Team versus Individual Play in Finitely Repeated Prisoner Dilemma Games

John H. Kagel

Peter McGee*

Abstract

In finitely repeated prisoner dilemma games, two-person teams start with significantly less cooperation than individuals, consistent with results from the psychology literature. This quickly gives way to teams cooperating more than individuals. Team dialogues show increased payoffs from cooperation, along with anticipating opponents' recognition of same, provides the basis for cooperation, even while fully anticipating defection near the end game. A strong status quo bias in defecting across super-games limits unraveling. Defecting typically occurs' one round earlier across super-games, consistent with low marginal, or even negative, benefits of more than one-step-ahead defection. (JEL D03, C92, C73)

*Kagel: Department of Economics, Ohio State University, 410 Arps Hall, 1945 High Street, Columbus, OH 43210 (email: kagel.4@osu.edu) McGee: Department of Economics, University of Arkansas, Business Building 402, Fayetteville, AR 72704 (email: pmcgee@walton.uark.edu). We have benefited from discussions with Guillaume Fréchette, P. J. Healy, Chester Insko and Tim Wildschut, as well as comments at talks at Ohio State University, New York University, the 2014 Decision Theory, Experiments and Applications workshop, the 2014 Texas Experimental Association meetings and the 2014 North American ESA meetings as well as a referee. Special thanks to Jim Peck for helpful discussions. Valuable research assistance was provided by Xi Qu, Matt Jones, Dimitry Mezhvinsky, and Andrzej Baranski. This research has been partially supported by National Science Foundation grant SES-1226460. Opinions, findings, conclusions or recommendations offered here are those of the authors and do not necessarily reflect the views of the National Science Foundation.

We report the results of an experiment comparing the behavior of two person teams in a finitely repeated prisoner's dilemma game (FRPD) with individuals playing the same game. There are two motivations for the experiment. First, given the high frequency with which economic decisions are made by teams as opposed to individuals, there is interest in what, if any, differences will emerge between the two. Second, and of equal interest, is to use the team dialogues to better understand the basis for the typical pattern of play reported – cooperation in early stage games followed by defection as the end game draws near, with very limited unraveling over time.

The experiment sits at the intersection of two strands of research – social psychology experiments that compare team versus individual play in FRPD games, and economic experiments designed to investigate models aimed at explaining the typical pattern of play in these games. The main finding from the social psychology experiments is that teams are significantly less cooperative than individuals, a result referred to as the "discontinuity effect" (see Wildschut et al., 2003 and Wildschut and Insko, 2007, for surveys of the literature). This is reported for two, three, and four person, financially incentivized teams, and is accentuated when there are face-to-face discussions between teams before deciding what actions to take. Explanations for the discontinuity effect are offered in terms of greater fear, distrust and greed in inter-group relations than in inter-individual relations, or that group discussions facilitate more rational comprehension of the situation, thereby leading to superior backward induction and less cooperative play.[1] We expand on these explanations when reviewing the social psychology literature on the discontinuity effect.

[1] Results from this literature have been used to argue that groups are likely to reduce social welfare in the prisoner dilemma game (Charness and Sutter, 2012). As will be shown, this conclusion seems premature.

Our experimental results, which do not involve communication with opponents, are consistent with the discontinuity effect in that there is significantly less cooperation for teams in the first super-game. The team dialogues show that this is largely the result of "safety" considerations, namely fear of getting the "sucker" payoff because the other team defects. However, after the first super-game, teams cooperate at the same or higher levels than individuals as early concerns with "safety" give way to a willingness to take some risks in order to earn the higher profits from cooperation. There is no corresponding social psychology data that we are aware of to compare this last result, as the psychology experiments typically involve a single super-game.[2]

Economic experiments have focused on evaluating the Kreps et al. (1982) model of behavior in FRPD games. This elegant model shows that if perfectly rational agents believe that there are sufficient numbers of conditionally cooperative types in the population ("crazy" types), it is in their best interest in early stage games to play cooperatively, only to defect as the end game draws near. The percentage of conditionally cooperative types needed to support this model can be surprisingly small; in fact, the model does not even require that there actually be any conditionally cooperative types in the population, as actions are driven by beliefs (Reny, 1992). This argument serves to rationalize, at least qualitatively, the typical pattern of play in FRPD experiments after subjects have gained some experience. There have been a large number of experiments investigating more detailed predictions of this model with mixed results, discussed in the next section of the paper.[3]

[2] This is based on personal communication with Professors Tim Wildschut and Chester Insko.

[3] There are a number of possible variations on the possible "crazy" types that yield different detailed predictions from those developed in Kreps et al. These have not been investigated in any great detail.

Our experimental results are consistent with strong elements of rational play defined as maximizing expected payoffs subject to players' beliefs, but are not consistent with the particular mechanism underlying the Kreps et al. model. In contrast to the static nature of the Kreps et al model there is strong growth in early stage game cooperation across super-games. The team dialogues make it clear that this is driven by the increased payoffs from cooperation, in conjunction with anticipation that their opponents are likely to reciprocate, and the low cost of these cooperative overtures should they be ignored (i.e., defection in subsequent stage games). That is, unlike the original Kreps et al. model, agents beliefs regarding what their opponents are likely to do are not fixed, but rather change between super-games based on their past experience and prior beliefs. Further, there are a number of cooperative patterns inconsistent with the Kreps et al. reputation model, particularly in the first several super-games for teams. Unraveling of cooperation across super-games is limited by a strong no change bias in when to defect, along with, typically, one-stage-game-ahead early defection when it occurs. The team dialogues indicate that this reflects limited backward induction in determining when to defect, failing to account for others adjusting in the same way, consistent with the strong status quo bias in when to defect across super-games. Simulations show this is supported by low (or even negative) marginal returns for defecting two rounds earlier than the last time on a cooperative path. The net result is that for both teams and individuals' starting playing cooperatively at the start of a super-game consistently yields higher average earnings than not cooperating.

Overall patterns of play are similar between teams and individuals with some differences in terms of the detailed pattern of play: Both teams and individuals defect one period earlier over seventy percent of the time when defecting earlier than the last time they were on a cooperative path. Regression results show that the same factors impact stage one cooperation rates between the two. Both teams and individuals occasionally cooperate in the last stage game, but teams do so significantly less often, with the team dialogues indicating cooperation in this case results from mistakes, confusion, or naiveté.

The remainder of the paper is organized as follows: Section I reviews prior research on FRPD games that provides the background for the issues explored here. Section II outlines our experimental design and procedures. Section III reports the experimental results in relationship to the issues raised in Section I. This has two distinct parts – comparing individuals with teams and using the team dialogues to gain insight into subjects' beliefs driving their behavior. Section IV briefly summarizes our main results and conclusions.

I. Prior Research:

There has been much work done on FRPD games in both the economics and social psychology literature. The goal in this section is not to provide an exhaustive review of the literature, but to summarize results from papers most closely related to the work reported here. Within economics, the major puzzle is to explain why these finitely repeated games do not completely unravel, but rather consistently show early stage game cooperation. Within the social psychology literature much of the focus has been on the "discontinuity effect", the fact that teams tend to cooperate less than individuals.

The evidence for teams cooperating less than individuals is strongest when communication between opponents is permitted, along with within-group discussions (Wildschut et al., 2003; Wildschut and Insko, 2007).[4] Typical procedures here are to first have within-team discussions, followed by discussions between representatives of each team, followed by teams independently deciding on what to do, with corresponding procedures for individuals.[5] A number of

---

[4] Note the discontinuity effect is by no means limited to experimental designs involving inter-party discussions.

