (R3|C1)	C_{ap} (C3|C2)	C_{ar} (C4|C1)
U10-1	.38	.06	.83	1.00	.67	.00
U10-2	.72	.40	.26	.62	.57	.54
U10-3	.70	.22	.29	.71	.50	.86
F10-1	.58	.10	*	*	*	*
F10-2	.60	.32	*	*	*	*
F10-3	.48	.16	*	*	*	*
U5-1	.88	.54	.74	.40	.63	.12
U5-2	.76	.44	.77	.75	.77	.81
F5-1	.90	.62	*	*	*	*
F5-2	.94	.44	*	*	*	*
D5-1	1.00	.74	.46	.78	.60	.62
D5-2	.98	.64	.44	.90	.28	.81
N5-1	.68	.34	.09	1.00	.29	.94
N5-2	.94	.80	.70	.97	.70	.97
U10x-1	.88	.38	.63	.76	.78	.88
U10x-2	.84	.48	.72	.78	.94	.91
U10x-3	.16	.04	.63 (.77[a])	.00 (.00[a])	.50 (.69[a])	.00 (.00[a])
F10x-1	.36	.04	*	*	*	*
F10x-2	.52	.04	*	*	*	*
F10x-3	.04	.12	*	*	*	*
U5x-1	1.00	.68	.75	.85	.81	.97
U5x-2	.98	.80	.48	.85	.80	.92
F5x-1	.62	.34	*	*	*	*

[a] Numbers in parentheses below U10x-3 report second-stage propensities when plays off the E1 equilibrium path are included in the calculation.

participants are at the bottom.[13]

5. Results

Our analysis focuses on the capacity of punishment/reward strategies to support cooperation in the first stage. We consider first the incidence of cooperation in the first stage, and then the extent to which players support cooperation with rewards and punishments in the second stage. These observations are followed by a discussion of the evolution of behavior within sessions.

First-Stage Cooperation

As is discussed below, there is substantial evidence that group behavior is more nearly consistent with equilibrium play after some initial period of adjustment. Therefore, the data summarized in Table 2 are for the last half of each session. The first-stage columns of Table 2, labeled R1/(R1+R2) and C1/(C1+C2), list the proportions of instances in which R and C players chose cooperative decisions. As is clear from the table, the incidence of cooperative behavior varies widely across sessions, with R (for whom defection is not a dominant strategy) tending to cooperate much more consistently than C. Participants in R roles cooperated less than half the time in only five of the 23 sessions (three of those were in the F design). Incentives for C to defect were much more obvious, and C choices are much more dispersed.

Nevertheless, an interesting pattern emerges. First, note the very low incidence of cooperation in some of the U10 sessions, particularly in U10-1 and U10x-3. These cases of low cooperation are surprising, given the fact that the cooperative outcome is supported by a Pareto-undominated punishment. Somewhat higher levels of cooperation were observed in the U5 (nickel design) sessions with inexperienced subjects, which is reasonable since C has a lower

---

[13] The asymmetry of the array of treatments is noteworthy, as it reflects the unanticipated difficulty we encountered in finding a baseline condition where cooperative behavior was clearly supported by contingent play. As mentioned above in the text, low rates of cooperation initially observed in the U10 design led us to explore performance in variants of the "nickel" design. But high rates of cooperation even without the possibility of contingent play in subsequent F5-design sessions led us back to a series of sessions in the U10 and F10 designs. Despite the imbalance of the design, we ceased investigation when the (to us somewhat unexpected) conclusions reported below in section 5 emerged from the analysis with sufficient clarity. Sessions were conducted in the following order: U10-1, U5-1, U5-2, D5-1, D5-2, N5-1, N5-2, F5-1, F5-2, U5x-1, U5x-2, U10x-1, U10x-2, U10-2, U5x-1,F10-1, F10-2, F10-3, U10x-3, F10x-1, F10x-2, F10x-3.

incentive to defect from the R1/C1 outcome in the nickel design. But as is clear from inspection of the results of sessions F5-1 and F5-2, the proportions of C1 choices for U5-1 and U5-2 were no higher on average than the corresponding rates in a treatment with flat second-stage payoffs. The small gain to C1 for defecting in the first stage of the nickel design treatments appears to facilitate cooperation, independent of whether or not cooperation is supported by an equilibrium punishment in the second stage. Moreover, some of the D5 and N5 sessions, summarized in the middle of Table 2, are characterized by very high cooperation rates, despite the fact that the punishment outcome is Pareto dominated or non-Nash. The incidence of C1 actions in D5-1 and D5-2, for example, exceeds that for any of the inexperienced U5 sessions. Similarly, the rate of first-period cooperation by C players in N5-2 equals or exceeds the analogous cooperation rates for any of the U5 sessions, and indeed for any of the other 22 sessions summarized in Table 2. These observations provide the first conclusion, which contradicted our initial expectations.

