

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

INCREASING COOPERATION IN PRISONER'S DILEMMAS BY ESTABLISHING A PRECEDENT OF EFFECIENCY COORDINATION IN GAMES

Colin F. Camerer

Mark Knez

Graduate School of Business University of Chicago



SOCIAL SCIENCE WORKING PAPER 1080

January 2000

**Increasing Cooperation in Prisoner's Dilemmas by Establishing a
Precedent of Efficiency in Coordination Games**

Colin Camerer

Marc Knez

Abstract

Coordination games have multiple Nash equilibria (i.e., sets of strategies which are best responses to one another). In "weak-link" coordination games players choose a number 1-7. Their payoff is increasing in the minimum number (or weakest link) and decreasing in the difference between their number and the minimum. Choosing 7 is an "efficient" equilibrium because it gives everybody a higher payoff than any other coordinated choice. Higher-payoff equilibria are riskier, however, so the game expresses the tradeoff between group efficiency and personal risk present in many social and organizational settings. We tested whether choosing efficiently in a weak-link game increases cooperative play in a subsequent prisoner's dilemma (PD) game. This cross-game transfer resembles transfer of cooperative norms in small firms (which are more like coordination games than PDs) as firms grow larger and become like PDs. In two experiments, if a group of players share a history of playing the weak-link game efficiently, that efficiency precedent can transfer to a subsequent PD game, improving the level of cooperativeness. The effect of transfer is much larger in magnitude (increasing cooperation from 15-30% to 71%) than the effects of most variables in previous PD studies. However, the transfer effect depends on descriptive similarity of strategies in the two games, since it largely disappears when the strategies are numbered differently in the weak-link game and the PD.

Forthcoming, Organizational Behavior and Human Decision Processes. Address reprint requests to: Colin Camerer, H&SS, M/C 228-77, California Institute of Technology, Pasadena, CA 91125 USA. Phone: 626/395-4054. E-mail: camerer@hss.caltech.edu

Increasing Cooperation in Prisoner's Dilemmas by Establishing a Precedent of Efficiency in Coordination Games

Blau (1964) describes the development of cooperation within exchange relationships as a slow process, starting with minor transactions in which little trust is required because little risk is involved, and eventually expanding into much riskier situations that require significantly higher levels of trust to generate cooperative behavior. If this sort of trust-building works, it does so because a counterpart's behavior in the less risky situation serves as a *precedent* for his or her behavior in the riskier situation. Since organizations are complex webs of exchange relationships (Baker, 1994), these precedent effects are important in organizational environments.

In this paper, we use game theory to formalize how cooperative relationships develop through precedents. Then we run experiments to see whether a precedent of efficient play in a coordination game will raise the level of cooperative behavior in a finite horizon prisoner's dilemma (PD) game. We first illustrate our idea with a simple example, and later present the more complex games used in the experiments.

Consider the PD game and the coordination game displayed in Tables 1a and 1b. (The upper left payoff in each cell is the row player's payoff; the lower right payoff is the column player's payoff.) The "efficient" outcome is the one which is unambiguously better for both players. In both games efficient outcome occurs when both players select C (cooperate). However, in the PD game it is a dominant strategy to play D (defect); individually rational players who are trying to maximize their own payoffs will both choose D, leaving both players worse off.

In the coordination game, in contrast, choosing C is a best response to an expectation of C; so if one player expects the other to cooperate (with sufficiently high probability) then she will cooperate, and both player are better off. In the language of game theory, the efficient (C,C)

outcome is a Nash equilibrium in the coordination game, but the inefficient outcome (D,D) is the only Nash equilibrium in the PD game. Note that the only difference between these two games is the payoff a player receives if she defects when the other players cooperate (known as the *temptation* payoff). In the PD, the temptation payoff (of eight) is higher than the payoff of five to reciprocating cooperation, while in the coordination game the temptation payoff (two) is lower than the reciprocation payoff.

Our prediction is that the shared experience (or *precedent*) of playing the efficient outcome in a coordination game will increase the likelihood of cooperative (efficient) play in a finitely repeated PD game.

Why Study Transfer of Precedent?

Before describing where the precedent hypothesis comes from, it is useful to ask why it is important to study the effectiveness of precedent across social situations. There are at least three reasons. First, many firms use training practices which depend on precedent for their effectiveness. In recent years businesses have become particularly interested in how to get people within their organizations to trust one another more and coalesce as teams. Organizations use trust-building exercises, such as asking one person to fall backwards with her eyes closed while others prepare to catch her before she hits the ground, to build trust. The implicit hypothesis is that the falling person will learn to trust fellow employees in general, so she will trust them more in business situations (trusting underlings to finish an urgent report on time or prepare adequately for a big client presentation). A similar example is taking groups of employees on “Outward Bound”-type outdoor adventures that require teamwork (such as rock climbing, or white-water rafting). The hypothesis, again, is that learning to work like a team in the wild will establish a precedent which transfers to working like a team in the concrete jungle back home. We know of no scientific evidence on whether these exercises work. Experiments like ours provide one very simple way to evaluate whether transfer of efficiency precedents works, which may provide some clues about whether these business exercises work.

A second reason to study transfer of precedent as we do is that the transition from coordination games to PD games mimic's the path of organizational challenges as firms grow from small to large, in an extremely simplified way. Suppose that a basic interaction in organizations involves how much individuals contribute to the organization's goals, and that players prefer to contribute little money or time (they play a PD; Mannix, 1991). The key idea is if players are empathetic toward one another, or if players are sufficiently able to monitor, and punish or sanction, employees who defect, then PD games actually have the same utility payoffs as coordination games. Under those conditions, self-interested players prefer to reciprocate cooperation rather than defect, corresponding to Table 1b payoffs rather than Table 1a payoffs (Camerer & Knez, 1996a, for details). Conditions of high empathy and easy monitoring and punishment are more likely to hold in small firms, since employees will know one another personally, and everyone knows who is pitching in and who isn't. As firms grow larger, new employees are less likely to empathize with one another, and the ability to catch and punish defectors probably falls, so that the games become PD's. Therefore, the same organizational activities which begin as coordination games at a small scale---one employee will do her part if she expects others to do theirs (because she cares about the others, or is afraid of getting caught)--disintegrate into PD's as firms grow larger. If small firms are able to achieve a precedent of efficiency by coordinating, and the precedent transfers to subsequent games, they may be able to sustain efficient behavior even when the basic game their employees play changes from coordination to a PD. Our experiments mimic precisely this process of switching from coordination to PD over time, and tests whether a precedent of efficiency transfers. To the extent that the experiments are externally valid-- which is, of course, always a matter of debate-- the results will suggest whether firms that start out small, and are able to establish a culture or norm of cooperation, will be able to sustain that norm when it is no longer sustained by empathy or fear.

Obviously, the analogies between our very simple experiments and naturally-occurring organizations are imperfect. The idea is to use the experimental findings to inspire a search for similar patterns in field data. The simple experiments are also a flexible baseline which can be enriched (like a blank canvas that can be painted): If the experiments appear to lack generalizability because they do not have a particular feature of naturally-occurring organizations, if that feature can be clearly articulated then it can be added to the experimental design in future work.

