Infinitely Repeated Prisoner Dilemma Games: Comparing Teams with Individuals*

David J. Cooper

Florida State University

John H. Kagel

Ohio Stater University

6/4/2019

(Draft)

Abstract

Teams are less cooperative than individuals in infinitely repeated PD games to begin with, but are more cooperative with experience. Strategies from team discussions are similar to maximum likelihood estimates at an aggregate level. However, discussions identify strategies that are substantially more complicated than standard strategies. Around 65% of teams using Always Defect to begin with, primarily out of safety concerns, with the remainder not considering the potential benefits of cooperating. Stage 1 cooperation rates are the same, or higher, in a comparable set of finitely repeated games, for both super-game 1 and the last common super-game.

Key words: Infinitely repeated prisoner dilemma games, teams and individuals, strategies identified from team chats.

JLE numbers:

*We have benefitted from discussions with Tim Cason and Guillaume Fréchette. Jinsoo Bae, Puja Bhattacharya, and Kirby Nielsen have provided valuable research assistance. Research support was partially provided by NSF grants SES 1226460 and SES 1630288. Opinions, findings, conclusions or recommendations offered here are those of the authors and do not necessarily reflect the views of the National Science Foundation.

The results of an infinitely repeated prisoner's dilemma (IRPD) game experiment are reported, comparing individuals with two person teams. The motivation for the experiment is two-fold: First, many, if not most, strategic interactions in economics involve groups of individuals, so it is important to see if their behavior differs significantly from individuals. Looking at two-person teams is a step in that direction. Second, within team discussions provide direct insight into the strategies and motivations for choices made. These, are compared with maximum likelihood estimates for strategies in these games (Dal Bó and Fréchette, 2011, Fudenberg, Rand and Dreber, 2012). Results are also compared to an earlier series of finitely repeated prisoner dilemma (FRPD) games (Kagel and McGee, 2016), with the same stage game payoffs, and the same expected number of stage games in each super-game.

In the IRPD games, teams start out with significantly lower cooperation rates than the teams. Within team chats show this is rooted in safety considerations (fear of the getting the sucker payoff), along with a minority of "myopic" teams, focused on the higher stage game payoffs for defection compared to cooperating. These lower cooperation rates, along with the focus on safety, coincide with results comparing teams with individuals in PD games reported in the psychology literature. With experience, cooperation rates for teams are higher than for individuals in later super-games. Strategies based on coding of team discussions coincide with this increased cooperation, and are broadly consistent the Strategy Frequency Estimation Method (SFEM) employed in Dal Bo and Frechette (2011), although there are some important differences. Among the latter, the SFEM identifies a small, but positive percentage of agents choosing Always Cooperate, albeit with a substantial standard error. But the team chats fail to show any teams willing to always cooperate. This difference can be attributed to teams playing Grim, who happen to meet each other. Also identified are a number of more complicated strategies, including one labeled generalized suspicious tit-for-tat: teams that start off defecting, intending to cooperate if their opponent does, but with a variety of patterns that may or may not result in cooperation. Always Defect captures a significant percentage of team and individual strategies. Coding of team chats show a significant number employing more cooperative strategies over time, albeit with occasional backsliding, but no team is identified as moving from a more cooperative strategy to Always Defect.

Comparing FRPD games to IRPD games with the same stage game payoffs, and the same expected number of stage games, stage 1 cooperation rates are the same or higher in the FRPD games, in the first super-game and in the last common super-game. This contrasts with to Dal Bó (2005), who finds that with some experience stage 1 cooperation rates are lower in the FRPD games. This difference can be explained in terms of the much higher continuation value used here ($\delta = .9$) compared to $\delta = .1$ and $\delta = .3$ in Dal Bó.[1] Cooperation rates drop precipitously across FR stage games, while remaining roughly the same across comparable stage games in the IR games.

The outline of the paper is as follows. Section I reviews past research with PD games. Experimental design and procedures are reported in Section II. Section III compares differences between teams and individuals for IRPD games, focusing first on outcomes over time, followed by analysis of strategies employed. Section IV briefly compares outcomes in the IRPD games with the FRPD games.

I Prior Research

There have been numerous studies of both FRPD and IRPD games over the years. Dal Bó and Fréchette (2018) provide an extensive survey of the experimental literature on IRPD games. Embry, Fréchette and Yuksel (2017) provide a brief survey of FRPD games. Experiments have been conducted with and without "noise". With noise subject choices are replaced by a randomly chosen outcome, with a known probability (Fudenberg, Rand, and Dreber, 2012; Ayogi, Bhaskar and Fréchette, 2019; Cason and Moi, 2018). Most of the economics research on PD games involve individuals. The exceptions are Kagel and McGee (2016) who compare two-person teams and individuals in FRPD games, and Cason and Moi (2018) who compare three person teams and individuals in IRPD games. In Kagel and McGee teams agree jointly on the choice between Cooperate (C) and Defect (D); Cason and Moi use majority rule.