> ***Conclusion 1:*** *The structure of second-stage equilibria is not the primary determinant of the rate of cooperation in the first stage. Very low rates cooperation rates are sometimes observed when the punish option is a Pareto undominated equilibrium. Very high rates of cooperation are sometimes observed even when the punish option is not a Nash equilibrium.*

Second-Stage Decisions

Recall that second-stage behavior can be described by four independent measures: R's propensity to punish, R_p = (R4|C2); C's propensity to accept punishment, C_{ap} = (C3|C2); R's propensity to reward, R_r = (R3|C1); and C's propensity to accept reward, C_{ar} = (R4|C1).

Observed punish/reward proportions for the last two sequences of each session are shown in the four columns on the right side of Table 2. The proportions are also illustrated in Figures 4, 5 and 6. Our discussion is based on these figures, which have the same structure as Figure 2. Each session is summarized by a "punish/reward" line, which connects a point, $r = (R_p, R_r)$, representing R's punish/reward behavior, and a point, $c = (C_{ap}, C_{ar})$ representing C's acceptance behavior. The line is bolded if participants are experienced. In these figures, evidence suggesting punish/reward behavior in the equilibria E1-E3 is generated as the propensity sums, $R_p + R_r$ and $C_{ap} + C_{ar}$, exceed 1, i.e., as the points $r = (R_p, R_r)$ and $c = (C_{ap}, C_{ar})$ are located further above the diagonal.

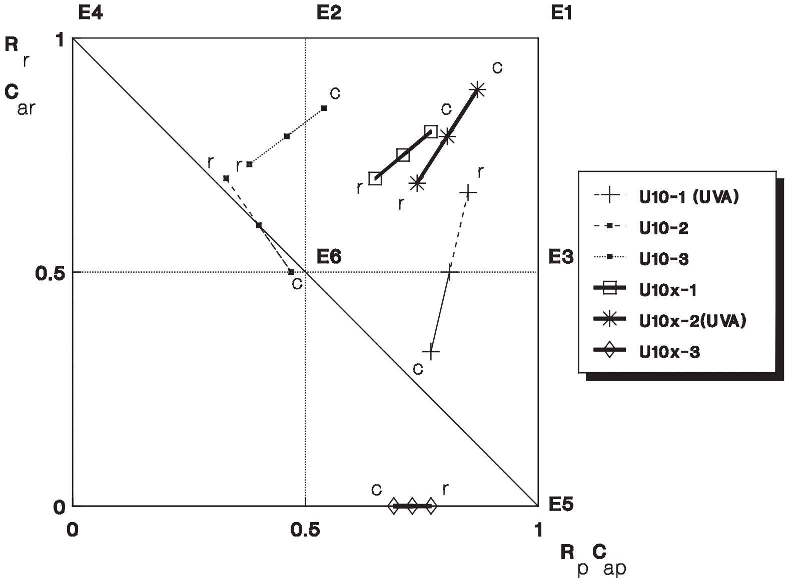

Figure 4. U10 Design Punish and Reward Propensities.