A third contribution of our paper is to study a new way in which the PD might be “solved.” Many studies have found that changing the way the game is described, allowing players to communicate, and so forth, have modest, albeit interesting, effects in the rate of cooperation and defection (Ledyard, 1995; & Sally, 1995). Our study finds a huge change in the rate of cooperation due to transfer of precedent: In two-player PD's, 15% of players chose the most cooperative action when playing PD's for the first time (and 30% did so in the first period of a second round of PD's), while 71% of players who played a coordination game efficiently subsequently chose the most cooperative action in PD. Putting aside questions of interpretation, this effect is simply much larger than the effects of most variables, which have been exhaustively studied (Sally's, 1995, review).

Weak-link Coordination and PD Games

We use the “weak-link” coordination game shown in Table 2 to study precedent. In the weak-link game, each subject picks an action from the set of integers $\{1, 2, \dots, 7\}$. The table shows the payoffs from selecting a particular action (a row), given the minimum action chosen by the other subjects (a column). For example, player A picks a five while B picks a three. The minimum action is three. A earns \$.70 and B earns \$.90. Player B “loses” \$.20 by picking an action two units above the minimum action selected. Notice that each player wants to select exactly the minimum of the other players, but everyone wants the minimum to be as high as possible. As a result, high actions are risky because other players may select lower actions.

This game is called a weak-link game because, like a chain whose strength depends on its weakest link, the common component of each player's payoff depends on the lowest action anyone takes. Camerer and Knez (1996a,b) suggest that weak-link games are useful models of social and project-based organizational processes which are very sensitive to the level of the worst input (Bryant, 1983; Becker & Murphy, 1992). A social example is meeting a large group of people at a crowded restaurant which will not seat anybody until the last person arrives. An economic example is price competition among identical firms, who must match the cheapest price (and lose more customers when their original price is higher). Organizational examples include chemical recipes, and safety in "high-reliability organizations" (such as nuclear power plants), which have accidents if a single mistake occurs.

Many examples involve synchronicity in time. Suppose (i) a team's product is not finished until the slowest player finishes; (ii) everyone prefers the product to be finished earlier; and (iii) no player wants to finish their portion before the slowest player does. Then the players are playing a weak-link game. These conditions hold in any organization or work unit for which the completion of a project requires that a set of highly interdependent activities are performed in a timely fashion. Examples include investment banking activities (doing deals), software development, building satellites and---familiar to academics---contributing chapters to an edited book. In these examples, free-riding is not the source of inefficient outcomes; it's not possible to free-ride, per se, because if free riders do not perform their tasks they do not benefit. Instead, the source of inefficiency is a combination of a lack of mutual understanding about the game and mutual expectations of how others will behave in a highly interdependent work environment. The weak-link game allows us to study mutual expectations while controlling for mutual understanding.

The weak-link game has two key game-theoretic properties. First, each of the outcomes on the diagonal satisfies the mutual best response property of Nash equilibrium (but the off-diagonal outcomes do not). Therefore, any number X has the property that, in theory, if players

expect it to be the minimum, they will choose it and it will be the minimum. Second, every equilibrium is “Pareto-dominated” by all higher-numbered equilibria (i.e., better for everybody); therefore, everyone choosing seven is the best equilibrium of all. These properties create two simultaneous coordination problems: (i) Players would like to coordinate on some equilibrium (since any players who choose above the minimum are worse off than if they had matched it); and (ii) players would like that equilibrium to be as high an action as possible. While the weak-link game requires extremely good coordination---because payoffs depend on the minimum action (rather than the median or average; cf. Van Huyck, Battalio, & Beil, 1991)---this sensitivity is useful because it tends to create between-group variation in efficiency. Between-group variation is helpful for testing how groups with different efficiency precedents behave in a subsequent PD game.

Experimental research on the weak-link game has found that sufficiently large groups of subjects have difficulty coordinating their actions on the efficient equilibrium. Van Huyck, Battalio, & Beil (1990) find that in seven experimental sessions, with groups of size 14 to 16 subjects, at least 80% of the subjects select a one or a two by the 10th period of play. Camerer and Cacho (1996) replicate this result with nine-subject groups. Of course, it may not be too surprising that with groups this large, subjects are unable to coordinate on the efficient equilibrium since it only takes the selection of a low action by one player to disrupt the efficient equilibrium.

Van Huyck et al. (1990) also report the results of 14 two-player weak-link games, where subjects play in dyads for seven periods. In 12 of the 14 sessions subjects do play the efficient equilibrium by the seventh period. However, the efficient equilibrium is played in the first period in only one of the fourteen pairs. Hence, pairs of subjects do not start out playing efficiently, but learn to do so reliably in several periods. In Camerer and Knez (1994), 20 three-player weak-link games were played for five periods. The efficient equilibrium was reached in only four of the 20 three-player groups, and five of the groups converged to the least efficient equilibrium (choosing

all one's) by the fifth period. Comparing Van Huyck et al.'s (1990) two-player results and these three-player results shows that increasing the number of players from two to three significantly reduces subjects' ability to coordinate on the efficient equilibrium.

Weber, Camerer, Rottenstreich, and Knez (in press) studied weak-link games in which one randomly-chosen "leader" encouraged the group to choose large numbers, to see whether limited one-way communication would improve efficiency. It did not: Large (8-10 person) groups were not able to achieve efficiency, and small (two-person groups) reached efficiency, but did so even without the leader's speech. Furthermore, players committed a "fundamental attribution error" by blaming the leaders for the failure of large groups to coordinate efficiently, and crediting leaders for the success of small groups.

Our experimental design compares weak-link games to two closely related PD games. The first, displayed in Table 3a, superficially resembles the weak-link game because each player selects an action from the set {1, 2, 3, 4, 5, 6, 7}. Action seven corresponds to the efficient (cooperative) outcome and action one corresponds to the inefficient outcome. The three-action PD game, displayed in Table 3b, is the seven-action game with actions 1, 4, and 7 relabeled 1-3.

Both of these games are multi-step versions of PD. Actions between one and seven represent varying degrees of cooperation. Increasing your number helps the other player more than it hurts you, so if both players choose the highest number they are both better off. That is, if each player "cooperates" and selects a seven (or three) then they both receive their Pareto-efficient payoff of \$2.10. If one player selects one ("defects") and the other player cooperates by selecting a seven (or three), then the defector receives his temptation payoff of \$2.70, and the other player receives the sucker payoff of \$.90. If both players select one (the dominant strategy) they both earn \$1.50, which is the unique Nash equilibrium of the PD game.

There are many, many studies of PD games (for reviews see Dawes, 1980; Murnighan, 1994; Sally, 1995; and Ledyard, 1995). Most of this research tests variables which are hypothesized to change the amount of cooperation. Structural factors include the structure of

payoffs in the game and the expected length of the relationship. Psychological factors include individual values, knowledge, communication, and group identity (Murnighan, 1994). The last two factors come into play in the current study. Here is how.