[5] Much of this research has involved financially incentivized agents.

clever experimental designs have been employed to try and tease out the reasons why teams are more competitive. Within that literature, there are two competing explanations for why teams are less cooperative. One explanation is that intergroup relations are characterized by greater fear and greed than inter-individual relations leading to less cooperative play. The second perspective is that group discussion facilitates rational comprehension of the forces at work in mixed motive situations like FRPD games, with the greater rational comprehension favoring greater backward induction (hence less cooperation) on the part of teams.[6]

One of the shortcomings of this literature, from an economist's perspective, is that these experiments have typically involved a single super-game between a pair of agents, as opposed to the typical economic experiment where agents engage in a number of super-games, and are re-matched following each super-game.[7] Among other things, this means that there has been no investigation of whether the "discontinuity effect" will persist over time. By randomly re-matching agents between super-games within a given experimental session, we are able to investigate this, as well as whether or not cooperation unravels faster over time for teams as opposed to individuals.

Much of the economics literature has focused on investigating different formal models rationalizing the typical pattern of play reported in FRPD games: An initial period of cooperation followed by cooperation breaking down near the end of each super-game (Selten and Stoecker, 1986, Andreoni and Miller, 1993). With fully rational, own income maximizing agents and common knowledge of

---

[6] We have found no attempt to sort out between these explanations based on team discussions. This may be because they are typically free form and face-to-face as opposed to the more restrictive procedures employed here.
[7] The one-shot super-game nature of most of this research is not obvious from reading the surveys on the discontinuity effect. We are thankful to Professors Tim Wildschut and Chester Insko for clarifying this.

rationality, the standard backward induction argument predicts defection in each stage game, yet this is rarely observed. A number of alternative explanations for early stage cooperation have been discussed in the literature (Bicchieri, 1989; Reny, 1992). The Kreps et al. (1982) reputation model has provided the focus for much of the experimental research. In this model, if there is incomplete information about the types of players one is likely to face, with a high enough probability that some of these agents will be committed tit-for-tat (TFT) players, then cooperation in early plays of the game is consistent with fully rational behavior, along with defection as the end game draws near. The model is static, failing to account for the fact that early stage game cooperation often increases at first across super-games, followed by some unraveling (earlier defection) in later super-games.

Early experiments provided qualitative support for the Kreps et al. model: Andreoni and Miller (1993) compared FRPD games with one-shot PD games, reporting substantially more cooperation in early stage play in the FRPD games, and close to the same level of end game cooperation as in the one-shot games, consistent with agents believing there are "altruistic types" in the population. In addition, FRPD games in which there was a 50% chance of playing against a computer playing TFT (where the TFT strategy was announced and explained to subjects) resulted in cooperation being sustained at very high levels for substantially more early stage games than with all human competitors.[8]

Subsequent experiments have not provided such strong support for the model. Cooper et al. (1996) report results for a 10-period FRPD game which exhibits the typical pattern of early cooperation followed by defection, along with higher cooperation rates in early play than in a one-shot game. While this aggregate pattern of play is qualitatively consistent with the Kreps et al. model,

[8] Also see Camerer and Weigelt (1988) who report strong support for the Kreps et al. reputation building model in the context of a borrower-lender game.

using more detailed data they report contrary evidence: (i) at times cooperative play follows non-cooperative play by the same player or his opponent, which should not occur, and (ii) the high levels of early cooperation observed require substantially higher levels of conditionally cooperative types than found in end period play.[9] Cooper at al. note that the inability of their data to fit the Kreps et al. model better could reflect the specific type of "irrationality" underlying it, as a substantially wider variety of equilibria can result from alternative "irrational" types (alternative types often addressed in the infinitely repeated game literature, e.g., Fudenberg and Maskin, 1986), and/or behavior that is driven in part by mistakes on the part of players.[10]

Selten and Stoecker (1986) focus on learning in a simultaneous move FRPD game, with the focus on the unraveling in end game play. To investigate this issue, they employ a Markov learning model where subjects change their intention to deviate from cooperation depending on their experience in the previous super-game. Subjects were asked to write down reasons for each period's decision, with these descriptions, in conjunction with observed patterns of play, used to determine the period in which subjects intended to defect. In their model, defection between super-games either does not occur, or shifts one period earlier or later than in the previous super-game. Defection is likely to occur earlier if a player's opponent deviated earlier than the player intended to, or deviated in the same period the player intended to, with increased likelihood of defecting earlier if the player's opponent deviated before he did. Defection is likely to occur one period *later* if a player defected before their opponent did in

---

[9] Similarly, Jung et al. (1994) look at the Kreps et al. reputation building model for limit-entry pricing in the context of the chain-store paradox, finding substantially weaker support for the model than Camerer and Weigelt report.

[10] Others reporting results contrary to Kreps at al. are Cox et al. (2012), who study a finitely repeated sequential PD game, and Reuben and Suetens (2012), who study a finitely repeated sequential PD game with an uncertain end point.

the previous super-game. They find strong support at the individual subject level for their learning model in later super-games.

The Selten and Stoecker model is essentially one of bounded rationality. There are other bounded rationality models applied to FRPD games. Bereby-Meyer and Roth (2006) report an experiment in which the speed of adjustment to the mature pattern of play reported in noisy FRPD games is captured by reinforcement learning models from psychology. Jeheil (2005) develops a boundedly rational model in which agents establish equivalence classes across rounds for when their opponent is likely to defect and best respond to this. In this model, agents always defect in the last stage game and, like the Selten and Stoecker model, there is no characterization of the underlying behavioral forces that drive the typical learning process. Neyman (1985) develops a model of cooperation in FRPD games under the assumption that there are bounds to the complexity of the strategies that players can use.

Our experiment takes place at the intersection of the social psychology and economics literatures. To our knowledge it is the first study of team play in finitely repeated PD games with several super-games against different opponents. This can help to determine if less cooperative play on the part of teams (the "discontinuity" effect) persists with experience. The team chats provide a natural way of obtaining insights into the beliefs underlying the typical pattern of play in these games. Further, from the point of view of whether teams are more "rational" than individuals, we compare the frequency of end game cooperation between the two, and whether unraveling occurs faster for teams.

II. Experimental Design and Procedures:

Subjects played a ten stage, simultaneous move, FRPD with stage-game payoffs reported in Figure 1. Payoffs were denominated in experimental currency units (ECUs) which were converted into dollars at the rate of $1 = 250 ECUs. Payoffs were computed over all plays of all the super-games and paid in cash at the end of an experimental session along with a $6.00 participation fee. Each member of a team received his team's payoff.

[Insert Figure 1 here]

In the team treatment, subjects were randomly matched with a partner at the beginning of an experimental session, with partners remaining the same throughout the session. Teams played against teams, and individuals played against individuals. In what follows we will use the term agent to refer to either a two person team or an individual. Following each FRPD game, agents were randomly and anonymously re-matched under the restriction that no two agents would be re-matched in consecutive super-games. Teams in sessions1-3 each played seven FRPD super-games. This was increased to nine and ten FRPD super-games in sessions 4 and 5, as it was clear there was sufficient time for the extra games. All five individual subject sessions played ten FRPD games. Each session had between 8 and 12 individuals/teams for a total of 52 individual subjects and 51 teams.[11] Agents in both treatments were told they would play between 7 and 10 super-games.

Teams had 3 minutes to discuss and make their choices in the first two rounds of each super-game. This was reduced to 1.5 minutes after that. Default options if time ran out without a coordinated choice are enumerated in the instructions, which can be found in the online appendix.[12] Similar time limits were imposed for individual play, but these were never binding.

[11] One team session used a student assistant to ensure an even number of teams. The assistant informed his teammate that he was one of the experimenters and would agree to whatever his partner did. He also asked his partner to write out any thoughts he/she had about the game in the chat box. Data for this team is dropped except as needed to complete play when paired with another team.