Most of these prior PD experiments employ the Direct Response Method in which each agent chooses C or D in each round of the repeated game. They then use the Strategy Frequency Estimation Method (SFEM) introduced in Dal Bó and Fréchette (2011) to estimate the frequency of a given set of strategies. In our case team chats will be used to identify strategies, some of

[1] See Lugovskyy, Puzzello, and Walker (2018) who report game structures with both significantly higher, as well as lower, stage one cooperation rates for FRPD compared to IRPD games, with continuation values of .4 or less.

which have complicated dynamic elements that change within a given super-game. Outcomes will also be compared to STEM estimates of strategies employed.[2]

There is an important line of research in social psychology concerned with differences between individuals and teams in PD games. These are quite different in structure than those commonly employed in economics: They typically involve a single super-game, with prior knowledge that agents will be paired with each other for somewhere between t and t+n stage games, with the actual stopping point somewhere in that interval.[3] Further, this exercise is typically repeated only one time within a given experimental session. The key finding from these experiments is that teams are *less* cooperative than individuals (referred to in the psychology literature as the "discontinuity effect"). This is generally attributed to intergroup relations being characterized by greater fear and greed than inter-individual relations, as opposed to teams being more insightful regarding the finite nature of the interactions. (See Wildschut et al., 2003 and Wildschut and Insko, 2007 for surveys). Consistent with this, Kagel and McGee report that teams are significantly less cooperative than individuals in stage one of the first super-game, with team discussions supporting fear of getting the sucker payoff as the primary factor behind the lower cooperation rates. However, in later super-games teams are as, or more, cooperative than individuals in stage 1, recognizing the benefits of early stage-game cooperation.[4]

II. Experimental Design and Procedures:

Subjects played a simultaneous move, indefinitely repeated PD (IRPD) game with a 0.9 continuation value, with the 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 stage-games and paid in cash at the end of an

---

[2] Cason and Moi (2018) and Dal Bó and Fréchette (2018b) conduct experiments where agents choose between extensive set of fixed strategies, which are then played out round by round against an opponent who has done the same. Agents are permitted to change strategies between super-games. Closely related to this, Romero and Rosokha (2018) have subjects construct their own strategies, which are then played out round by round against an opponent who has done the same.

[3] These experiments typically employ financial incentives. They have been run with and without between agent discussions, after which agents made their decisions. The data shows that teams are generally less cooperative than individuals, with the largest differences occurring with unrestricted, face-to-face, communication (Wildschut et al., 2003) compared to restricted, or no, between agent discussions.

[4] As is typical in PD games in economics, Kagel and McGee do not permit between agent discussions.

experimental session. Each member of a team received the team's payoff. Earnings averaged $44.98 per subject.

Figure 1: Stage Game Payoffs

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

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 refer to agents playing a game, with agents being either two person teams or individuals. Following each IRPD game, agents were randomly re-matched under the restriction that no two agents would be re-matched in consecutive super-games. There were six individual and six team sessions.

Following the last stage game within a super-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 that "there is a 90% chance of another round for that match and a 10% chance you will move on to another match with another team/individual".

There were 6 individual subject sessions with between 14 and 18 subjects in each session. Subjects were told there would be between 12-15 matches (i.e., super-games) within an experimental session. Four sessions had 13 super-games, the other two had 12.

Instructions for the team sessions were essentially the same as for individual sessions, except that teams were told to "make decisions jointly", along with a chat box to send messages back and forth. They were told their messages would be recorded, that only they and their partner would see these messages, and to use the "chat box to discuss your choices, and come to an agreement regarding what choice to make". Subjects were instructed to be civil to each other and to not use profanity, nor to identify themselves, instructions that were followed throughout.

There were 6 team sessions with between 8 and 12 teams in each session. For each super-game, the first two stage games allotted 2 minutes for teammates to discuss their choices and come to an agreement. This was reduced to 40 seconds for subsequent stage games.[5] The first two team sessions had 6 and 7 super-games respectively, as they were scheduled to last 2 hours. Subsequently, session time was increased to 2.5 hours, along with modest reductions in the time teammates could discuss their choices.[6] The remaining four team sessions had between 9 or 12 super-games.

A total of five different random seeds were used for the team and individual sessions, with one or more of the individual (team) sessions conducted under each of the five seeds. This leads to variability across sessions in the number of stage games in each sequence of super-games with different seed values. Dummy variables will be employed to control for these differences in the statistical analysis. In addition, results are analyzed under the restriction that for each seed, the number of super-games considered are the same for the teams and individuals with the same seed. That is, even though teams generally have fewer super-games than individuals, the analysis is limited to the same set of super-games for a given seed.