Many studies have found that players who expect their counterparts to cooperate in a PD game are more likely to cooperate themselves (Dawes, 1980). This results suggest that one may be able to raise levels of cooperate behavior by raising players' expectations of cooperative behavior. Indeed, one of the variables which increases cooperation most strongly is communication about the game itself and preplay promising (although merely talking has little effect; see Dawes, McTavish, & Shaklee, 1977; Sally, 1995). A precedent of efficient play in an earlier game is like a kind of tacit communication about what will happen in the future, which may increase cooperation by raising expectations of cooperation by others. (We test this directly by asking subjects both how they want to play, and what they expect others to do.)

Research on group identification (Kramer & Brewer, 1984, 1986) suggests that increasing group identity will induce higher levels of cooperation, and can affect organizational behavior (Kramer, 1992). They argue that higher levels of interpersonal trust follow from common membership in a group, because members expect other group members to treat them well, and reciprocate (an *expectations* effect), members prefer to treat others well (a *preference* effect), or because both effects occur. Since precedent requires shared experience, it may inherently create group identity as well. We control for this effect by comparing subjects who play the weak-link game, followed by a PD, with subjects who play the five-period PD once, then play it again. If group identity, per se, creates more cooperation then subjects should cooperate as much in the second PD as they do in a PD which was preceded by the weak-link game.

Hypotheses

Equilibrium Predictions

What does game theory predict about behavior in these games? There are two kinds of hypotheses: Predictions about what players will choose in the games; and predictions about how behavior in the weak-link game will create a precedent which affects behavior in the PD games.

In the coordination game there are multiple equilibria: Any pair (X,X) is an equilibrium. In games with multiple equilibria, players face the problem of converging on shared expectations about which equilibrium will be played. (If their expectations, and consequently their behavior, are not shared, they are "out of equilibrium" and have an incentive to change their strategies.) Following Schelling (1960) game theorists have described focal principles which distinguish certain equilibria by their "salience" or "psychological prominence." A salient equilibrium outcome, by definition, is one which players will play if they expect others to play it too (if their expectations are self-fulfilling). Sources of salience include the structure of payoffs (Pareto-efficiency, equal-sharing) or strategy labels and context (Crawford, 1997). Experimental studies of coordination games have established a wide variety of focal principles, which subjects use rather cleverly to coordinate their behavior (Crawford, 1997; Camerer, 2000, chapter 7).

In the weak-link games, different focal principles predict different outcomes. Choosing the strategy which maximizes the minimum payoff (maximin), and "risk-dominance," both predict choices of one. Payoff-dominance predicts choices of seven. Maximizing against the belief that the choice of the other player will be random predicts all choices are equally likely. Game theory simply does not make a precise prediction about what is likely to happen in these games.

In the PD games in Tables 3a-b, the unique Nash equilibrium prediction is that self-interested players will choose one if the game is played once. However, repeating the game five times (with the same dyad partner) changes the prediction. Modern theories showed that cooperation in the finitely-repeated PD can be an equilibrium if players are not perfectly certain

that other players are rational and self-interested (so that it pays for players who would defect in one-shot PD's to cultivate reputations for being cooperative, Kreps et al., 1982). Thus, one prediction of game theory is cooperative choices of seven until close to the end (which is the pattern usually observed in repeated-PD experiments, Murnighan & Roth, 1983; Cooper et al., 1996). However, defecting throughout (choosing one) is also an equilibrium. Indeed, since defecting throughout is safer for players, but worse for everybody, the repeated game has a payoff structure like the coordination game in Table 1b. The key challenge for players is to coordinate their beliefs on whether their partner is likely to choose the cooperative path, in which case they will want to cooperate also, or whether their partner is likely to choose the defection path, in which case they will defect too.

Precedent Hypotheses

The fact that game theory makes multiple predictions about behavior in the repeated PD sets the stage for the role of precedent. The focal principles listed above predict what will happen in a multiple-equilibrium game based solely on the structure of its payoffs, or on contextual labels which focus attention on psychologically-prominent strategies. Precedent--having played a particular equilibrium in the past--is another focal principle. As Lewis (1969;36) stated, "Precedent is merely the source of one important kind of salience: conspicuous uniqueness of an equilibrium because we reached it last time." Several game-theoretic studies have shown the empirical power of precedent (Van Huyck, Battalio, & Beil, 1991; Knez, 1998).

Note that precedent does not imply that people will simply repeat what they did in the past. Instead, it says that expecting others to do what they did in the past (and expecting that they will think you will do what you did in the past, etc.) can coordinate expectations about that of many equilibria will happen, if these expectations are self-fulfilling. Precedent creates a socially-understood convention, rather than (simply) an individual preference.

The central hypothesis of this paper is that a precedent of efficient play in a weak-link game will increase the likelihood of efficient (or cooperative) play in a subsequent five-times-repeated PD.

Which precedents transfer from one game to another undoubtedly depends on some shared perception of the similarity of the games. There is no general theory of what games players regard as similar, and creating one is extremely important, but far too ambitious a task to begin here (Camerer, 1998; Warglien, Devetag, & Legrenzi, 1999, for a start). However, such a theory should distinguish two kinds of similarity a pair of games could have: Descriptive similarity, and payoff similarity. (This distinction parallels the important distinction between surface structure and "deep" structure in problem-solving; e.g., Singley and Anderson, 1989). Descriptive "similarity" is defined by the number and identity of players, the actions available and how they are labeled, the rules of the game, and so forth. Payoff similarities is defined by the payoffs associated with action combinations, the equilibrium properties of the outcomes, and socially defined properties of payoffs (such as Pareto-dominance, equity, first-mover advantage, and so forth). Indeed, game theory is often interpreted as a taxonomy of games grouped into structurally similar classes by their set of equilibria or payoff properties (all PD's have certain properties, all signaling games have other properties, and so forth). A complete theory of similarity would address each of these dimensions separately and how they interact with one another to determine perceptions of similarity.

The seven-action weak-link game and the seven-action PD game have descriptive similarity, because their strategies are labeled in the same way (integers 1-7). They also have payoff similarity, because both have multiple equilibria in which the efficient equilibrium (choosing seven) is riskier and the inefficient equilibrium (choosing one, until near the end in the PD) is safer. However, the seven-action weak-link game and the three-action PD game only have payoff similarity. They are not descriptively similar because the weak-link actions 1-7 do not correspond precisely to the PD actions 1-3. Comparing the effect of previous experiences with

seven-action weak-link games on behavior in subsequent seven- and three-action PD's therefore gives an empirical sense of whether transfer requires both descriptive and payoff similarity, or can be sustained if there is only payoff similarity. (In future work, it would be nice to see whether descriptive similarity alone can generate transfer.)

Experiment 1: 2-player Games

METHODS

Participants

Subjects were University of Chicago undergraduates in sessions 1-6 and Caltech undergraduates in sessions 7-8. The Chicago students were recruited from announcements in economics classes and signs posted around campus. They were promised \$3 for participating plus an unspecified sum that depended on the decisions they made and the decisions of others. The Caltech undergraduates were given the same information about earnings and were recruited from an electronic subject mailing list.