[12] Overall, 97.5% of all team choices involved active coordination between teammates on choices made. The instructions can be found at http://sites.google.com/site/econpjmcgee/AppendixKM.pdf.

Following the end of each stage game agents had up to 30 seconds to view the results before moving on to the next stage game. Following the last stage game, agents were notified that their match had ended and that they would start another match with another randomly chosen agent. Neutral language was used throughout; e.g. agents chose between option A or B in each stage game, and were told they would be "paired with the same other team (individual) for a set of 10 repeated choices."

III. Experimental Results:

Results are reported in two parts, first comparing patterns of play between individuals and teams, making use of the team chats as needed. Second, further analyzing the team chats to better understand the beliefs underlying the behavior reported, and to sort out between explanations for the behavior. The analysis is limited to the seven super-games common to all sessions.

III .A Comparing patterns of play between individuals and teams

Figure 2 reports average levels of cooperation for teams and individuals over the seven super-games. The data exhibits the usual pattern in both cases with cooperation rates at their peak in the early stage games followed by a rather precipitous drop as the end stage draws near.

Table 1 reports average stage one cooperation rate for each super-game along with z-statistics for differences between the two treatments. The focus is on stage-one cooperation rates, as cooperation in later rounds is very much dependent on what happens in the first stage game, which creates complicated interdependencies that are difficult to account for. Further, once two or more stage games have passed in which one agent has defected, in the overwhelming number of cases both agents defect for the remainder of the super-game.

Average stage one cooperation rates are significantly higher for individuals in the first super-game. However, by the second super-game the rates are essentially the same, with teams having higher cooperation rates in the remaining super-games, a difference that is statistically significant in super-games 5 and 6.

[Insert Figure 2 and Table 1 here]

Appendix A1 reports the results of a probit investigating the factors behind cooperation in the first stage game across super-games, employing variables shown to impact first round play in infinitely repeated super-games (Dál Bo and Fréchette, 2012). Key results are that (i) cooperation in round 1 of the first super-game measures an inherent tendency to cooperate that carries over to later games, (ii) agents are more likely to cooperate if the agent they were paired with in the previous super-game cooperated, and (iii) the main effect for a team treatment dummy shows higher overall rates of cooperation ($p < 0.10$). The first two characteristics hold equally for teams and individuals, i.e., interaction effects between these two variables and teams are not statistically significant at conventional levels. Factors (i) and (ii) have the same strong impact in infinitely repeated super-games (Dál Bo and Fréchette, 2012) as found here.

Conclusion 1: Consistent with the discontinuity effect reported in the social psychology literature teams are less cooperative than individuals in the first super-game. However, they are as, or more, cooperative than individuals in later super-games, so that overall teams are more cooperative than individuals.

Fully rational, own income maximizing agents should never cooperate in the last stage game as it is a dominant strategy to defect. In contrast to this, there is some cooperation for both teams (9.8%) and individuals (26.9%), a difference that is significant at the 5% level.[13] Cooperation in the last round is sometimes treated as evidence for altruism, typically reciprocal altruists committed to TFT

[13] These percentages report the frequency with which an agent defected in the end game for one or more super-games. Agents who cooperated more than one time are counted once in the data. No team cooperated more than once, but one individual cooperated twice and another cooperated three times.

(Andreoni and Miller, 1993; Cox et al., 2012). However, the team chats suggest that it results from mistakes, confusion or naiveté. For example, one team had decided to defect in the next to last stage game in order to earn the higher payoff ("…we get 175? we won't ever play them again"). But one member of the team mistakenly chose to cooperate (the other chose to defect), with no time to correct the mistake, with the computer selecting the cooperative player's choice. Or to take another case: After choosing to defect in 8 out of 9 rounds, one team chose to cooperate in the end game "just for the hell of it". That mistakes, confusion or naiveté account for most of the end game cooperation is also supported by the fact that over half of these occur in the first two super-games for both teams and individuals. Finally, note that for these agents who cooperated in the last stage game, there is at least one other super-game where that same agent unilaterally defected, or defected simultaneously with their opponent. The latter is inconsistent with committed TFT types.

Given that end game cooperation is more than likely driven by mistakes, confusion or naiveté teams, in having two-heads to work with, should have lower rates of cooperation, to the point that they meet or beat the truth wins (TW) norm (Lorge and Salomon, 1955). The TW norm holds that for problems that have a clear, correct answer which can be easily explained, a team should do as well or better than the best member of that team, as the one with the correct answer is able to explain the solution to her teammate.[14] This is investigated with a simulation randomly matching decisions by individual agents in the subject population, and determining the frequency with which at least one of the two chose not to cooperate in the last stage game, making that the synthetic team's

[14] Note that the psychology research on this issue shows that teams rarely meet, much less beat, the TW norm (Davis, 1992). Investigations of team versus individual behavior in economics rarely address this question, in part because in most cases the insight needed to solve the problem is sufficiently complicated that it would be quite difficult to explain the solution to one's partner(s).

choice. The simulation predicts that on average between 5.9% and 7.8% of the teams would have cooperated under the TW, with the observed cooperation rate (9.8%), well within the 90% confidence interval for the TW.[15]

Conclusion 2: There is significantly less cooperation in the last stage game for teams than for individuals. The team chats indicate that cooperation resulted from mistakes, confusion, or naiveté, which no doubt holds for individuals as well. Simulations show that the lower end stage cooperation for teams falls well within the 90% confidence interval of the demanding truth wins norm. This is indicative of greater "rationality", and/or clarity of thought, on the part of teams as there is little evidence for committed conditional cooperators, or altruists, in the subject population.

Teams unravel a bit more and faster than individuals. This is measured by the round in which an agent defects, conditional on being on a cooperative path at the start of a super-game. The latter is defined as sustained cooperation over rounds 1-4, typically with both agents cooperating in all rounds.[16] Table 2 reports the average number of defections in each super-game along with the round in which the defection occurred. In the first super-game, the median round for defections was 10 for individuals and 9 for teams ($p > 0.10$, Mann-Whitney test). For both teams and individuals, there is slow, and far from complete, unraveling across super-games, with the median for teams always one step ahead of individuals until the last super-game, where it is two steps ahead (round 7 versus round 9; $p < 0.01$, Mann-Whitney test). This greater unraveling for teams can be attributed to having more experience with defection as they get onto a cooperative path significantly more often than individuals, as well as starting out defecting

[15] The simulation consisted of samples of 51 teams, drawn from 52 individuals (14 cooperators; the rest defectors) with replacement and repeating the simulation 250,000 times. Simulated teams were counted as cooperating when both of the individuals drawn had cooperated.

[16] The exceptions to this criterion are discussed below.

earlier than individuals. Teams starting out defecting earlier more than likely results from their being better able to recognize the dominant strategy to defect in the last stage game, and applying one-step ahead reasoning to this.

[Insert Table 2 here]

Table 3 reports the extent to which agents responded to defection *between* super-games conditional on being on a cooperative path in rounds 1-4. There are three categories, listed at the top of the table for what agents were doing when the defection occurred in the previous super-game: (1) they were cooperating (Were Cooperating), (2) both defected at the same time (Both Defected), or (3) defecting when their opponent was cooperating (Unilateral Defection).[17] The percentages show how agents responded the next time they were on a cooperative path – defecting in an earlier stage game (Earlier), in the same stage game (Same), and in a later stage game (Later). Note that some of these observations are censored, in particular when an agent is defected on in an earlier stage game. However, for teams, the chats almost always indicate when the team was intending to defect, which is used to overcome the censoring problem. Also note that in almost all cases, once agents were on a cooperative path and a defection occurred, both agents defected for the remaining stage games.