As noted, data will be employed from a corresponding series of FRPD games using the same payoffs, where the number of stage games (10) equaled the expected number in the infinitely repeated super-games (Kagel and McGee, 2016; KM). These will be used to compare cooperative outcomes with the IRPD games, as well as to provide a sufficient number of sessions to cluster standard errors at the session level in the statistical analysis. These sessions used essentially the same procedures as those employed here, with the same subject population (Ohio State University students).[7]

III Theoretical Considerations

When agents are sufficiently patient, repeated PD games have a great many equilibrium outcomes, as shown by various folk theorems.[8] These folk theorems show that cooperative

---

[5] The default options, in case agreement was not reached within the time allotted, can be found in the instructions (web site goes here). In the overwhelming number of cases, teammates reached agreement on what action to take. Individuals had up to 1 minute to make their choices, but this was never a binding constraint.

[6] In the first two sessions, three minutes were allotted for reaching agreement in the first two stage games, with 1 minute after that. Subjects said that this was more than sufficient time to reach agreement, leading to the reduction in time in subsequent sessions. Following the end of each stage game teams (individuals) had up to 10 (15) seconds to view the results before moving on to the next stage game, and ended earlier once all agents hit the next button. The team chat box was open after a team made their choice.

[7] KM had 5 teams sessions with 7 or more super-games in each session, and 5 individual sessions with 10 super-games in each session.

[8] Friedman (1971), Auman and Shaply (1994), and Fudenberg and Maskin (1986).

outcomes are possible, but non-cooperative play is an equilibrium as well. Specific criteria have been discussed as to whether cooperation is more likely or not. Blonski and Spagnolo (2004) proposed that cooperation is more likely when payoffs satisfy "risk dominance" defined as whether tit-for-tat (TFT) is a best response to ½ – ½ probability distribution over TFT and always defect (AD) in a game with only these two strategies. The Dal Bó and Fréchette (2018) survey shows that when risk dominance is satisfied cooperation generally increases as subjects gain experience. They also show that an even more predictable criteria that cooperation will increase over time is when the size of the basin of attraction for always defect (BAD) is .5 or less, where BAD is defined as the set of beliefs that make AD optimal when play is limited to opponents using Grim and AD (Dal Bo and Frechette, 2011, 2018). Stage game payoffs, along with the continuation value of .9 employed here, satisfy risk dominance with BAD = .42.[9]

The experimental literature shows that teams generally, but not always, tend to be more "rational" than individuals as defined in economics (Charness and Sutter, 2012, Kugler, Kausel, and Kocher, 2012). This is of no help in predicting whether teams will be more cooperative than individuals given the folk theorem. However, the psychology literature on PD games, as flawed as it is from an economists' perspective, would predict lower cooperation rates for teams, at least to begin with, consistent with the discontinuity effect, and to be driven largely by fear and/or greed. As for IRPD and FRPD games with the same payoffs, economic theory would predict that given time for unraveling to occur, both teams and individuals will have lower cooperation rates in stage one in finitely repeated games, and that in general cooperation will decrease across stage games within a given finitely repeated super-game, with no systematic decrease across stage games for infinitely repeated games.

Coding of team chats will be used to identify team strategies employed. These will be aggregated to closely match strategies typically employed using the strategy estimation technique when comparing between the two. Using the team chats we also look for previously unrecognized strategies, as well as how fixed these strategies are within a given super-game as teams have no commitment device to stick with the chosen strategy, and often change choices on the fly.

---

[9] Lugovskyy, Puzzello, and Walker (2018), explore the effects of variations in stage game payoffs on cooperation, along with BAD values less than or greater than .5 on cooperation.

IV Experimental Results.

4.1: Cooperation Over Time: We use the notation SGx to refer to the xth supergame in an experiment (i.e. SG1 for the first supergame, SG2 for the second supergame, etc.) and Stx to refer to the xth stage game within a supergame (i.e. St1 for the first stage game, St2 for the second stage game, etc.).

The discussion in this subsection focuses on mutual cooperation, stage games where both players choose cooperation. Partially this is a matter of convenience: individual cooperation and mutual cooperation are highly correlated, and it is cumbersome to describe both. This choice also reflects a reality of the dataset. Normally, pairs of players are either mutually cooperating or mutually defecting – for 89% of all stage games, and 92% of stage games subsequent to St1, the two players make the same choice, mutually cooperating or mutually defecting. The success of a pair at achieving cooperation is largely determined by whether they mutually cooperate in the first stage game and when they change their status between mutual cooperation and mutual defection in later stage games.