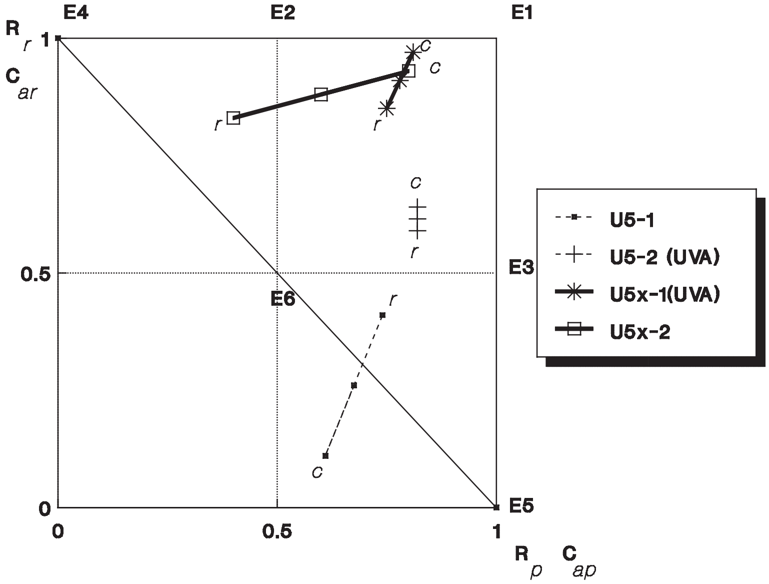

Figure 5. U5 Design Punish and Reward Propensities.

Consider first Figures 4 and 5, which illustrate second-stage behavior in the U10 and U5 design sessions. A punish/reward strategy is not consistently selected, even when the threat point is Pareto undominated. In U10-2, for example, second-stage behavior for both R and C is very nearly random, regardless of first-stage outcomes. In session U10x-3, behavior is closest to equilibrium E5, where second-stage decisions again do not depend on first-stage outcomes.[14]

Nevertheless, Figures 4 and 5 indicate a tendency to use punish/reward strategies; 15 of 20 ordered pairs of points are above the "payoff relevance" diagonal. Moreover punish and reward behavior is quite strong in some sessions. In U5x-1, all components of both ordered pairs exceed .75. Other U-design sessions with high rates of punish/reward play include U10x-2 and U5x-2. Table 2 reveals that these sessions also exhibit relatively high first-period cooperation rates.

High rates of punishment and reward were also observed in some of the D5 and N5 sessions, despite the fact that the punishment outcomes were not Pareto undominated equilibria for the second stage. This is seen in Figure 6, which constructed like Figures 4 and 5, except that the thin dashed lines are the D5 sessions, and the bolded lines are the N5 sessions. Note in particular session N5-2, where the c and r points overlap in Figure 5 at (.70,.97). Despite the non-Nash punishment option, the average rate of punish/reward behavior is as high as for any of the U design sessions.

Consider now the link between punish/reward behavior and first-stage cooperation rates. The 14 sessions with a non-flat second stage can be divided into "high punish/reward" and "low punish/reward" groups. The "high punish/reward" group consists of the seven sessions whose punish/reward line midpoints are closest to E1. Each of these midpoints falls in the upper right quadrant of Figures 4-6, indicating that behavior should tend dynamically to E1. The remaining seven sessions comprise the "low punish/reward" group. The midpoint of the punish/reward line for each of these sessions lies outside of the upper-right quadrant in Figures 4-6 (although the midpoint of this line lies on the upper-right quadrant boundary in U10-1). Cooperation rates in

---

[14] The coordinates for U10x-3 in Figure 4 understate the extent to which participants coordinated on equilibrium E5 in this session. As mentioned in note 6, responses to plays off the E1 equilibrium path (e.g., responses to R2 plays in the first-stage) are relevant to evaluation of noncontingent strategy equilibria. From the first column of Table 2 it is evident that the preponderance of data involved R2 plays. The numbers in parentheses listed below second-stage responses in Table 2 indicate that although R_r and C_{ar} are unchanged by the additional responses (to R2 plays), R_p and C_{ap} increase by .14 and .19. Including this change moves the center point of U10x-3 to (0,.74) in Figure 4.

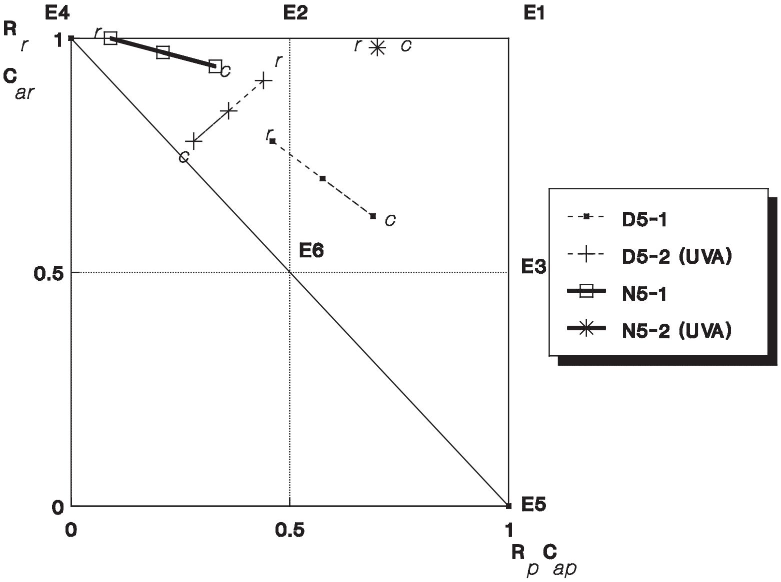

Figure 6. N5 and D5 Design Punish and Reward Propensities.

these groups are compared with those observed in the nine F-design trials, where punish/reward behavior is not possible.