[Insert Table 3 here]

Several things stand out. First, the infrequency with which agents defected earlier regardless of how their previous cooperative path ended, with Same (no change) as large or larger in all cases, consistent with a strong status quo bias (Samuelson and Zeckhauser, 1988). Defecting earlier is strongest when both agents defect simultaneously, as opposed to having been defected on while cooperating in the previous super-game. This is in direct contradiction to Selten and Stocker's (SS) adaptive learning model which predicts the same or greater

[17] Category (1) excludes mutual cooperation in the last stage game.

frequency of defecting earlier when defected on while cooperating in the previous super-game.[18] Further, defecting later than in the previous super-game is strongest following agents being defected on while cooperating, in contrast to the SS learning model which predicts that Later will be most common following unilateral defection. However, defecting later in this case is partly based on our accounting procedures in conjunction with the greater profits to be had from cooperation, as the examples in this footnote illustrate.[19]

Earlier defections typically occurred one round earlier than the last time an agent was on a cooperative path: 70.0% and 77.8% for teams and individuals, respectively. All but one of the remaining defections was two rounds earlier than in the previous super-game.[20] The team dialogues show that the one stage game ahead defections resulted from focusing on best responding to what happened the last time on a cooperative path, failing to consider the fact that other teams were doing the same. For example:

> 5: ok so next time i think we should try B on turn 9 in the same situation
> 7: yeah I was thinking about that
> 5: since the previous two times they had B for the last one anyway
> 7: right
> 7: so we'd gain 70 on turn 9

---

[18] Part of this may be due to learning (or lack of it) as SS had many more super-games. Also it is not clear if SS conditioned on being on a cooperative path, as is done here.

[19] For example, one team successfully defected in round 9 in two earlier super-games and planned to defect in round 8 the next time on a cooperative path. However, while cooperating they were defected on in stage 6, at which point they briefly discuss how their opponents were "jerks" for cutting off cooperation so early. In the next super-game they continued to plan to defect in round 8 which, being later than 6 is counted as Later. Or take another team who was defected on in round 7 when they had planned on defecting in round 8: "man people are starting to be greedy *@#*" "doing it in the 7th round doesnt even get u any extra money" "people are just stupid" (all quotations are direct, including typing and grammatical errors). This team continued to plan to defect in round 8, rather than earlier, but were defected on prior to round 8 in the next super-game.

[20] There was one instance of dropping 3 rounds earlier for a team. Table 3 does not include what happened following a super-game that ended with mutual cooperation (one time each for teams and individuals). For the teams, the next time they were on a cooperative path, they defected in rounds 7 and 8, respectively. For the individuals, both defected in round 9 the next time on a cooperative path.

.

Teams defecting two rounds earlier explicitly considered other teams also recognizing the need to defect earlier, and best responded to this.

> 10: chances are, they either got screwed over or screwed someone over on round 10 of last block which means they'll be thinking they should screw us over in the 9th which is why we whould go with B in the 8th

To the extent that explicit backward induction is at work, it is effectively limited to one or two-stage-ahead thinking.[21]

Conclusion 3: Conditional on being on a cooperative path there were minimal differences between teams and individuals in defection patterns. The fact that teams unraveled more than individuals by the last super-game is largely attributable to teams defecting earlier in the first super-game, as well as being on a cooperative path more often, with its increased opportunities for defection. The most common response to an immediate past defection is to not change the intended period of defection next time on a cooperative path, consistent with a status quo bias. When defecting earlier, agents typically defected one period ahead, with the team dialogues indicating a failure to consider the possibility that their opponents might be adjusting as well.

Given the limited frequency of defecting earlier than the last time on a cooperative path, along with the majority of defections occurring one round earlier, the question is how costly was this limited unraveling? That is, were there strong economic forces that should have induced teams to defect two or three or even more rounds earlier given that they were defected on in the previous super-game? To address this we ran simulations calculating the expected cost to a team's failure to defect one round earlier than in the previous super-game in

[21] There is a related phenomenon in the psychology literature where people when asked to judge themselves relative to a reference group focus on self-assessment (How good am I?) as opposed to other-assessments (How good are others?) (Windschitl, Kruger, and Simms (2003).

which they were defected on, or defected simultaneously with their opponent, and then repeated this for defecting two rounds earlier.[22] These expected values were calculated using the distribution of defections (other than their own) for the next super-game in which the team in question was on a cooperative path.[23] In calculating expected payoffs, we assumed that once a defection occurred there was mutual defection in all subsequent stage games. Expected payoffs were calculated over the last eight rounds of a super-game, with the calculations confined to teams, since censoring problems limit the analysis for individuals.[24]

Separate calculations were performed averaging over cases where there was mutual defection in the previous super-game (DD), and when defection occurred while a team was cooperating (CD) in the previous super-game.[25] The biggest gain for defecting one round earlier occurs following mutual defection in the previous super-game, with an average increase of 36.6 ECUs (5.0%), relative to making no adjustment in the actual (or intended) round of defection, with the marginal benefit for defecting two rounds earlier 12.4 ECUs (1.6%). When defected on while cooperating, average earnings increased 12 ECUs (1.6%) for defecting one round earlier in the subsequent super-game as opposed to making no adjustment, while the marginal benefit for defecting two rounds earlier *reduces* payoffs by 7.2 ECUs (0.9%). Failing to respond to a defection in the previous super-game was most costly for mutual defections (DD) in Round 10. In 7 out of 8 cases these teams defected earlier the next time they were on a cooperative path, indicating strong responsiveness to the cost of defecting in the last round, when

[22] These calculations are restricted to super-games in which the team in question got on a cooperative path (reported in Table 3).
[23] Here we are assuming that teams were updating their beliefs regarding how the population of potential opponents will play over time based on their own experience and prior beliefs, as is common in learning models.
[24] The analysis is for eight rounds because it is possible that, in defecting two rounds earlier than they actually did, cooperation could have ended after the second stage game.
[25] All of these calculations are conditional on being up on a cooperative path in both the previous and current super-games.

their opponent was almost sure to defect if they had not done so earlier. In defecting one round earlier, and defecting more often following DD than CD, teams at least captured the largest benefits to defecting earlier as opposed to holding to the status quo.

Conclusion 4: Simulations show the marginal benefit of two-round-ahead defections following mutual defection and being defected on while cooperating, as compared to one-round-ahead, is small or even negative, consistent with the high frequency (70.0% or more) of one-round-ahead defections. The fact that teams defect earlier more often following DD as opposed to CD (12 versus 5) is consistent with the greater opportunity costs of failing to defect earlier following DD. Teams are most responsive to simultaneous defection in the last round of the previous super-game where the opportunity costs of failing to defect earlier are greatest. These calculations demonstrate an *absence* of strong economic forces that would have promoted more substantial unraveling from one stage game to the next.

The analysis of defection rates is conditional on agents being on a cooperative path. The full set of cooperative paths underlying this analysis is shown in Table 4. Our initial definition for "on a cooperative path" required mutual cooperation for four or more rounds (P_{CC} in Table 4).[26] While this is by far the most consistent pattern overall reported in Table 4, the chats, identified a number of unorthodox patterns, particularly the first time a team got up on a cooperative path (reported under Other in Table 4): Teams planning to defect in the first stage game followed by cooperation in the second stage game and for a number of games thereafter if the team they were paired with cooperated in the first stage game. For example, here is a team discussing what they planned to do in the next super-game:

[26] The choice of four rounds here is, admittedly, arbitrary but seems natural under the circumstances and corresponds to the number of rounds employed in Selten and Stoecker (1986).

14: you want to do B (defect) again?
9: it's a new team
9: i dont know
9: but to be safe
9: better go with b right?
14: i think so, yes.
9: go with b first and see what the other team pick for the first round
14: if they choose A (cooperate)... that means they want to be nice... so round 2 we'll choose A to apologize

.