The dataset is not well-suited to the use of simple non-parametric statistics given the changing length of supergames and the different number of supergames played in different sessions. Nonetheless, to give a first crude impression of the data, some simple non-parametric tests are reported on first. For the IRPD data, these are generally Wilcoxon matched-pairs signed-ranks tests, with observations paired by seed class between treatments. An observation is the average for an entire session, making these very weak tests so that the tests are biased in favor of Type II, rather than Type I, errors. The FRPD data does not have any obvious matching of sessions between treatments, so rank-sum tests are used.

Figure 2 reports mutual cooperation rates over the first ten stage games of a supergame. The upper and lower panels show data from the IRPD and FRPD respectively. Within each panel, data is broken down by whether the agents were individuals or teams and either early (SGs 1 – 3) or late (SGs 4 – 7) supergames. Figure 2 cuts off the late supergames after SG7 as this was the first point where a substantial number of sessions terminated, but our statistical analysis otherwise uses all data prior to the final common supergame, defined as the last supergame played by both individuals and teams within a seed class.[10]

[10] One team session terminated after SG6, but another four terminated after SG7.

Before examining treatment effects, it is worth noting an important general feature of the IRPD data. Subsequent to the first stage game, St1, the frequency of mutual cooperation typically is quite flat. There is a bit of random noise, some of which reflects supergames of shorter length finishing and dropping out of the dataset, but there is no obvious trend. Using a similar econometric specification like the one reported in Table 3 below, a small, statistically insignificant, negative trend in mutual cooperation rates can be detected across stage games.[11] That is, subsequent to St1, mutual cooperation is sticky. This implies that treatment effects are largely driven by what happens in St1. Analyzing cooperation based on choices in St1 also has the advantage of not being affected by the differing length of supergames in the IRPD. We therefore begin our discussing by examining mutual cooperation in St1.

In the IRPD games, mutual cooperation in St1is *higher* for individuals than teams in SG1 (19.2% vs. 10.3%), but teams overtake individuals with time. By the final common supergame, mutual cooperation in St1 is lower for individuals than teams (36.5% vs. 55.2%). Comparing SG1 and the last common supergame, mutual cooperation in St1 increases significantly for teams ($z = 2.20$; $p = 0.028$) but not individuals ($z = 1.05$; $p = 0.292$), and the increase is significantly larger for teams than individuals ($z = 1.78$; $p = 0.075$).

Comparing the IRPD and FRPD data, shows a similar pattern with respect St1 cooperation rates: Mutual cooperation in St1 is higher for individuals than teams in SG1 (46.2% vs. 12.0%), but by the final common supergame, SG7, this difference has flipped with individuals having less mutual cooperation in St1 than teams (38.5% vs. 52%). This change in mutual cooperation is not significant for individuals ($z = 0.00$; $p = 1.000$), but is for teams ($z = 1.76$; $p = 0.078$), with the difference in differences between teams and individuals statistically significant ($z = 1.79$; $p = 0.074$).

Observation 1: Mutual cooperation in St1 rises faster for teams than individuals for both the IRPD and FRPD games.

---

[11] Based on estimates from all stage games except St1 and all supergames up to and including the final common supergame, in a probit model where the dependent variable is a dummy for mutual cooperation and the independent variable of primary interest is the current stage game. Controlling for mutual cooperation in St1 of the current supergame, the number of stage games in the previous supergame, the choices of opponents in St1 of the previous supergame, the current supergame, and the seed class, the estimated marginal effect for stage game is -.0013 with a standard error of .0008, which is not statistically significant at standard levels.