Table 3 summarizes information regarding first-stage cooperation rates, by punish/reward classification. R1 propensities are reported in the upper part of Table 3, while C1 propensities are reported in the lower part of the table. In each part, mean cooperation rates for each group appear in the left-most column. The propensities in the left column indicate that cooperation rates are higher when participants consistently employ a punish/reward strategy. Moreover, the differences are large: the incidence of R1 outcomes in the high punish/reward group, at .91, is nearly 30 percentage points higher than in either the low punish/reward group or the flat 2nd-stage group. Similarly, the incidence of C1 outcomes in the high punish/reward group, at .62, is approximately twice the rate of C1 outcomes in either the low punish/reward group or the flat second stage group. Note also that there is little difference between the cooperation rates in the low punish/reward group and in the flat second-stage design group.

These differences are statistically significant. Entries in the boxes on the right side of Table 4 report Mann-Whitney test statistics for a series of pair-wise tests of the null hypothesis that mean cooperation rates for a group listed as a row entry do not exceed cooperation rates in a corresponding column entry. An asterisk by a test statistic indicates that the null hypothesis

Table 3. First Stage Cooperation Rates and Test Statistics, Grouped by Punish/Reward Behavior.

$\bar{R}$		Mann-Whitney Test Statistics		
		High Punish/Reward	Low Punish/Reward	Flat 2nd Stage
.91	High Punish/Reward	--	35*	32*
64	Low Punish/Reward		--	51
.56	Flat 2nd Stage			--

$\bar{C}$		Mann-Whitney Test Statistics		
		High Punish/Reward	Low Punish/Reward	Flat 2nd Stage
.62	High Punish/Reward	--	35*	32*
.32	Low Punish/Reward		--	51
.24	Flat 2nd Stage			--

* Indicates rejection of the null hypothesis at a 98% confidence level, two-tailed test. For the High Punish/Reward - Low Punish/Reward comparisons, the 98% critical value is 35 (7,7 d.f.). For comparisons involving the Flat Second Stage, the 98% critical value is 38 (7,9 d.f.)

may be rejected at a confidence level at or in excess of 98% (two-tailed tests).[15] The pattern for R players is the same as the pattern for C players: Cooperation rates in the high punish/reward sessions are significantly higher than in either low punish/reward sessions, or in sessions with a flat second stage. Cooperation rates in the low punish/reward sessions are not significantly higher than in sessions with a flat second stage. These results motivate our second conclusion.

---

[15] For a description of the Mann-Whitney test, see e.g., Conover (1980), pp 216-227. The comparisons made in the text withstand the inclusion of session U10-1 into the high punish/reward group, albeit at a somewhat weaker level. Mean R1 and C1 differences decrease to roughly 20 percentage points across high and low punish/reward groups, and the null hypothesis of no increase in cooperation rates in the high punish/reward group over the low punish/reward group may be rejected at only a 90% confidence level (Mann Whitney test statistic T = 32 equals the 90% critical value, 6,8 d.f., two-tailed test.).

Table 4. Average First-Stage Propensities, by Treatment

	Session	First Stage $\frac{R1}{R1+R2}$	 $\frac{C1}{C1+C2}$
Inexperienced Trials	U10-avg.	.60	.23
	F10-avg.	.53	.17
	U5-avg.	.82	.49
	F5-avg.	.92	.53
	D5-avg.	.99	.69
	N5-avg.	.81	.57
Experienced Trials	U10x-avg.	.60	.33
	F10x-avg.	.31	.07
	U5x-avg.	.93	.74
	F5x	.60	.36

> ***Conclusion 2:*** *Participants frequently employ punish/reward strategies, and when they do, first-stage cooperation rates are significantly higher. Thus, second-stage behavior is often associated with factors that are not "payoff relevant."* [16]