Design

Table 4 summarizes the design. In all sessions, several two-subject dyads played two consecutive rounds of a five-period game. The games are denoted by W (weak-link), and P(3) or P(7), for three- or seven-action PD. In treatment WP(7), for example, each dyad played the weak-link game in round one and then played the seven-action PD game in round two. In treatment PP(7), subjects played the seven-action PD in both rounds.

Using this notation, the precedent hypothesis is that there will be more cooperation in the second-round PD in treatment WP(7) than in the second-round PD in treatment PP(7) (provided subjects play efficiently in the weak-link game). If precedent transfer is only due to descriptive similarity, there will be more second-round cooperation in WP(7) than in PP(7), and equal amounts of second-round cooperation in PP(7) and in WP(3) (where there is payoff similarity but no descriptive similarity). If precedent transfer is due to payoff similarity, there will be equal amounts of transfer in those two comparisons.

Procedure

Ten to fourteen subjects participated in each session. Subjects sat in a room together and were given a common set of instructions, which were read aloud (so subjects could be sure that others had the same instructions as they had). Subjects were told there were two rounds of play, and each round had five periods. They were not told about the structure of round two until round one was over. While subjects were organized into dyads, they did not know who their dyad partner was. The instructions (shown in the appendix) were written abstractly.

In general, we followed conventions in experimental economics rather than conventions in psychology (Camerer, 1996; Loewenstein, 1999; Hertwig & Ortmann, in press, for more discussion). Subjects were actually paid their earnings from the games (in addition to a \$3 show-up fee); the game was repeated to allow learning and equilibration; there was no deception; and the game was described abstractly. Abstract description does not seem to matter empirically because an earlier experiment found no differences between the weak-link game played by choosing 1-7 (as in our experiment) or described by a cover story involving timing of contributions to a group project (Weber, et al., in press).

In each period subjects chose an action from the set of actions {1, 2, 3, 4, 5, 6, 7} (or {1, 2, 3} in the three-action PD). In addition to making choices, each subject guessed the action of the other player in their dyad, and earned \$.10 if they guessed correctly. After each choice, they were told the minimum choice of their dyad, computed their earnings from that period, and awaited the next round.

RESULTS

The results are shown in Table 5 below. The Table pools the two sessions in each treatment, and shows the total number of subjects selecting a particular action in a particular period of play. For example, in the first period of the weak-link games played under treatment WS(7), 22 of 24 subjects selected a seven while the other two subjects selected a five.

In the weak link games in the first round of treatments WP(7) and WP(3), all but two dyads of subjects played the efficient Pareto-dominant equilibrium (choosing seven's) by the fifth period. (These data replicates the results of Van Huyck et al., 1990.)

In the PP(7) experiments, in contrast, by the fifth period of play in the first-round PD, most of the subjects chose an inefficient one (15/22) while only two subjects (in different groups) chose seven. Hence, a homegrown precedent of efficient play was generated in all but two of the weak-link dyads, while a precedent of inefficient play was generated in the first-round PD dyads.

We first test the precedent hypothesis by comparing the distributions of actions chosen by subjects in the first period of the second round, across the WP(7) and PP(7) treatments. The results are consistent with the precedent hypothesis, because PD play in round two after weak-link play WP(7) is significantly different than first-period play in both the first and second rounds of the PP(7) control groups, at $p < .001$ and $p < .05$, respectively (using a conservative Kolmogorov-Smirnov test, see Table 6).

The effect of first-round precedent is also quite large in magnitude: Most of the subjects (71%) choose efficiently in the first period of round two of WP(7), compared to small minorities of 15% in the first period of round two of PP(7) and 30% in the first period of round one of PP(7).

The comparison between second-round choices in WP(3) and PP(3) tests whether precedent has an effect when the two games are only payoff similar, but are not descriptively similar. There is a minor effect of playing the weak-link game first--50% choose the efficient action in the first period of WP(3), compared to 29% and 42% in the round one and round 2 of PP(3)---but the effect is not statistically significant.

Our precedent hypothesis is about shared expectations of cooperative actions. Therefore, it is important that experience in the weak-link game has a strong effect on guesses about what others will do, which are made before period one in the second round (shown in parentheses in Table 5). Once again, the precedent effect on expectations is relatively strong for seven-action games ($p = .02$ and $.10$ for respective comparisons with first- and second-round PP(7) results) and

weak for transfer from seven-action weak-link games to three-action PD's. This is important evidence for the precedent interpretation. If players were simply repeating the numbers they chose earlier when they begin playing the PD, they would not necessarily say they thought others would repeat the same numbers, but they do.

Experiment 2: 3-Player Games

Experiment 1 showed that there is an effect of precedent: A history of efficient play in the seven-action weak-link game is correlated with an increase in efficient play in a subsequent PD game. However, this seems to be primarily due to descriptive similarity rather than establishing a precedent of efficiency that transfers across payoff-similar games, because the effect of seven-action weak-link precedent on subsequent three-action PD's is weak.

A drawback of the dyadic results is that coordinating in the weak-link game is “too easy” for dyads, so we did not observe any variation in efficiency across dyads. Based on earlier work (Camerer & Knez, 1994), we hypothesized that using three-player groups would generate a wider range of variation in efficiency levels in the weak-link part of the experiment. Wider variation would allow a cross-group comparison of how precedents of efficiency (or inefficiency) in coordination games affects play in subsequent PD's, to complement the cross-condition results from experiment 1.

Subjects

Subjects were 33 University of Chicago undergraduates recruited and paid as in Experiment 1.

Design

The design had two treatments. In the first treatment, subjects were grouped into triads. Each triad played five periods of a three-person weak-link game in the first round, followed by five periods of a three-person PD game in the second round. As in Experiment 1, we refer to this treatment as 3WP(7) (where the 3 denotes triadic behavior rather than dyadic).

The second treatment was conducted after the results of the first. (Thus, these two treatments are not cells in a proper design, but they are reported together for coherence and brevity.) The second treatment was designed to generate a larger percentage of efficient triads than in the first treatment. To do so, we exploited the facts that (a) two-person dyads behave quite efficiently, and (b) precedent has empirical force in coordination games, to see if efficient triads could be “grown” by adding one subject to a dyad who had played together for several periods. We refer to this treatment as 2,3WP(7).

Procedure

Experiment 2 procedures were almost identical to experiment 1, except that players in a dyad or triad were told the choices by all subjects in their group, rather than simply the minimum choice.

In 3WP(7) sessions 1-2, the procedures were the same as in Experiment 1, except that 21 players were organized into triads. Each triad first played a three-person weak-link game with the same payoffs as in Table 2 (except a player's payoff depended on her choice and the minimum of the other two players' choices).

In the one 2,3WP(7) session, a total of 12 subjects participated. In the first round of the experiment, eight of the twelve subjects were organized into dyads. Each dyad played a weak-link game for five periods. We call these subjects *incumbents*. The remaining four subjects did nothing during this first round of play (although they read the instructions for the first round along with the subjects participating in dyads as they were being read out loud by the experimenter).