Figure 2: Mutual Cooperation Rates

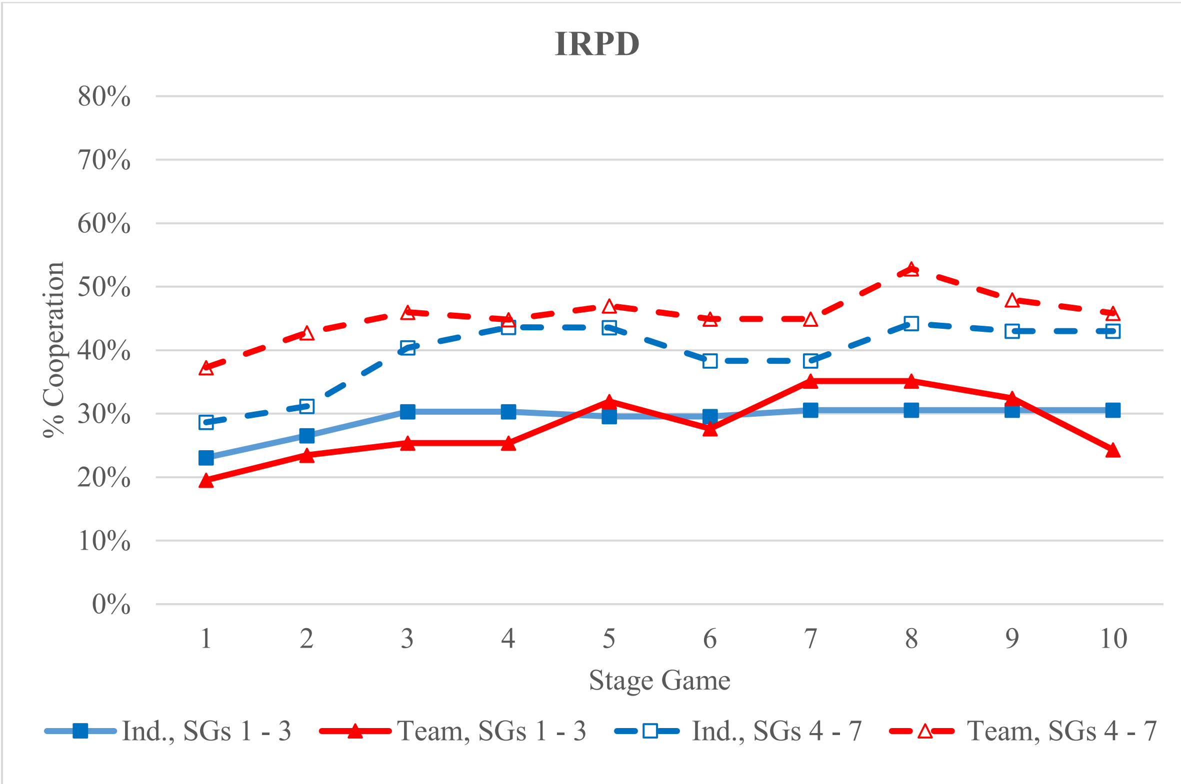

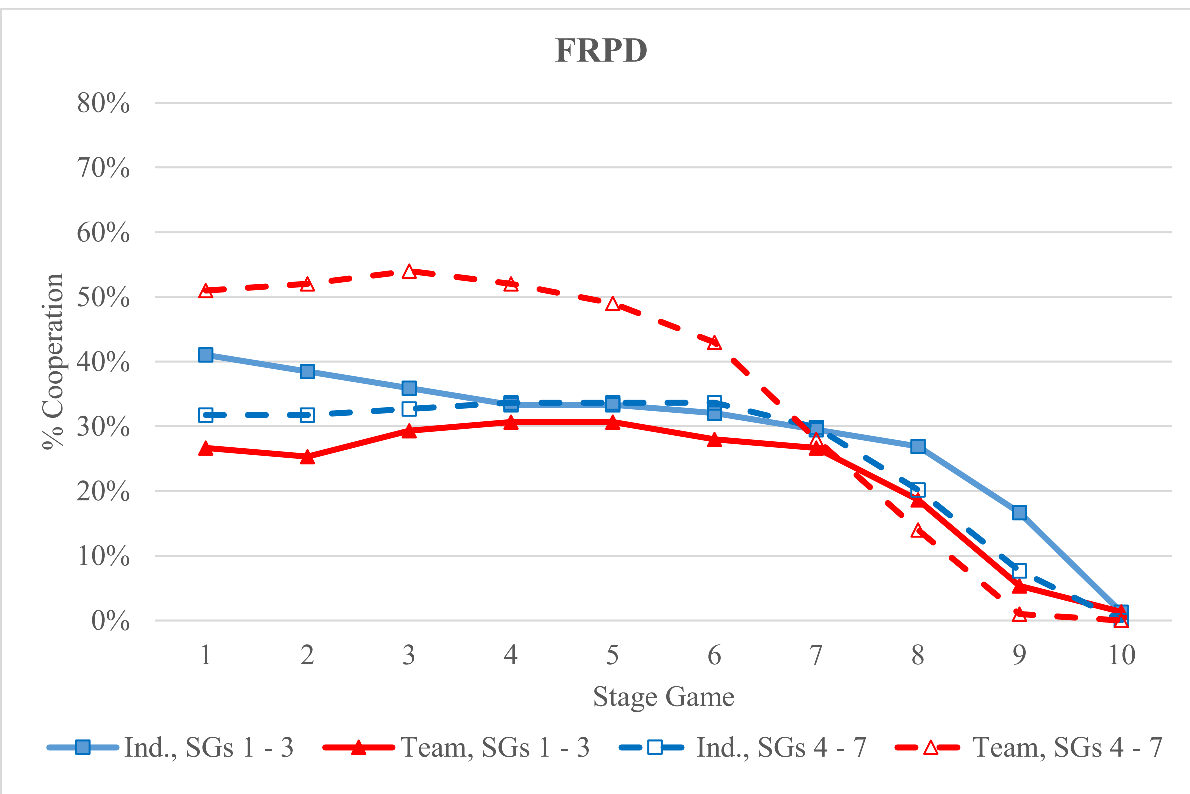

Figure 3 examines what happen subsequent to St1. Rather than focusing on levels of mutual cooperation, this figure examines the stability of mutual cooperation across stage games within supergames, distinguishing between the frequency of switching between mutually cooperating (Lagged Mutual) and parings with both defecting or one cooperating with the other defecting (Not Lagged Mutual). The data is broken down by the type of game (IRPD vs. FRPD), the type of agent (Individual vs. Team). Because most of the action in IRPD games takes place early in supergames, the first five stage games are broken out separately from all infinitely repeated stage games.