Session Dynamics

The marked variability in behavior across sessions represents an important qualification to conclusion 2. Although participants select punish/reward strategies, apparently uncontrolled session-specific characteristics appear to determine when this occurs. To shed some light on the importance of session-specific factors to the incidence of contingent strategy play, this subsection offers some observations pertaining to cohort effects and convergence dynamics. We will begin

---

[16] This conclusion is not inconsistent with results of public goods experiments in which "payoff irrelevant" nonbinding communications between subjects can sometimes affect subsequent cooperation (e.g., Isaac and Walker, 1988, and Palfrey and Rosenthal, 1991).

by comparing behavior in the first and second halves of each session. This division controls for changes in the composition of the decision makers, since in each session half all participants play the game five times as a R player and in five as a C player.[17] At the end of this section we will discuss individual data in more detail, and in particular, the individual responses to their own observed histories of others' play.

(Note, Figures 7 and 8 can be found in the published version.)

Consider the transition of first-stage cooperation rates across session parts, shown in Figures 7 and 8. In each figure, the average of R and C cooperation rates for the second session-half is listed on the horizontal axis.[18] For each session, initial period (·); first half (+); and second half (*) cooperation rates are illustrated as a series of vertically aligned markers. For example, the left-most asterisk in Figure 8 is for the session with the lowest second-half cooperation rate. This session, which incidentally was U10x-3, began with a cooperation rate of about 35% in the initial matching, which decayed to about 25% for the first half of the session and then to 10% for the final half. Trend lines illustrating cooperation rates for initial periods, first session halves, and second session halves are shown as dotted, thin and thick lines, respectively. By construction, the trend line for the second session halves is a diagonal, which serves as a reference. Initial and first-half markers falling above the diagonal indicate that cooperation diminished in the course of the session. Corresponding markers below the diagonal indicate an increase in cooperation as the session progressed.

As is evidenced by the close proximity of the first-half (thin) trend lines to the thick diagonal in each chart, mean cooperation rates across session halves are highly correlated. For the F design sessions, the nonparametric Spearman rank order correlation coefficient between first

---

[17] Smaller divisions of the data, even into session quarters, lose this balance. The role reversals complicate dynamics of session dynamics. Each role-shift interrupts the evolutionary process by injecting a different subgroup of cohort participants with a different history into the respective decision making positions.

[18] Despite substantial differences in R and C cooperation rates within sessions, averaging these measures for comparisons across sessions is innocuous. The statistical comparisons reported in Table 3 generate essentially identical test statistics when the average of R and C cooperation rates are used as the basis of comparison. Moreover, the correlation between R and C cooperation rates is high. The Spearman rank order correlation coefficient of .9 easily allows rejection of the null hypothesis of no correlation (99% critical value is .67, two-tailed test, 14 d.f.). For a description of the Spearman rank order correlation coefficient, and the associated statistical test, see Conover, pp. 250-256.

and second session halves is .97, which easily allows rejection of the null hypothesis that cooperation rates in session halves are uncorrelated (the 99% critical value = .82, two-tailed test, 9 d.f.). Similarly, for the non-F design sessions shown in Figure 8, the Spearman rank order correlation coefficient of .84 also easily allows rejection of the no-correlation null hypothesis (99% critical value = .67, two-tailed test, 14 d.f.). The dispersion of first stage-cooperation rates indicates that cohort effects are substantial. These high correlations imply that such effects are also persistent: Groups that are cooperative in the first half of a session tend to persist in their behavior in the last half of the session. Similarly, groups that are uncooperative in the first half of a session tend to remain so in the latter session half.

Perhaps even more remarkable is that these cooperative propensities appear to be largely established at the time of a cohort's decisions in the first matching. From Figures 7 and 8, this relationship is suggested by comparison of the initial matching (dotted) trend line in each chart with the second half (thick) diagonal. Corresponding Spearman rank order correlation coefficients for the F and non-F sessions are .87 and .65 respectively. Each is sufficiently large to allow rejection of the null hypothesis that cooperation rates for the initial matching and for the second half of each session are uncorrelated at least a 98% confidence level (98% critical values for the F and non-F design sessions are .77 (9 d.f) and .62 (14 d.f.), respectively, two-tailed tests).