In the second round, one subject who did not participate in the first round was added to each of the dyads from the first round, creating a triad. We call the added subjects *entrants*. Each entrant was told the history of actions selected in all periods of the first round by the two incumbents in their newly-formed triad, and the incumbents knew the entrants were given that history. (Camerer & Knez, 1994, found that subsequent behavior was sensitive to this kind of

history information, which is more evidence for the precedent view instead of the theory that players just repeat old choices.)

Each triad then played a three-player weak-link game for five periods. Then, in a third five-period round, the triads played a three-player (seven-action) PD game, with payoffs shown in Table 7. The three-player PD is like the two-player PD except that a player's payoffs depend on her own choice, and on the sum of the choices of the other two players.

We hypothesize three effects. First, we expect higher numbers (and more efficiency) in the weak-link games in the 23W(7) treatment, compared to the 3WP(7) treatment. We expect this because dyads are likely to play efficiently, and adding one entrant who knows the incumbents' history is likely to lead to all three triad members playing more efficiently than groups which began as triads. Second, if the first hypothesis is confirmed, creating a stronger efficiency precedent in the 2,3WP(7) condition, we expect to see more cooperation in the subsequent third-round PD than in the corresponding second-round PD in the 3WP(7) condition. Third, we expect dispersion in the numbers chosen by different triads (both within the 3WP(7) condition, and pooling across triads in the 3W(7) and 2,3WP(7) conditions); and we expect that triadic efficiency in the weak-link game will be correlated, across triads, with cooperation in the subsequent PD.

Results

The results of experiment 2 are displayed in Tables 8a-b. (For brevity, only periods one and five of rounds one and two are displayed in Table 8b.)

First look at the 3WP(7) results in Table 8a. Consistent with the results reported earlier (Camerer & Knez, 1994), triads had difficulty reaching the efficient equilibrium in the weak-link game. Out of seven triads, one coordinated on 7 and two groups came close.

Unlike in experiment 1 (when almost all dyads converged on seven), differences in behavior among the triads permit a cross-triad comparison of efficiency in the weak-link game

with first-period behavior in the subsequent three-person PD, to test for transfer of precedent. There is modest evidence of transfer, because the correlation between actions selected in the last period of the weak-link game and actions selected in the first period of the PD game is .25, which is marginally significant ($p = .063$, one-tailed z-test).

Now turn to Table 8b, which summarizes results from the 2,3WP(7) session. All four of the two-player groups played the efficient equilibrium by the final period of the first round. In round 2, after a third player was added, all four three-player groups played the efficient equilibrium in every period of play. These results are in striking contrast to the behavior observed in 3WP(7) in Table 8a, in which only one of the seven three-player groups reached the efficient equilibrium.

The comparison between de novo triads in Table 8a, and “cultivated” triads in Table 8b shows the power of precedent within a sequence of plays of the same game (even when the number of subjects grows). Weber (1999) took the idea further by starting with dyads and gradually adding entrants who knew the previous player's history, one by one, until the group size reached 12. He was able to “grow” large groups, if the rate of adding entrants was slow enough, which played much more efficiently than equal-sized groups which started out large. (Recall that all the 14-16-person groups in Van Huyck, Battalio, & Beil's (1990) study converged to the worst equilibrium, a minimum of one.) In addition, subjects who were able to endogenously control the rate of growth tended to add entrants far too quickly (leading to inefficient large groups); they did not seem to appreciate how quickly coordination failure could occur, and how hard it was to reverse. Weber interprets these results as a partial explanation for why small organizations grow too quickly and subsequently fail.

Table 8b also shows that the efficient triads in the 2,3WP(7) condition play the subsequent three-player PD quite cooperatively (75% choose the most cooperative strategy seven). This pattern is consistent with a transfer of precedent, but since there is no variation in

triadic weak-link behavior, and no three-person PP(7) control group, we cannot conclude that precedent per se causes the high level of PD cooperation.

By pooling across triads in the 3WP(7) condition and the 2,3WP(7) condition, we get a healthy amount of variation in triadic weak-link efficiency, to test whether different efficiency precedents affect behavior in the subsequent PD. Pooling across the conditions, the correlation between individual choices in the final weak-link period and the first PD period is .275 ($p = .028$, one-tailed z-test). Taking triads as the unit of analysis, 11 of 15 (73%) subjects who had played in an efficient group in the weak-link game chose a seven in the first period of PD, while 4 of 18 (22%) subjects from inefficient triads chose a seven ($\chi^2 = 8.62, p < .004$). The corresponding figures for subjects' expectations about what others would pick in the final-round PD are 12 of 15 (80%) for efficient triads and 3 of 18 (17%) for inefficient ($\chi^2 = 13.24, p < .001$). Of course, these pooled results should be taken with a grain of salt because the conditions leading up to the final period of weak-link play are different in the two conditions (the 2,3WP(7) subjects played more periods, and had a history of efficiency from their dyadic past).

Discussion and Conclusion

The goal of this paper was to formalize and test the intuition of Blau (1964) and others that cooperation develops over time in exchange relationships as exchange partners move from less risky situations to more risky situations.

We start by asking what type of prior experiences in strategic situations (games) support the trust building process. One place to look for “experience effects” is to see whether experience in previous PD games affects later PD games. Bettenhausen and Murnighan (1985, 1991) have done this and found evidence that norms of cooperation develop, which transfer across different PD games. However, their results beg the question of what generated cooperation in the earlier games, and raises a new question: Does cooperation transfer across different types of games? We address the second question by testing whether efficient cooperation in coordination games

creates a precedent of efficiency that transfers to PD games, increasing cooperation (relative to control groups).

The coordination games we studied are “weak-link” games in which players choose actions from one to seven. A player's payoff increases with the minimum action anybody chose, and decreases with the distance between their own action and the minimum. In these weak-link games, there is a “Nash (mutual best-response equilibrium)” in which everybody chooses seven, playing efficiently. We tested the hypothesis that the shared experience of playing the efficient equilibrium in the weak-link game would create a precedent of efficient play strong enough to generate expectations of cooperation in a finitely-repeated PD game, and whether those expectations would lead to cooperation.

There are several justifications for using coordination games to generate a precedent of cooperation. First, in these coordination games the efficient outcome is an equilibrium (unlike in one-shot PD's) but reaching it is by no means easy. It requires players to have sufficient faith that their counterparts will select the efficient action. Since efficiency is not assured, one might think that substantial trust is built when efficiency is reached.

Second, there is a game-theoretic justification for using coordination games to generate “cooperative” experiences. The folk theorem of repeated games says that when a game is played repeatedly, equilibria can arise that are not equilibria in the stage game. A famous example is “tit-for-tat” in the repeated PD (play cooperatively at first, then mimic what the other player did in the previous period), which supports the cooperative outcome in the repeated PD game, even though defection is predicted in the one-shot PD. While cooperation can be an equilibrium in the finitely-repeated PD (until near the end), defection all the way through is also an equilibrium. Given the multiplicity of equilibrium implied by repetition of the game, players face the problem of selecting one of these equilibria---that is, they face a coordination problem. In fact, it is easy to show that the infinitely repeated PD game has the payoff properties of a weak-link game (Miller, 1992; Camerer & Knez, 1996a,b). If players transfer precedents from one game to a new

game that has similar payoff properties, they should transfer efficiency from the weak-link game (if they achieve it) to the repeated PD.