These teams were fully aware that by defecting in the first stage game they were likely to face punishment in order to get onto a cooperative path. But as the quote indicates, they adopted this strategy out of "safety" concerns.[27] There are other, even less traditional, patterns for getting on a cooperative path. In a handful of cases a team planned to cooperate two times before defecting for the remainder of the super-game if their opponent did not reciprocate. A few others alternated between defecting and cooperating for the first three stage games before settling down for a long sequence of mutual cooperation. These alternative patterns, most of which occurred early on, were counted as getting onto a cooperative path. They are completely outside the types of early stage cooperation typically anticipated in the economics literature. But for the initial defector they are fully rational given their beliefs and fears, as are the agents who, after punishing the initial defection, choose to cooperate once their opponent wanted to "be nice." While initially defecting with the hope of getting onto a cooperative path would not seem like a winning strategy, it involves the kind of experimentation and uncertainty regarding what others might do commonly reported in early rounds of economic experiments. There is a methodological aside to these alternative patterns as well: We had not thought about them in advance, which is one reason

[27] This is similar to, but not quite the same, as suspicious tit-for-tat as once the "score was even" by cooperating in the second stage game while their opponent defected, the defector in stage 1 was planning to cooperate in stage 3 and beyond provided their opponent reverted to cooperation (otherwise they would defect).

we prefer unstructured team talk, even though coding unstructured communication is quite labor intensive.[28]

Table 4 shows that these alternatives to the mutual cooperation over the first four stage games (P_{CC}) account for just over a third of the teams who first got on a cooperative path, but much less so for individuals.[29] However, for teams, these alternatives to the P_{CC} pattern were largely eliminated following the first-time onto a cooperative path.

Conclusion 5: There is a high initial frequency of unorthodox patterns for getting onto a cooperative path compared to what one might expect (e. g., Kreps et al., 1982), with a number of agents defecting in the first stage game only to get onto cooperative path after that. However, as the team chats make clear, these "crazy" types were looking to establish cooperation, but were worried about getting the sucker payoff. These patterns are fully compatible with expanded notions of how cooperation could "rationally" develop in FRPD games (Reny, 1992; Cooper et al., 1996). These unorthodox patterns are also consistent with the learning and adjustment process typically present in almost all experimental studies.

[Insert Table 4 here]

III.B Team Dialogues in Relationship to Behavior

Table 5 reports coding categories for team dialogues. There are three broad categories with a number of sub-categories. The broad categories were coded conditional on whether a team was cooperating or not, along with several categories regardless of cooperating or defecting. Coders could assign multiple codes to the same stage game. Two economics graduate students coded the

---

[28] In addition, unstructured dialogues have less potential for generating demand induced effects.

[29] Note there is some overlap in the initial frequency with which teams are counted in Table 4 because an agent who first cooperated in super-game *t* might be paired with an agent who first cooperated in a later super-game. A total of 10 (out of 51) teams and 13 (out of 52) individuals never got on a cooperative path.

dialogues. Categories were initially established by the authors after reading some of the dialogues. The coders then independently coded a single (common) session, after which they met with one of the authors to refine their common understanding of the categories. They then independently coded the rest of the sessions, after which there was a meeting to reconcile obvious discrepancies. Coders were in agreement 76% of the time over all sessions.[30] In the analysis that follows a coding is counted if either of the two coders assigned the code in question.

The goal behind the coding is to better understand the beliefs and strategies underlying teams' actions. Our assumption is that these beliefs and strategies are largely reflective of individuals as well. This assumption is based on the common elements in the probit estimates for first stage cooperation rates, along with the similarities in behavior reported in the previous section. The coding is imperfect and time-consuming, but to the extent that agents' beliefs are central to understanding behavior, they provide a natural way of tapping into these beliefs.

[Insert Table 5 here]

One key factor we wanted to identify was the basis for teams' decisions to cooperate or defect in the first super-game. Within the social psychology literature one explanation for teams low cooperation rates rests on opting for the "safest choice" (code D1 in Table 5).[31] That is, choosing to defect, in order to guarantee the payoff of 75 as opposed to the possibility of cooperating and getting the sucker payoff of 5. For round 1 of the first super-game, 91.7% (22/24) of the

[30] The same code assigned to a *different* round of the same super-game was counted as a disagreement, accounting for a number of the disagreements reported.

[31] As noted earlier, the social psychology literature on team play in PD games does not involve analysis of team discussions.

defecting teams were assigned code D1.[32] The following provides an example of one of these dialogues:

16: Pick B (defect) every time, yes?
1: what do you want to go with?
16: If we choose A (cooperate) we get 105 or 5
16: if we pick B we get 175 or 75
16: seems to me B is the choice in every situation… I don't want to jeopardize that minimum (earnings) with some 5 point takes

.

Of the teams cooperating in round one of super-game one, the most common coding was C1 - cooperating in order to elicit cooperation with its increased earnings. 70.6% (12/17).[33] An example of a C1 chat follows:

2: what do you think we should do
17: so i say pick a (cooperate)
2: ok thats fine. i hope the others arent greedy
17: bc that would give us higher payoff average
17: if we fall in the A-A zone
2: alright im game, lets do it

.

The increase in team cooperation rates beginning in super-game two is associated with teams recognizing the advantages of early round cooperation, either while stuck in mutual defection (code D5) or following a super-game in which they were stuck in mutual defection (code D1 in super-game *t* followed by C1 in game *t+1*). The following is an example of deciding to cooperate while stuck in a mutual defection sequence (D5):

18: B (defect) agian?
18: or do you want to lose money to get them to mutually choose a (cooperate)?
6: we'll ride b the rest of the way out this block but i think the best option is to go A the first 2 blocks, see if the other team catches on and if so choose A mutually for the remainder of that block.

.

[32] Six teams defected with no code recorded for round one.
[33] In 7 out of 12 of these cases, C1 and C2 (what to do in cases cooperation is not reciprocated) were coded simultaneously.

An example of a team reassessing its strategy after defecting throughout in the previous super-game on account of "safety concerns" (D1 followed by C1):

> 24: try A and see if we can get them on 105?
> 19: yeah let's do that[34]

.

In both cases teams start cooperating because of the mutual advantages of doing so, hoping the other team "catches on." And, as show below, teams recognized the relatively low cost of these efforts, planning to not cooperate in case the other team failed to reciprocate.

A second objective of the coding was to better understand the factors underlying defection patterns, conditional on being on a cooperative path. Although we had a code, X4, for complete unraveling, it was never assigned, as there were never any discussions approaching the full backward induction argument. Code X3 was designed to capture partial unraveling, discussions of whether and when to defect earlier than in previous super-games. As already noted, a significant insight from these dialogues is that in defecting one round earlier than in a previous super-game, teams typically did not consider that their rivals might be doing the same.

Under the X3 coding, there were occasions where teams identified the round in which they were likely to be defected on but failed to best respond:

> 2: round 8 is usually where we get screwed
> 17: TRUE
> 2: i kinda wanna go B (in round 8). that way if we do then we still get 75

.

[34] For an example of a team that partially stumbled on the cooperative strategy after defecting throughout in the previous super-game: "want to go A (cooperate) for fuN?" "ahh, Idk whatever let's do it" … "nice!" "that was interesting …now what" "if we can manage a monopoly with other team and stick to A we get more money out of it!"

Given the frequency with which PD games are discussed in introductory classes, across a number of social science classes, the dialogues were coded for cases where one or both members of the team had some prior experience with PD games of one sort or another (code X1). This was done regardless of whether the prior experience was from a PD experiment, or from classroom instruction, in an effort to see if our results were compromised by this prior knowledge. Eleven (out of 51) teams were coded as X1.