Despite the prominence of cohort effects in both the F and non-F design sessions, the adjustment patterns of the sessions exhibit some regularity, as can be seen by comparing trend lines for initial periods and first session halves, with the diagonal trend line for the second session half. In the F design sessions, shown in Figure 7, the dotted initial period trend line is both above, and almost parallel to the diagonal, indicating cooperation consistently decays from the initial decision to the second session half. Comparing the individual initial period dots to second session half cooperation rate stars, it is seen that decay occurred in eight of the nine instances. As suggested by the intercept of the trend line, the decay in cooperation from the initial period to the second session half was about 17 percent on average. In contrast, the trend line for first session half in the F designs lies closer to the diagonal trend line for the second session half, suggesting that there is relatively little decay in cooperation after the first matching.

The high levels of cooperation over the last half of the session, combined with the relative lack of adjustment across session halves, suggests that decay toward the Nash equilibrium remains incomplete in the F design. This result is not terribly surprising in light repeated prisoner's dilemma supergames reported by Selten and Stoecker (1986). There, cooperation decayed only very slowly in a symmetric design. Even slower rates of decay may be expected in the asymmetric design investigated here, since defection is not a dominant strategy for R. To the contrary, in the F10 design R should not defect unless the perceived probability that C will play 2 is at least .83 (50/60). Thus, very infrequent instances of cooperative play by column are sufficient to make continued cooperation by row rational. In the nickel design, the comparable minimum probability rises to .92 (60/65). It is not surprising that the three sessions with the highest second half cooperation rates were the three sessions conducted in the F5 design.

The trend lines shown in Figure 8, suggest a rather different adjustment pattern for the non-F design sessions. Here the (dotted line) initial period, (thin line) first session half and (thick diagonal) second-session half trend lines radiate out as progressively more steeply sloped spokes, from a common hub of about 65% cooperation. The increasing slope of the lines indicates a fairly smooth transition: In sessions with high cooperation rates in the second session half, cooperation rather consistently increased from the first period, to the first session half, to the second session half. Conversely, cooperation rates fairly consistently decreased in sessions with low second-half cooperation rates.

The ordering of the trend lines suggests that, despite the heterogeneity of initial cooperation rates across sessions, there is some regularity in adjustment. The role of contingent strategies in eliciting higher cooperation rates in the non-F design sessions can be more clearly seen by comparing Figure 8 in light of the contingent strategy play transitions, shown in Figure 9. This figure is formatted as Figures 7 and 8, except that the vertical and horizontal axes are denominated in terms of contingent-play frequencies. The evolution of contingent play, in Figure 9, parallels exactly the evolution of first stage cooperation for the non-F design sessions: Initial matching (dotted line) and first session half (thin line) contingent-play trend lines are flatter than the second session half diagonal. This tendency for sessions with high rates of contingent-strategy play in the first half to have higher contingent-play rates in the second half is consistent with the directional arrows drawn in Figure 2.

The progressive growth in the slope of the trend lines in Figures 8 and 9 suggests that changes in cooperation rates within a session are closely tied to changes in the incidence of contingent strategy play. Specifically, as punish/reward behavior diminishes, so do cooperation rates, and as cooperation rates increases, so does the incidence of punish/reward behavior. This relationship is, in fact, very strong. Regressing the change in first stage cooperation on the change in contingent-play rates,

$$\begin{array}{l} \Delta\ 1st\ Stage\ Cooperation \;=\; \underset{(1.00)}{-0.01} \;+\; \underset{(6.35)}{0.59}\ \Delta\ Contingent\text{-}Strategy\ Play \\ n = 14, \quad d.f. = 13, \quad \overline{R^2} = .75 \quad F_{1,\,12} = 40.17 \end{array} \qquad \textbf{(1)}$$

In the equation, *t*-test statistics are printed in parentheses below the estimates. As suggested by the $\overline{R^2}$ of .75 and the highly significant slope term, the growth in contingent strategy play alone explains a very high proportion of the growth in first-stage cooperation rates in the non-F design sessions: Each one percent increase in the frequency of contingent strategy play from the first session half to the last elicits a .59% increase in the incidence of first stage cooperation. The strength of this correlation forms a third conclusion, which is a supplement to conclusion 2.