Finally, there is a practical justification for examining transfer from coordination games to PD's. Organizations are complex webs of exchange relationships in which participants (workers and managers) implicitly compete with their peers for resources and promotions (effectively playing PD's). At the same time, they are highly dependent on these peers in their day-to-day, interdependent work activities in which they have common goals. All suffer if one person makes a mistake (effectively playing weak-link games; Baker, 1994). Indeed, discussions of organizational culture implicitly assume that expectations of peer behavior generalize across these types of situations to form a set of "organizational expectations." Little is known about how these expectations are generalized from one organizational situation to another. Experiments with sequences of different games are one simple way to begin studying the generalization process. Experiments with coordination games followed by PD's are also the right way to study organizational development over time, if small organizations resemble weak-link games but large organizations resemble PD's (for reasons mentioned in the introduction).

To study transfer of precedents from coordination to PD games, we conducted two experiments with a total of 123 subjects. The results show that positive precedent effects are possible. In dyadic experiments, players who have participated in seven-action weak-link games, and reached efficiency, tend to choose much more cooperatively in subsequent seven-action PD's, compared to a control group of subjects who simply participate in two PD's in a row. There is also a substantial effect of precedent across three-person triads (which vary more in efficiency than dyads): Triads which played more efficiently in weak-link games tended to expect more cooperation, and also to be more cooperative, in subsequent PD's. At the individual level, the correlation between final-period weak-link actions and first-period PD actions was .25-.30 (modestly significant).

We have one hunch about why these precedent effects are not larger in magnitude. There may be a “reverse precedent effect” which causes behavior to go in the opposite direction of an inefficient precedent. Playing inefficiently in the weak-link game reminds players of how costly inefficiency can be. This shared experience could galvanize their collective desire to play cooperatively, provided they get a chance to start over, which they do in the second-round PD. In this interpretation, behavior generates a precedent of inefficiency, which is overruled by a meta-precedent---“do what we did before if it worked, and do the opposite if it failed” (cf. Seely, Van Huyck, and Battalio, 1999).

A very important qualification from our data is that the effect of a precedent of efficiency in seven-action coordination games on three-action PD's is much weaker (and insignificant) than the effect on seven-action PD's. This difference suggests transfer works best across games which are descriptively similar (have the same labels or surface structure, as in transfer of problem-solving skills; Singley & Anderson, 1989). Transfer is weaker when games that have similar payoff structures but different descriptions.

Even if the lion's share of the precedent effect is due to surface transfer across descriptively similar games, documenting the effect is a contribution to the vast literature on cooperation in the PD (and in related public goods games). Indeed, the size of the effects we reported are quite large compared to the effects of other variables that have been studied in previous PD experiments. For example, in this journal Pillutla and Chen (1999) reported a significant effect of labeling the PD actions differently (in an economic versus a non-economic way). Their labeling change increased first-period cooperation from 32% to 39%, an effect which is only a fraction as large as the corresponding increase, from 15-30% to 71%, in our two-player games. Thus, even if the precedent effect we observe is simply due to transfer of expectations that players will keep choosing similarly-numbered strategies, it represents a huge increment in cooperation.

A similar point is made by Ahn et al. (1998). Like us, they found that efficiency in 2-person weak-link (“stag hunt”) games influenced the rate of cooperation in subsequent one-shot PD's (which were descriptively similar to the weak-link games). Subjects who had encountered no efficient outcomes in eight previous stag hunt games cooperated in subsequent PD's with 16% probability, while subjects who had encountered eight previous efficient outcomes cooperated 65% of the time. They also found that the precedent effect is much stronger in magnitude than the effect of changing payoffs (changing the “temptation” or greed premium to defecting against a cooperator, and the “sucker” or fear loss from cooperating against a defector).

While it is premature to conclude that precedent effects are necessarily reliable or cognitively deep (due to payoff similarity), our results and Ahn et al.'s find precedent effects that are much larger in magnitude than many other variables which have been exhaustively studied (such as labels and payoffs) in the PD. On an empirical basis, then, transfer of precedents across games certainly deserves to be studied more thoroughly as a determinant of PD behavior. The pressing open questions are whether players are really transferring a precedent of efficiency (as opposed to merely repeating, and expecting repetition of, previous behavior), and what features of similarity between games affect the amount of transfer.

Finally, in the introduction we noted that many discussions of organizational behavior implicitly assume transfer across organizational tasks or games, by discussing cross-situational concepts like organizational expectations, learning, and culture. We mentioned that many firms actively employ trust- and team-building exercises away from the office grind. They must think behavior in these exercises will transfer to workplace activities. Our results suggest that transfer is mostly limited to those cases in which activities have similar descriptions, rather than to those with similar strategic structures and different descriptions (cf. Glaeser et al, in press, who find low correlations across different measures of trust). Taken seriously, this claim predicts that making trust-building exercises too different than office life may undermine their usefulness by

limiting transfer. We wouldn't bet heavily that our conclusion is correct, but we will bet that an experiment evaluating cross-experience transfer would prove interesting.

References

- Ahn, T., Ostrom, E., Schmidt, D., Shupp, R., & Walker, J. (1998). *Cooperation in PD games: fear, greed and history of play*. Working paper, Indiana University.
- Baker, W. (1994). *Networking smart*. New York: McGraw-Hill.
- Becker, G., & K. Murphy (1992). The division of labor, coordination costs, and Knowledge. *The Quarterly Journal of Economics*, **107**, 1137-1160.
- Bettenhausen, K., & Murnighan, J. (1985). The emergence of norms in competitive decision-making groups. *Administrative Science Quarterly*, **30**, 350-372.
- Bettenhausen, K., & Murnighan, J. (1991). The development of an intragroup norm and the effects of interpersonal and structural challenges. *Administrative Science Quarterly*, **36**, 20-35.
- Blau, P. (1964). *Exchange and power in social life*. New Brunswick: Transactions Publishers.
- Bryant, J. (1983). A simple rational expectations Keynesian-type model. *The Quarterly Journal of Economics*, **98**, 525-529.
- Cachon, G., & Camerer, C. F. (1994). Sunk costs and forward induction in experimental coordination games. *Quarterly Journal of Economics*, **111**, 165-194.
- Camerer, C. F. (1996). Rules for experimenting in psychology and economics, and why they differ. In W. Albers, W. Guth, P. Hammerstein, B. Moldovanu, & E. Van Damme (Eds.), *Understanding strategic interaction: Essays in honor of Reinhard Selten*, (pp. 313-327). Berlin: Springer-Verlag.
- Camerer, C. F. (1998). *Mental representations of games*. Working Paper. California Institute of Technology.
- Camerer, C. F. (1999). *Behavioral game theory: Experiments on strategic interaction*. Manuscript in progress.