First, note that this prior experience was *not* always particularly helpful as the following case illustrates:

1: did you hear about gaming theory
16: no, what's that?
1: i guess that is about the same scenario
1: of this experiment
16: oh, ok. i've done this experiment before in sociology with the prisoner's dilemma and its the exact same thing. and youre always supposed to pick B (defect)
.

This team did not cooperate throughout super-game 3, turning to cooperate, successfully, in super-game 4 and continuing to do so through the end of their session.

Two questions were posed for teams with prior experience: Were they more or less cooperative to begin with and was their unraveling process, once on a cooperative path, materially different from those with no prior experience? Teams with prior experience were more cooperative in the first play of the first super-game, 54.5% versus 37.5%, but the difference is not significant ($p = 0.30$, two-tailed Z-test). With respect to X3 codes, 54.5% of teams with prior experience were coded X3 (discussing defecting earlier than in a previous super-game) versus 50.0% of those with no prior experience. Further, there was no material difference in the unraveling process between the two: The median for the *earliest* round in which a team defected, conditional on being on a cooperative path, was the same for those with and without prior experience (round 7).[35]

The dialogues were also used to determine if teams had a clearly articulated strategy for dealing with their rival's defection when attempting to establish cooperation in *early* stages of a super-game - namely defect themselves in subsequent rounds (codes C2 and D5). The following is representative of a C2 dialogue in round 1 of super-game 1:

2: what do you think we should do
17: so i say pick a (cooperate)
2: ok thats fine. i hope the others arent greedy
17: bc that would give us higher payoff average
17: if we fall in the A-A zone
…
2: if they choose b (defect) and we only get five what should we do next time
17: pick B for the next round and see what they choose this time around?
17: and if they pick B again, we stick with B i guess

.

A representative D5 dialogue in round 4 of super-game 1 follows:

23: B (defect) again?
7: yeah
23: lame
23: we could be making so much more if we all chose A (cooperate)
23: I say the next block we do 2 rounds of A
23: to see if the next group knows
23: and if not we'll just go back to B?

.

Conclusion 6: The cost of signaling cooperation was relatively low as by super-game two, 66.7% of teams were coded as C2 or D5, having clearly articulated strategies for defecting in case their opponent failed to cooperate when they did. By the end of super-game four, 86.2% of the teams were coded as C2 or D5, with

[35] This is restricted to cases where a team unilaterally defected or defected simultaneously with the team they were paired with, conditional on being on the cooperative path.

the remaining teams always defecting. No team was recorded as discussing cooperation regardless of what the other team did,

Category X5 was used to identify the first instance when the team dialogues expressed the belief that their opponent was either clearly going to defect, or very likely to defect, near the end of a super-game. This was correlated across super-games along with the percentage of teams on a cooperative path, generating a Pearson correlation coefficient of 0.85 ($p < 0.05$). That is, teams were increasing their early stage cooperation rates across super-games while simultaneously explicitly recognizing that defection was all but certain near the end game.

The growth in cooperative play across super-games, particularly for the teams, is inconsistent with the particular mechanism underlying the Kreps et al. model, with the more "rational" types cooperating to begin with out of beliefs that there are committed TFT types. That is, in the original Kreps et al. model agents have *fixed* beliefs about the distribution of "crazy" versus "rational" types. The changes in the level of cooperative play over time indicate that at a minimum agents have uncertainty about the distribution of types, which is changing over time based on their past experience and prior beliefs. This is not to denigrate the spirit of the model, or its importance in driving experimental and theoretical research in a large variety of finitely repeated games. The results reported show that the outcomes are driven by "crazy" types relative to common knowledge of rationality, as standard theory defines it, but the details of the mechanism are different.

Early stage cooperation here rests in part on a trial and error learning process motivated by the higher earnings that cooperation offers. The predominant pattern consists of agents who start cooperating in the first stage game to determine if their opponents are also interested in obtaining the greater benefits from mutual cooperation. If there is cooperation, they are willing to accept risky play akin to Russian roulette with respect to end game payoffs, as the payoff to mutual cooperation is more than enough to offset the occasional sucker payoff. This process results in a number of weird, early stage, patterns of play relative to what the Kreps et al. model would predict (recall Table 4 and Conclusion 5): Teams who fear non-cooperators to the point of defecting to begin while planning to cooperate after that if their opponent cooperated, fully recognizing that they must effectively pay back for their increased earnings from initially defecting. In response to this cooperators who faced defection in the first stage punished the defectors in stage two, but were often willing to cooperate with them if the defectors signaled their good intentions with cooperation in stage two. There are even occasional initial cooperators who were willing to give these initial defectors a second chance to positively respond to their willingness to cooperate.

The team discussions make it clear that these alternatives to the Kreps et al. pattern of early stage game cooperation are more than just noisy best responses, as teams are trying to figure out how to capture the benefits of mutual cooperation.

The net effect is that cooperation paid: Agents who started a super-game with cooperation received higher payoffs across all but the first super-game for teams, compared to starting with defection. These increased earnings averaged 42.6 ECUs per super-game for teams and 69.2 ECUs for individuals, increases of 5.2% and 8.2%, respectively.[36] The higher payoff for individuals in the presence of lower initial cooperation rates results from defection occurring later once on a cooperative path.

Conclusion 7: The particular mechanism underlying early stage game cooperation identified here is different from the one specified in Kreps et al., but is consistent with the wide range of behaviors recognized in the literature resulting in early

[36] The higher payoffs for individuals who start with cooperation hold even after dropping the first super-game: 70.0 ECUs (8.6%) versus 50.8 ECUs (6.1%).

round cooperation (e.g., Reny, 1992, Cooper et al., 1996). In the present case it is based on a trial and error learning process motivated by the increased earnings to be had from establishing early stage game cooperation, and the fact that signaling one's intention to cooperate results in increased average earnings. This is supported by slow, uneven and far from complete unraveling, in part due to a strong status quo effect in deciding when to defect, with what earlier defection there is across super-games typically one-stage-game ahead defection.

IV. Discussion

This paper investigates team versus individual play in finitely repeated prisoner dilemma games. Several results stand out. First, teams are less cooperative than individuals in the first super-game, consistent with the psychology literature on the so-called discontinuity effect. However, after the first super-game, teams are more cooperative than individuals cooperating at higher rates in the first stage game. An important difference between our experimental design and the typical psychology experiment is that psychologists employ a single match with the same competitor, whereas we employ repeated matching with different competitors. This is important because once cooperation breaks down (or never gets started) it is typically quite difficult to return to cooperating. The repeated matches provide the possibility for a reset and/or an effort to try a new, superior, strategy.

The results also have something to say with respect to the explanations offered in the psychology literature for the lower cooperation rates for teams. There are two, broadly defined, competing explanations (Wildschut et al., 2003; Wildschut and Insko, 2007): (1) Intergroup relations are characterized by greater fear and greed than inter-individual relations leading to less cooperative play and (2) group discussion facilitates rational comprehension of the forces at work in FRPD games, i.e., greater rational comprehension favoring greater backward induction (hence less cooperation) on the part of teams. Our results show that both factors are at work: In the first super-game, where the majority of teams start out defecting, non-cooperators are overwhelmingly coded as doing so out of "safety" concerns (91.7%). However, there is also support for the idea that team discussions facilitate rational comprehension, as teams are significantly more likely to play the dominant strategy–defection–in the last stage game (91.1% versus 73.1% for individuals).