> ***Conclusion 3:*** *Cohorts differ markedly in their initial propensity to cooperate, and the effects of these initial differences are persistent. Nevertheless, independent of initial cooperative tendencies, the growth in contingent strategy play explains most of the growth in first stage cooperation.*

We close this section with three additional observations regarding the incidence of contingent-strategy play. First, the heterogeneity of cooperation rates among cohorts is not independent of the first-stage structure. As noted in the F design sessions, the three highest second session half cooperation rates were observed in the F5 design. Similarly, in the non-F design sessions, the six highest final second session half cooperation rates occurred in sessions with a nickel design first stage structure. Also, the three lowest second-half cooperation rates in the non-F design occurred in sessions with a dime first stage structure. Using the Mann Whitney test, the null hypothesis than first stage cooperation rates are independent of the first stage structure can be rejected at a 98% confidence level for both the F and non-F design sessions (For the F designs, $T_F = 6$ equals the 98% c.v., two-tailed test, 3,6 d.f. For the non-F sessions T_{NF} = 26 exceeds the 98% c.v. of 28, two-tailed test, 6, 8 d.f.)

That contingent strategies tend to increase cooperation when levels are initially high suggests an ad hoc relationship between *first stage* design structure and the use of contingent strategies: In this environment, contingent strategies are more likely to be employed when the structure facilitates first stage cooperation. Again using the Mann Whitney test, the null hypothesis that the first stage structure does not affect the incidence of contingent-strategy play in the non-F design sessions may be rejected at a 90% confidence level (T_{NF} = 31.5 exceeds the 90% c.v. of 32, two-tailed test, 6, 8 d.f.) *Ex post*, this relationship is quite reasonable, as it is probably much easier to surmount the coordination problem by using punishments and rewards to maintain initially high levels of cooperation, that to do the opposite, that is, to use punishments and rewards to elicit high levels of cooperation.[19] The relationship is likely not stronger because of the increased proportion of experienced sessions in the (generally less cooperative) U10 design sessions relative to the nickel design sessions.

This brings us to a second observation: Some evidence suggests that experience increases both first stage cooperation rates, and the incidence of contingent strategy play.[20] But experience effects do not dominate first-stage structure and other cohort effects. Pooling across all 14 non-F design sessions, neither the null hypothesis that first-stage cooperation rates are unaffected by experience levels, nor the null hypothesis that experience does not affect the incidence of contingent strategy play can be rejected with a nonflat second stage. (Mann-Whitney test statistics for first-stage cooperation and contingent strategy play are T_{coop} = 36, and T_{cont} = 41, respectively. Neither exceed the 80% c.v. of 35, two-tailed test, 6, 8 d.f.)

---

[19] Other design factors may also help improve the selection of contingent strategy play. Particularly intriguing is the high incidence of punish/reward behavior in the N5 design sessions. The combination of a first stage structure that tends to generate high rates of cooperation at the outset of a session, and the unique nature of second stage equilibrium may cause participants to stumble onto a contingent-strategy two-stage path, at the outset of the session. Once there, it may be relatively easy for R participants to communicate an intention to punish deviations by playing R4, even though this action is not part of a sequential equilibrium strategy.

[20] In the first stage, the highest two cooperation rates observed among sessions in the U10 design were with experienced subjects. Similarly, the two U5x sessions generated two of the four highest rates of first-stage cooperation. In the second stage, the two sessions with the highest incidence of contingent-strategy play in the U10 design were sessions with experienced participants, and punish/reward rates in the two U5x sessions were two of the top three observed in the nickel design.

Finally, we note that contingent strategy play is not a consequence of participants stumbling on to a noncontingent pattern of responses that happens to be consistent with equilibrium E1. Rather, participants' evolution toward contingent strategy play comes at the expense of obvious adjustment alternatives. For example, one commonly cited adjustment pattern is that participants play the simple best response to the actions observed in the preceding game (see, e.g., Merlo and Schotter, 1994). Correlating the incidence of best response plays with the incidence of contingent strategy play in the second session halves indicates that the incidence of best response plays diminishes as the incidence of contingent plays increases.[21] The Spearman rank order correlation coefficient of -.56 is easily large enough to reject the null hypothesis of no correlation (95% critical value = .53, two-tailed test, 14 d.f.). A similarly inverse relationship exists between the incidence of "fictitious play," or best responses to the average of preceding observations and contingent strategy play in the second half of sessions. The Spearman rank order correlation coefficient of -.68 is again easily large enough to reject the null hypothesis of no correlation (99% critical value = -.67, two-tailed test, 14 d.f.).