- Camerer, C. F., & Knez, M. (1994). Creating expectational assets in the laboratory: Coordination in "weak-link" games. *Strategic Management Journal*, **15**, 101-119.
- Camerer, C. F., & Knez, M. (1996a). Coordination in organizations: A game-theoretic Perspective. In Z. Shapira (Ed.), *Organizational decision making*. New York, NY: Cambridge University Press.
- Camerer, C. F., & Knez, M. (1996b). Coordination, organizational boundaries, and fads in business practices. *Industrial and Corporate Change*, **5**, (1), 89-112.
- Cooper, R., DeJong, D., Forsythe, R., & Ross, T. (1996). Cooperation without reputation: Experimental evidence from prisoner's dilemma games. *Games and Economic Behavior*, **12**, 187-218.
- Crawford, V. (1997). Theory and experiment in the analysis of strategic interaction. In D. Kreps, & K. Wallis (Eds.), *Advances in economics and econometrics: theory and applications, I* (pp. 206-242). Cambridge: Cambridge University Press.
- Davis, D., & Holt, C. (1992). *Experimental economics*. Princeton, New Jersey: Princeton University Press.
- Dawes, R. (1980). Social dilemmas. *Annual Review of Psychology*, **31**, 169-193.
- Dawes, R., McTavish, J., & Shaklee, H. (1977). Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation. *Journal of Personality and Social Psychology*, **35**, 1-11.
- Dawes, R., van de Kragt, A., & Orbell, J. (1988). Not me or thee but we: The importance of group identity in eliciting cooperation in dilemma situations: experimental manipulations. *Acta Psychologica*, **68**, 83-97.
- Gibbons, R. (1992). *Game theory for applied economists*. Princeton, New Jersey: Princeton University Press.

- Glaeser, E. L., Laibson, D., Scheinkman, J. A., & Soutter, C. L. (in press). What is social capital? The determinants of trust and trustworthiness. *Quarterly Journal of Economics*.
- Harsanyi, J., & Selten, R. (1988). *A general theory of equilibrium selection in Games*. Cambridge: MIT Press.
- Hertwig, R., & Ortmann, A. (in press). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences*.
- Knez, M. (1998). Precedent transfer in experimental conflict-of-interest games. *Journal of Economic Behavior and Organization*, **34**, 239-249.
- Kramer, R. (1992). Cooperation and organizational identification. In J. K. Murnighan (Ed.), *Social Psychology in Organizations* (pp. 244-269). Englewood Cliffs: Prentice-Hall.
- Kramer, R., & Brewer, M. (1984). Effects of group identity in resource use in a simulated commons dilemma. *Journal of Personality and Social Psychology*, **46**, 1044-1057.
- Kramer, R., & Brewer, M. (1986). Social group identity and the emergence of cooperation in resource conversation dilemmas. In H. Wilke, C. Rutte, & D. M. Messick (Eds.), *Experimental studies of PDs* (pp. 205-234). Frankfurt, Germany: Peter Lang Publishing Company.
- Kreps, D., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely-repeated prisoner's dilemma. *Journal of Economic Theory*, **27**, 486-502.
- Ledyard, J. (1995). Public goods experiments. In J. L. Kagel & A. E. Roth (Eds.), *Handbook of Experimental Economics*, (pp. 111-194). Princeton, New Jersey: Princeton University Press.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge, Massachusetts: Harvard University Press.
- Loewenstein, G. (1999). Experimental economics from the vantage point of behavioral economics. *Economic Journal*.
- Mannix, B. (1991). Resource dilemma and discount rates in organizational decision making groups. *Journal of Experimental Social Psychology*, **27**, 379-391.

- McGrath, J. (1964). *Social psychology: A brief introduction*. New York: Holt, Rinehart and Whinston.
- Miller, G. (1992). *Managerial dilemmas: The political economy of hierarchy*. Cambridge, MA: Cambridge University Press.
- Murnighan, J. K. (1994). Game theory in organizational behavior. In B. Staw & L. Cummings (Eds.), *Research in organizational behavior (Vol. 16)* (pp. 83-123).
- Murnighan, J. K., & Roth, A. (1983). Expecting continued play in prisoner's dilemma games: A test of several hypotheses. *Journal of Conflict Resolution*, **27**, 279-300.
- Pillutla, M. M., & Chen, X. P. (1999). Social norms and cooperation in PDs: The effects of context and feedback. *Organizational Behavior and Human Decision Processes*, **78**, 81-103.
- Sally, D. F. (1995). Conversation and cooperation in social dilemmas: Experimental evidence from 1958 to 1992. *Rationality and Society*, **7**, 58-92.
- Schelling, T. (1960). *The strategy of conflict*. Cambridge, MA: Harvard University Press.
- Seely, B., Van Huyck, J. B., & Battalio, R. B. (1999). Credible assignments and coordination failure in laboratory public goods games. Working paper. Texas A&M Department of Economics.
- Singley, K., Anderson, J. R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press.
- Steiner, I. D. (1972). *Group process and productivity*. New York: Academic Press.
- Stern, P. C. (1976). Effect of incentives and education on resource conservation decisions in a simulated commons dilemma. *Journal of Personality and Social Psychology*, **34**, 1285-1292.
- Van Huyck, J. B., Battalio, R. B., & Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, **80**, 234-248.
- Van Huyck, J. B., Battalio, R. B., & Beil, R. O. (1991). Strategic uncertainty, equilibrium selection principles, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, **CVI**, 885-910.

- Van Huyck, J. B., Gillette, A. B., & Battalio, R. B. (1992). Credible assignments in coordination games. *Games and Economic Behavior*, **4**, 606-626.
- Warglien, M., Devetag, M. G., & Legrenzi, P. (1999). I modelli mentali dei giochi: Focalizzazione e rappresentazioni selettive. *Sistemi Intelligenti*, **11**, 85-112.
- Weber, R. (1999). Information and strategic behavior in experimental games. Unpublished Ph.D. thesis, Pasadena, California: California Institute of Technology, Division of the Humanities and Social Sciences.
- Weber, R., Camerer, C. F., Rottenstreich, Y., & Knez, M. (in press). The illusion of leadership in weak-link coordination games. *Organizational Science*.

Authors Note

Helpful comments were received from Ron Burt, James Coleman, Chip Heath, Rod Kramer, one referee and both editors, participants at the Stanford GSB conference on Strategy and Organization, October 1994, and participants at the University of Chicago GSB Social Organization of Competition Workshop, January 24, 1995.