The team dialogues are used to better understand the mechanism underlying the typical pattern of play in FRPD games – early stage cooperation followed by defection as the end game draws near. The dialogues indicate that teams signal their willingness to cooperate in the first stage game hoping the other team "catches on." The cost of signaling is relatively cheap and teams make plans to defect in subsequent stage games if cooperation is not reciprocated. Unraveling is slow and incomplete as there is a strong status quo bias in deciding when to defect across super-games and, when defecting earlier, it is typically one-stage-game-ahead defection, as teams fail to account for others learning and adjusting in the same way. Simulations show that this seemingly myopic strategy is sensible, as the marginal benefit from defecting two stage games ahead is small or even negative. Further, defecting earlier than in the previous super-game occurs most often when it is most beneficial, namely following simultaneous defection in the last stage game. The net effect is that following the first super-game, both teams and individuals who start with cooperation consistently earn higher payoffs than those who do not. Although individuals who start with cooperation consistently get higher average payoffs –as they unravel less– more teams get to enjoy these benefits.

A number of interesting questions remain to be explored in FRPD games using the team technology. First, teams appear to develop the mature pattern of play characteristic of FRPD games faster than individuals and unravel more than for individuals over time. As such, it would be interesting to see how far teams unravel with more experience than is reported here. Will they hit a stationary point or continue to the point of complete unraveling? No doubt this will take bringing experienced subjects back into the lab, or having subjects who are hardy enough to stay alert for a four hour session or longer. Second, given that teams start out cooperating less than individuals, only to cooperate more with a modicum of experience, it would be interesting to try to replicate these results using between-agent discussions, which generates the strongest discontinuity effect reported in the psychology literature.[37]

[37] See Kagel (2015) for a first pass at answering this question.

Appendix

Results from Round 1 Cooperation Rate Probits: The dependent variable takes the value 1 if an agent cooperated in round one of each super game. For explanatory variables we looked to the key elements characterizing first round play in Dál Bo and Fréchette (2012). Comparable forces are at work here: The initial cooperation rate dummy is equal to 1 if the agent cooperated in round 1 of super-game 1, included to capture agents' inherent tendency to cooperate or not. A team dummy is introduced to account for differences in cooperation rates between teams and individuals (value of 1 for teams). Super game is a linear time trend across super games. The immediate past experience dummy takes on a value of 1 if the agent they were paired with in previous super game cooperated.

Table A1 reports a specification without interaction effects between the teams dummy and the other explanatory variables. Both the initial cooperation rate and immediate past experience dummies are significant at the 1% level. The teams dummy is positive and significant at the 10% level, and the time trend variable is positive and significance at the 5% level. A specification including interaction terms between the teams dummy and each of the other right hand side variables showed none to be statistically significant on their own, at anything approaching conventional levels, with essentially no overall effect as well (χ^2 (3) = 0.55).

[Insert Table A1 here]

Further dialogue results: What follows is for a team that for two super-game they are best responding to what happened to them last time they were up on a cooperative path only to be defected on one round earlier: The first time this team gets up on a cooperative path they defect in round 10 where they are met with mutual defection. The next time they get up on a cooperative path they again chose to defect in round 10 ("but i think we should choose B on the last one b/c they might choose it since we won't have anymore rounds"). But their opponent defects in round 9. At which point they recognize that it might make sense to defect earlier

> 9: that's unexpected
> 14: we should start doing that lol
> 9: absolutely!
> 14: b/c everyone's going to choose B on the last one so why not do it on the second to last when it's unexpected?

Which they put into effect in the next super-game they get up on a cooperative path only to have their opponent defect in round 8. There is no further discussion of when to defect until round 7 of the next super-game at which point they simply agree to defect, finally getting ahead of their opponent.[38]

[38] Unfortunately this is their last super-game so we cannot determine what they do after this.

References

Andreoni, James and John Miller 1993. “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence.” *The Economic Journal,* 103 (418), pp. 570-585

Bereby-Meyer, Yoella and Alvin Roth. 2006. “The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation.” *American Economic Review*, 96 (4), pp. 1029-1042.

Bicchieri, Cristina. 1989. “Self-Refuting Theories of Strategic Interaction: A Paradox of Common Knowledge” in Wolfgang Balzer and Bert Hamminga, eds. *Philosophy of Economics*. London: Kluwer, pp. 69-85.

Camerer, Colin and Keith Weigelt. 1988. “Experimental Tests of a Sequential Equilibrium Reputation Model.” *Econometrica*, 56 (1), pp. 1-36

Charness, Gary and Mathiass Sutter. 2012. “Groups Make Better Self-Interested Decisions.” *Journal of Economic Perspectives*, 26 (3), pp. 157-176.

Cooper, Russell, Douglas DeJong, Robert Forsythe, and Thomas Ross. 1996. “Cooperation without Reputation: Experimental Evidence from Prisoner’s Dilemma Games.” *Games and Economic Behavior,* 12 (2), pp. 187-218.

Cox, Caleb, Matthew Jones, Kevin Pflum, Paul Healy. 2015. “Revealed Reputations in the Finitely-Repeated Prisoners’ Dilemma.” *Economic Theory*, 58 (3), pp. 441-484.

Dál Bo, Pedro, and Guillaume Fréchette, Guillaume. 2011. “The Evolution of Cooperation in Repeated Games: Experimental evidence.” *American Economic Review*, 101 (1), pp. 411-429.

Davis, James**.** 1992. “Some Compelling Intuitions About Group Consensus Decisions,

Theoretical and Empirical Research, and Interpersonal Aggregation Phenomena: Selected Examples, 1950-1990.” *Organizational Behavior and Human Decision Processes*, 52(1), pp. 3-38.

Fudenberg, Drew, and Eric Maskin. 1986. “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information.” *Econometrica*, 53 (2) pp. 533-554.

Jeheil, Philippe 2005. “Anlogy-based Expectation Equilibrium.” *Journal of Economic Theory.* 123, pp 81-104.

Jung, Yun Joo, John Kagel, and Dan Levin. 1994. “On the Existence of Predatory Pricing: An Experimental Study of Reputation and Entry Deterrence in the Chain-Store Game.” *RAND Journal of Economics.* 25 (1) pp. 72-93

Kagel, John 2015. “Finitely Repeated Prisoner Dilemma Games with Cheap Talk: Teams versus Individuals.” Ohio State University working paper.

Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson. 1982. “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma.” *Journal of Economic Theory*, 27 (2), pp. 245-252.

Lorge, Irving and Herbert Solomon. 1955. “Two Models of Group Behavior in the Solution of Eureka-Type Problems.” *Psychometrika*, 20 (2), pp. 139-148.

Neyman, Abraham. 1985. “Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner’s Dilemma Game.” *Economics Letters,* 19, pp 227-229.

Reny, Phillip. 1992 “Rationality in Extensive-Form Games.” *Journal of Economic Perspectives* 6 (4), pp 103-118.

Reuben, Ernesto and Sigrid Suetens. 2012. “Revisiting strategic versus non-strategic cooperation.” *Experimental Economics* 15 (1), pp. 24-43.

Samuelson, William and Richard Zeckhauser. 1988. “Status Quo Bias in Decision Making.” *Journal of Risk and Uncertainty* 1, pp. 7-59.

Selten, Reinhard and Rold Stoecker. 1986. “End Behavior in Sequences of Finite repeated prisoner’s dilemma supergames: a learning theory approach.” *Journal of Economics Behavior and Organization* 7 (1), pp. 47-70.

Wildschut, Tim, Chester Insko. 2007. "Explanations of Interindividual-Intergroup Discontinuity: A Review of the Evidence." *European Review of Social Psychology* 18 (1), pp. 175-211

Wildschut, Tim, Brad Pinter, Jack Vevea, Chester Insko, and John Schopler. 2003. "Beyond the Group Mind: A Quantitative Review of the Interindividual-Intergroup Discontinuity Effect," *Psychological Bulletin* 129 (5), pp. 698-722.