Figure 3: Transitions: Mutual Cooperation

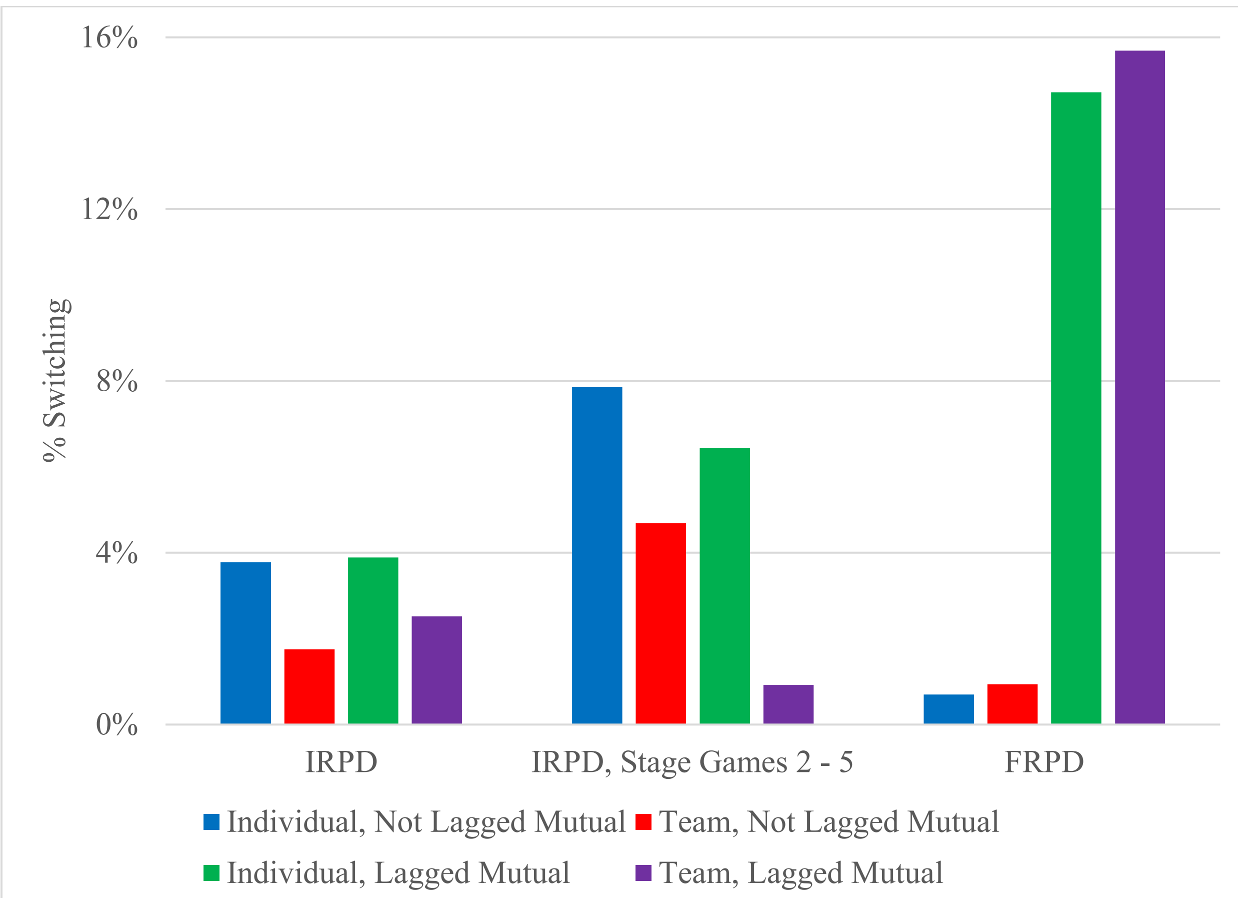

For IRPD games, if a pair of agents are mutually cooperating after St1, they normally continue to cooperate in the next stage game (96.2%), and if not cooperating, this also tends to persist into the next stage game (96.8%). The probability of switching, either to or from mutual cooperation, is consistently less with teams than individuals. The difference between teams and individuals is even stronger in early stage games (i.e. St2 – St5) when switches are most frequent. The difference between teams and individuals may appear small, but has relatively large cumulative effect: For individuals, by the fifth stage game 19.5% of pairs have switched either to or from mutual cooperation in St1, roughly double the team rate (10.3% switching). However, the net effect of switches to and from mutual cooperation is close to zero for both individuals and teams so that these differences in stability have little impact on the differences in **levels** of mutual cooperation reported earlier.