Table 1a: Prisoners' Dilemma

		Column	
		C	D
Row	C	5, 5	0, 8
	D	8, 0	2, 2

Table 1b: Coordination Game

		Column	
		C	D
Row	C	5, 5	0, 2
	D	2, 2	0, 2

Table 2: Weak-Link Game

MINIMUM VALUE OF X CHOSEN BY
OTHERS

		7	6	5	4	3	2	1
YOUR CHOICE OF X	7	1.30	1.10	0.90	0.70	0.50	0.30	0.10
	6	1.20	1.20	1.00	0.80	0.60	0.40	0.20
	5	1.10	1.10	1.10	0.90	0.70	0.50	0.30
	4	1.00	1.00	1.00	1.00	0.80	0.60	0.40
	3	.90	.90	.90	.90	0.90	0.70	0.50
	2	.80	.80	.80	.80	.80	0.80	0.60
	1	.70	.70	.70	.70	.70	.70	0.70

Table 3a: 2-Player PD Game with 7 actions

		Other Player's Choice of X						
		1	2	3	4	5	6	7
Your Choice of X	1	1.5	1.7	1.9	2.1	2.3	2.5	2.7
	2	1.4	1.6	1.8	2	2.2	2.4	2.6
	3	1.3	1.5	1.7	1.9	2.1	2.3	2.5
	4	1.2	1.4	1.6	1.8	2	2.2	2.4
	5	1.1	1.3	1.5	1.7	1.9	2.1	2.3
	6	1	1.2	1.4	1.6	1.8	2	2.2
	7	0.9	1.1	1.3	1.5	1.7	1.9	2.1

Table 3b: 2-Player PD Game with 3 actions

		Other Player's Choice of X		
		1	2	3
Your Choice of X	1	1.5	2.1	2.7
	2	1.2	1.8	2.4
	3	0.9	1.5	2.1

Table 4: Experimental Design – 2-player experiments

Session	# of Subjects	Treatment	Round 1	Round 2
1	10	WP(7)	Weak-link	7-action PD
2	14	WP(7)	Weak-link	7-action PD
3	10	PP(7)	7-action PD	7-action PD
4	10	PP(7)	7-action PD	7-action PD
5	10	WP(3)	Weak-link	3-action PD
6	12	WP(3)	Weak-link	3-action PD
7	12	PP(3)	3-action PD	3-action PD
8	12	PP(3)	3-action PD	3-action PD

Table 5: Total numbers of choices of actions 1-7 or 1-3

Treatment	<u>Round 1</u>					<u>Round 2</u>				
	Period					Period				
	1	2	3	4	5	1	2	3	4	5
WP(7)	Weak – link game					7-action PD				
No. of 7's	22	24	24	24	24	17 (20)	14	14	16	9
No. of 6's	0	0	0	0	0	0 (0)	0	0	0	0
No. of 5's	2	0	0	0	0	0 (0)	0	0	0	0
No. of 4's	0	0	0	0	0	1 (1)	1	1	1	0
No. of 3's	0	0	0	0	0	0 (0)	0	0	1	0
No. of 2's	0	0	0	0	0	0 (0)	0	0	0	0
No. of 1's	0	0	0	0	0	6 (3)	9	9	6	15
PP(7)	7-action PD					7-action PD				
No. of 7's	3 (7)	2	2	3	2	6 (10)	4	4	4	2
No. of 6's	4 (2)	1	1	2	0	1 (0)	0	0	1	0
No. of 5's	0 (0)	1	0	0	0	0 (0)	0	0	0	0
No. of 4's	2 (5)	2	1	0	1	3 (0)	2	2	2	0
No. of 3's	0 (1)	1	3	0	1	2 (1)	3	1	1	2
No. of 2's	2 (1)	0	2	4	1	2 (0)	1	1	2	0
No. of 1's	9 (4)	13	11	11	15	6 (6)	10	12	10	16
	Weak – link game					3-action PD				
No. of 7's	16	19	20	0	21					
No. of 6's	1	0	0	0	0					
No. of 5's	3	2	2	2	2					
No. of 4's	1	0	0	0	1					
No. of 3's	0	1	0	0	0					
No. of 3's	0	1	0	0	0	11 (13)	11	12	10	4
No. of 2's	0	0	0	0	0	1 (3)	3	3	3	2
No. of 1's	1	0	0	0	0	10 (6)	8	7	9	16
PP(3)	3-action PD					3-action PD				
No. of 3's	7(11)	7	7	7	6	10(15)	12	14	15	9
No. of 2's	3(2)	4	3	2	1	4(2)	4	6	5	4
No. of 1's	14(11)	13	14	15	17	10(7)	8	4	4	11

Note: Entries in parentheses are the number of times a subject guessed that the other player would select that action

Table 6: Two-sample Kolmogorov-Smirnov test statistics for treatments effects on number of cooperative actions selected.

7-action S games	PP(7:2.1)	WP(7:2.1)
WP(7:2.1)	0.41 (.05)	
PP(7:1.1)	0.15 (n.s.)	0.56 (.005)
3-action S games	PP(3:2.1)	WP(3:2.1)
WP(3:2.1)	0.08 (n.s.)	
PP(3:1.1)	0.17 (n.s.)	0.21 (n.s.)

Table 7: 3-Player PD Game

		Your Choice of X						
		1	2	3	4	5	6	7
	2	1.5	1.4	1.3	1.2	1.1	1	.9
	3	1.6	1.5	1.4	1.3	1.2	1.1	1
Sum of	4	1.7	1.6	1.5	1.4	1.3	1.2	1.1
X's	5	1.8	1.7	1.6	1.5	1.4	1.3	1.2
Selected	6	1.9	1.8	1.7	1.6	1.5	1.4	1.3
By Other	7	2	1.9	1.8	1.7	1.6	1.5	1.4
Two	8	2.1	2	1.9	1.8	1.7	1.6	1.5
Players	9	2.2	2.1	2	1.9	1.8	1.7	1.6
	10	2.3	2.2	2.1	2	1.9	1.8	1.7
	11	2.4	2.3	2.2	2.1	2	1.9	1.8
	12	2.5	2.4	2.3	2.2	2.1	2	1.9
	13	2.6	2.5	2.4	2.3	2.2	2.1	2
	14	2.7	2.6	2.5	2.4	2.3	2.2	2.1

Table 8a: Results of 3WP(7) treatment (sessions 1-2 pooled), experiment 2

Treatment	<u>Round 1</u> Period					<u>Round 2</u> period				
	1	2	3	4	5	1	2	3	4	5
3WP(7)	3-player weak-link game					7-action PD				
No. of 7's	12	8	9	9	7	6	3	1	0	0
No. of 6's	1	5	4	3	2	1	1	2	0	1
No. of 5's	3	2	3	1	1	4	3	1	0	0
No. of 4's	2	3	0	2	1	5	4	3	1	1
No. of 3's	2	3	2	1	1	0	0	1	2	0
No. of 2's	0	0	2	3	3	0	1	0	2	1
No. of 1's	1	0	1	2	6	5	9	13	16	18

Table 8b: Results of 2,3WP(7) treatment, experiment 2

Period	<u>Round 1</u>		<u>Round 2</u>		<u>Round 3</u>				
	1	5	1	5	1	2	3	4	5
Game	2-player weak-link		3-player weak-link		3-player, 7-action PD				
No. of 7's	6	8	12	12	9	8	6	6	3
No. of 6's	0	0	0	0	0	0	0	0	0
No. of 5's	1	0	0	0	0	0	0	0	0
No. of 4's	1	0	0	0	0	0	0	0	0
No. of 3's	0	0	0	0	0	0	0	0	0
No. of 2's	0	0	0	0	0	0	0	0	0
No. of 1's	0	0	0	0	3	4	6	6	9