Windschitl, Paul, Justin Kruger, and Ericka Simms. 2003. "The influence of egocentrism and focalism on people's optimism in competitions: when what affects us equally affects me more." *Journal of Personality and Social Psychology*, 85(3), 389-408.

Tables and Figures

Figure 1

Stage Game Payoffs

(in ECUs)

	A	B
A	105 105	5 175
B	175 5	75 75

Figure 2

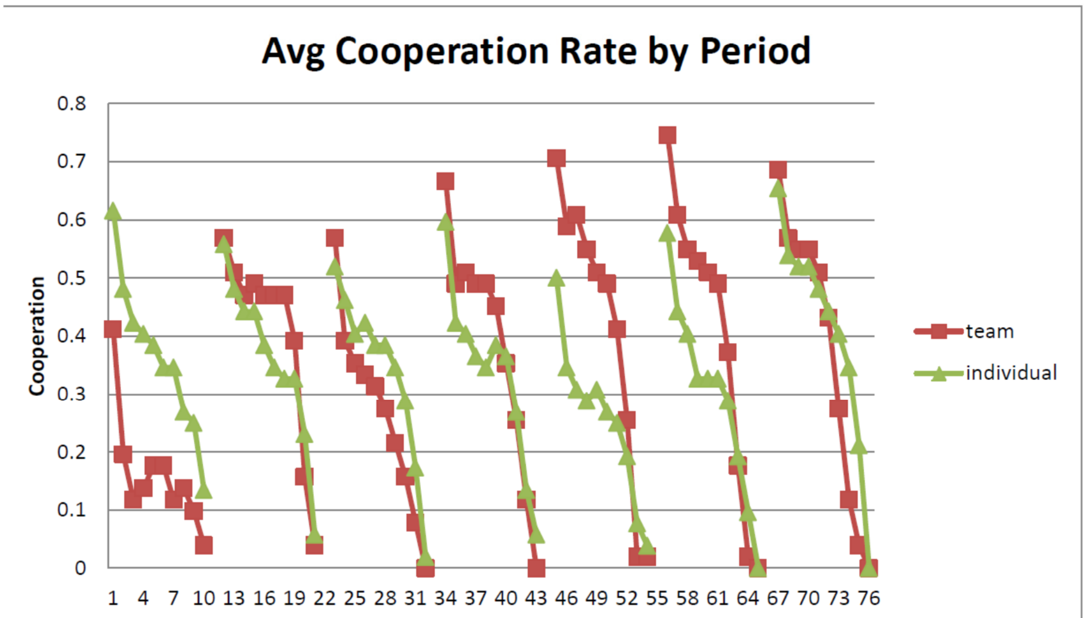

Table 1

Average Stage One Cooperation Rates

Super Game	Individuals	Teams	Difference: Individuals less Teams (z-statistic)
1	0.615	0.412	0.203 (2.07)**
2	0.558	0.569	-0.011 (-0.11)
3	0.519	0.569	-0.050 (-0.51)
4	0.596	0.667	-0.071 (-0.75)
5	0.500	0.706	-0.206 (-2.15)**
6	0.577	0.745	-0.168 (-1.81)**
7	0.654	0.686	-0.032 (-0.347)

** Significantly different from zero, at the 5% level, two-tailed test of differences in proportions

Table 2
Median Round Defected in Conditional on
Being on a Cooperative Path
(mean round in parentheses)

Super-Game Number	Individuals		Teams	
	Number of Defections[1]	Round	Number of Defections[1]	Round
1	8	10 (9.0)	5	9 (9.4)
2	13	10 (9.3)	14	9 (8.9)
3	13	9 (9.0)	10	8 (7.9)
4	9	9 (8.7)	15	9 (8.2)
5	10	9 (8.7)	18	8 (8.1)
6	11	9 (8.8)	16	8 (7.8)
7	15	9 (8.9)	14	7 (7.0)

[1] In cases where both agents defected in same round, both are counted. In cases where one agent defected first, it is counted as a single defection.

Table 3

Change in Round Defecting in Across Super-Games on a Cooperative Path

(number of observations)

Earlier Super-game	Were Cooperating[1]		Both Defected		Unilateral Defection		Pooled	
Following Super-game	Individual[2] (30)	Team (38)	Individual (25)	Team (28)	Individual (31)	Team (41)	Individual (86)	Team (107)
Earlier	20.0%	13.2%	36.0%	42.9%	9.7%	7.3%	20.9%	18.7%
Same	36.7%	50.0%	40.0%	42.9%	45.2%	73.2%	40.7%	57.0%
Later	33.3%	36.8%	4.0%	10.7%	32.3%	17.1%	24.4%	22.4%
Censored[3]	10.0%	0.0%	20.0%	3.6%	12.9%	2.4%	14.0%	1.9%

[1]Excludes cases with mutual defection in Round 10

[2]Includes four individuals who cooperated in the last stage game

[3]For teams, chats were used to eliminate all but 2 censored observations.

Table 4

Cooperative Path Patterns

Pattern	1st Time on Cooperative Path		After 1st Time	
	Teams	Individuals	Team	Individuals
P_{CC}	27	35	104	77
Other	14	4	7	10

Table 5

Coding Categories: Team Dialogues

Cooperate: Coding conditional on team cooperating (choice of A)

C1. If we cooperate other team might/will cooperate – includes cooperation will result in making more money or necessary to get the other team to cooperate.

C2. What to do if the other team fails to reciprocate cooperation in early plays of the game. Must include reference to defecting at some point in response to the other team's failure to reciprocate.

C3. It's in our best interest to cooperate without discussion of the logic behind cooperating. Essentially C1 above but without discussion of the underlying logic.

C4. Discussion of when to defect in later rounds (including coding the round in which planning to defect).

C5. Partner disagreeing with cooperation – advocating defection.

Defection: Coding conditional on teams defecting (choice of B).

D1. It's the safest choice

D2. Discussion of defection in terms of being a strategic response to the other team's defecting.

D3. Defecting but planning to cooperate if other does so. Often recognize must pay penance as the other team is likely to punish them for having defected. This is only coded for rounds 1-3.

D4. It's in our best interest – defection without any logic behind the doing so.

D5. Recognizing they can't cooperate until the start of a new match, along with the benefits of mutual cooperation. Includes discussion of what to do if the other team fails to reciprocate cooperation in early plays of the game. Analogue to C1 and C2 above.

D6. Partner disagreeing with defection - advocating cooperation.

Additional coding categories irrespective of choices:

X1. I know this game and the way it's supposed to be played; includes having played the same game in a previous experiment or learned about it in a class.

X2. When not cooperating discussing defection in later rounds of a match if and when table to achieve mutual cooperation. Coded just for the first time this occurred.[39]

X3. Partial unraveling - discussion of defecting earlier than in a previous super-game.

X4 Laying out the complete unraveling argument.

X5 Coding for first identifying beliefs that initial cooperators are clearly going to defect, or very likely to defect, at some point near the end of the super-game.

[39] This is the analogue to C4 when the team was not cooperating and was only coded for the first occurrence prior to having assigned C4 to a team. It was done after the initial coding of the data in order to fill an obvious gap in the analysis. It was done by one of the coders.

Table A1

Probit Regression: Cooperation Rates Across Super-games

Constant	-0.770 (0.257***
Initial cooperation dummy (= 1if coop in R1 of super game 1)	0.772 (0.213)***
Team dummy	0.399 (0.214)*
Super game (time trend)	0.059 (0.029)**
Past experience dummy (= 1 if opponent cooperated in R1 of super game t-1)	0.374 (0.094)***
Pseudo Log-Likelihood	-376.9
Pseudo R^2	0.097

*significant at 10% level; ** at 5* level: *** at 1% level

Note: Standard errors corrected for clustering at the subject level