This stability in mutual cooperation rates does not extend to FRPD games, as one would expect given the finite time horizon. Further, there are minimal differences on this score between teams and individuals.

Observation 2: Mutual cooperation in the IRPD is more stable with teams than individuals.

The non-parametric tests reported above are useful but extremely conservative and somewhat limited. We turn to less conservative regression analysis that corrects for the varying length and number of supergames as well as the differing experiences of subjects.

The regressions in Table 2 revisit the differences in stage 1 cooperation rates over time reported in Observation 1. The dependent variable is whether agents achieve mutual cooperation in the first stage game of a supergame, a binary variable, so a probit model is used and marginal effects are reported. There is one observation for each pair of agents playing a supergame, with corrections standard errors clustered at the session level. As standard controls, both regressions include dummies for the supergame and seed class,[12] the length of the previous supergame,[13] and whether the two agents experienced defection in the first stage game of the previous supergame.[14] For an apples to apples comparison between individual and team data, the analysis is limited to the last common supergame in each seed class. Standard controls for the lagged number of stage games in the previous super games and for whether there was mutual cooperation in St1 in the previous supergame are employed as well.

Beyond these standard controls, Model 1 only adds dummies for the team treatment in the IRPD and FRPD data. Team play has a weak positive effect on mutual cooperation in St1 for the IRPD. The surprise is that any effect at all is reported, since we know that the difference in

---

[12] There is a dummy for each seed class in the IRPD, with the FRPD the omitted category. The seed class dummies are co-linear with a dummy for "Individual, IRPD," hence no such variable is included.

[13] For SG1 this is set equal to 10, the expected supergame length.

[14] This is averaged across the two agents in a pair, with the mean value observed in SG 1 used so as not to drop SG 1 from the analysis.

cooperation rates between teams and individuals changes sign over time. Model 2 gets directly at these dynamic effects by adding interactions between the treatments (Individual, FRPD is the omitted category) and the supergame.[15] The interaction between the team treatment and the supergame is positive and significant for both the IRPD and the FRPD data confirming Observation 1 after controlling for a number of potential confounds: Mutual cooperation in St1 rises faster for teams than individuals for both the IRPD and FRPD games.[16]

Table 2: Regression Analysis, Mutual Cooperation in Stage Game 1

	(1)	(2)
Team, IRPD	0.105* (0.060)	-0.107 (0.077)
Team, FRPD	0.031 (0.092)	-0.221** (0.089)
Lagged # Stage Games	0.003* (0.002)	0.003 (0.002)
Experienced Defection St1, Previous Supergame	-0.234*** (0.040)	-0.208*** (0.034)
Supergame * IRPD		0.036** (0.016)
Supergame * Team, IRPD		0.039** (0.017)
Supergame * Team, FRPD		0.073*** (0.024)
Team Infinite - Finite	0.074 (0.155)	0.114 (0.103)
Supergame * Team Infinite - Finite		-0.034 (0.030)
Log-Likelihood	-599.43	-590.15
AIC	1238.87	1222.29
BIC	1338.56	1326.97
Observations	1,080	1.080

Notes: Three (***), two (**), and one (*) stars indicate significance at the 1%, 5%, and 10% level using a two-tailed test.

[15] Again, since the seed class dummies are co-linear with a dummy for "Individual, IRPD," it is not included.
[16] We have run similar regressions using either the time of the first mutual defection or cooperation across all stage games are the dependent variable. In both cases, mutual cooperation rises faster across supergames for teams than individuals. Given the stability of aggregate behavior across stage games, it is not surprising that results for these alternative dependent variables look similar to the results for mutual cooperation in St1.

The regressions reported in Table 3 are designed to provide support for Observation 2. The dependent variable is a dummy for whether a switch in mutual cooperation has taken place, either to or from mutual cooperation. This is a binary variable, so that a probit model is used, with marginal effects are reported. There is one observation per each pair of agents playing a *stage game* (rather than a *supergame*). To allow for the use of lagged variables, data from St1 is dropped. The standard errors are once again corrected for clustering at the session level. Both regressions include dummies for the supergame and seed class, the length of the previous supergame, whether the two agents experienced defection in the first stage game of the previous supergame, and the current stage game interacted with the treatment (IRPD or FRPD). Again only supergames up to the last common supergame between teams and individuals are used.