6. Closing Remarks

The most important finding of this study is that punishments and rewards can increase the incidence of cooperation in two-stage games. When participants consistently engage in punish/reward behavior in the second stage, cooperation in the first stage increases dramatically. In contrast, the most commonly discussed theoretic determinants of cooperation do not organize the data well. The structure of second-stage equilibria is not the most important predictor of whether participants use a punish/reward strategy. In some sessions when the punish option is a Pareto-undominated equilibrium, participants did not consistently engage in punish/reward behavior. In other sessions, very high rates of cooperation, supported by consistent punish/reward behavior, were observed even when the punish option is not a Nash equilibrium. The selection

---

[21] Calculating naive best-response decisions is complicated by role switches. In these calculations, we used the most recent decision observed by a person in their current role, even if this decision was observed five matchings previously. In addition, note that best-response and contingent strategy play are not entirely exclusive. Although C's first stage best-response to R is to always play C2, R's best response to C is to play R1 unless C played C2 in the preceding period. In the second stage, best response coordination decisions overlap punish/reward contingent play any time actions in a preceding period are repeated.

of punish/reward strategies appears to be more sensitive to the initial propensity of group members to cooperate, which is prominently affected by the structure of the first stage.

Finally, it is important to emphasize that these results do not preclude the possibility that the structure of punish equilibria may have a greater effect on cooperation rates in other environments. When the game is modified to make the notion of punishing clearer, for example, by allowing pre-play communications or by increasing the number of periods, participants may be much more sensitive to the structure of punishment outcomes.[22]

---

[22] In a different context, Davis and Holt (1994) find that increases in the number of stage games results in an increase in cooperation rates.

References

Benoit, J.P., and V. Krishna, "Finitely Repeated Games," *Econometrica*, July 1985, 53(4), 905-922.

Bernheim, B. Douglas, and Debraj Ray, "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, December 1989, 1(4), 295-326

Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross, "Communication in the Battle of the Sexes Game: Some Experimental Results," *Rand Journal of Economics*, Winter 1989, 20(4), 568-587.

Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross, "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, March 1990, 80(1), 218-233.

Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross, "Cooperation without Reputation," working paper, 1991.

Conover, W.J. *Practical NonParametric Statistics*, 2nd ed. New York: John Wiley and Sons, 1980.

Davis, Douglas D., and Holt, Charles A., "Equilibrium Cooperation in Three-Person, Choice of Partner Games," *Games and Economic Behavior*, 1994, *7*, 39-53.

Eaton, Jonathan and Maxim Engers, "Intertemporal Price Competition," *Econometrica*, May 1990, 58(3), 637-659.

Farrell, Joseph and Eric Maskin, "Renegotiation in Repeated Games," *Games and Economic Behavior*, December 1989, 1(4), 327-360.

Holt, Charles A. "An Experimental Test of the Consistent-Conjectures Hypothesis," *American Economic Review*, June 1985, 75(3), 314-325.

Isaac, R. Mark and James M. Walker, "Communication and Free-Riding Behavior: The Voluntary Contribution Mechanism," *Economic Inquiry*, October 1988, 26(4), 585-608.

Maskin, Eric, and Jean Tirole, "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves and Edgeworth Cycles," *Econometrica*, May 1988, 56(3), 571-599.

Merlo, Antonio, and Andrew Schotter, "An Experimental Study of Learning in One and Two-Person Games," Working Paper, New York University, 1994.

Nagel, Rosemarie, "Experimental Results on Interactive Competitive Guessing" working paper, University of Bonn., 1993.

Palfrey, Thomas, and Howard Rosenthal, "Testing for Effects of Cheap Talk in a Public Goods Game with Private Information," *Games and Economic Behavior*, May 1991, 3(2), 183-220.

Selten, Reinhard "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragetragheit," parts I - II, *Zeitschrift für die Gesamte Staatswissenschaft*, 1965, 121, 301-324 and 667-689.

Selten Reinhard and Rolf Stoecker, "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames, A Learning Theory Approach," *Journal of Economic Behavior and Organization*, 1986, *7*, 47-70.

van Damme, Eric, *Refinements of the Nash Equilibrium Concept*, Berlin: Springer Verlag, 1983.

Van Huyck, John, Ray Battalio, and Richard Beil, "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American, Economic Review*, 1990, *80*, 34-248.

Van Huyck, John, Ray Battalio, and Richard Beil, "Strategic Uncertainty, Equilibrium Selection Principles, and Coordination Failure in Average Opinion Games," *Quarterly Journal of Economics*, 1991, *106*, 885-910.