Table 3: Probit Models: Switches in Mutual Cooperation

	(1)	(2)
Team, IRPD	-0.017*** (0.006)	-0.016** (0.008)
Team, FRPD	0.006 (0.007)	0.007 (0.012)
Lagged # Stage Games	-0.000 (0.000)	-0.000 (0.000)
Mutual Cooperation St1, Previous Supergame	-0.003 (0.006)	0.001 (0.005)
Stage Game, IRPD	-0.001*** (0.000)	-0.001*** (0.000)
Stage Game, FRPD	0.006*** (0.001)	0.009*** (0.002)
Lagged Mutual Cooperation, IRPD		0.000 (0.005)
Lagged Mutual Cooperation, Team, IRPD		0.011 (0.012)
Lagged Mutual Cooperation, FRPD		0.265*** (0.039)
Lagged Mutual Cooperation, Team, FRPD		-0.001 (0.015)
Team, IRPD – FRPD	-0.023** (0.008)	0.023* (0.010)
Lagged Mutual Cooperation, Team, IRPD – FRPD		0.012 (0.024)
Log-Likelihood	-1651.04	-1499.15
AIC	3344.08	3040.30

BIC	3496.87	3193.09
Observations	10,679	10,679

Notes: Three (***), two (**), and one (*) stars indicate significance at the 1%, 5%, and 10% level using a two-tailed test.

Beyond the standard controls, Model 1 only includes dummies for the team treatment in both the IRPD and FRPD data. For the IRPD, the team dummy is negative and significant at the 1% level, indicating that for the IRPD games teams are more stable from one stage game to the next. Although the magnitude of this effect is not large, as noted earlier the cumulative effect is increasing over stage games. This team effect does not extend to the FRPD games where the parameter is smaller in magnitude and not significant. As shown at the bottom of the table (Team, IRPD – FRPD), the difference in stability between the IRPD and FRPD is statistically significant. Model 2 adds controls for the lagged outcome for the IRPD and FRPD, as well as interactions between the lagged outcome and teams. The lagged outcome has a strong effect for the FRPD, consistent with the movement away from cooperation in later stage games, with no significant effect for IRPD games. In both cases, the lagged outcome has no significant interaction effect with teams. The relative stability for teams in the IRPD differs little for switches too or from mutual cooperation as the lagged mutual cooperation by team variable is not significant for the IRPD games. The probit analysis is supportive of Observation 2 that behavior across stage games is more stable for teams in the IRPD games.

Observation 3: The probit analysis supports Observations 1 and 2.

Discussion: The higher cooperation rates for individuals is consistent with the "discontinuity effect" reported in the psychology literature. As noted the psychology literature attributes this is to greater fear and greed on the part of teams compared to individuals. Looking at the chats for teams playing Always Defect in the first super-game, 64.5% (20/31) were classified as doing so out of safety – concern about getting the sucker payoff. Below are two typical examples of these chats (where A and B choices have been changed to C and D):

"D is always the safe choice" "so we are guaranteed 75 each time"

"I think we should stick with D to be safe"

An additional 25.8% (8/31) of the chats indicated myopia or greed as teams focused on the higher payoff for defecting, not considering a possible cooperative outcome:

"choose D since our payoffs will be 175 or 75, instead of 105 or 5"

"I say we always do optimum D, as it has the highest payoff on average" "That's what I was thinking too definitely"

The remaining 9.7% (3/31) could not be classified (2 cases).

What this psychology literature fails to identify is the higher cooperation rate for teams with experience. As noted, this is a result typically studying a single supergame, with no rematching of between opponents.

IV.2 Identifying Strategies in the IRPD Games: For teams there are two methods for determining strategies – analyzing the team discussions and the Strategy Estimation Method (Dal Bó and Fréchette, 2011). One important difference between the two is that in the strategy estimation method one must specify a set of strategies to be estimated, whereas chats can reveal strategies that one might not otherwise think of, or be too complicated to estimate. That is, teams can change their strategy within a given super-game, which they do sometimes, as there were no commitment devices available in the lab, which would not be typically available in the field.

In coding team discussions, the two authors went through a sample of the team chats, after which we specified a number of strategies, and other factors, to identify and code. Two graduate students then coded a sample from the data, meeting jointly with one of the experimenters to reconcile any obvious differences in interpretation, as well as to introducing categories we had not considered.[17] Teams do not always discuss the strategy they are using within a given super-game, except for when they first adopt and/or change the strategy employed. In these cases coding is based on the last relevant discussion, provided behavior matches the strategy in question. The pre-specified strategies used in the ML estimation were: (1) Always defect, (2) Grim, (3) Tit-for-Tat, (4) Suspicious TFT, and (5) Always Cooperate. After coding up the chats, these were aggregated into one of these strategies to compare with the MLEs. So for example, a number of teams using Grim with forgiveness for two or three initial defections which were aggregated with standard, single defection Grim to compare with the

[17] This is a two-step process – students coded one session of chats, then met with one of us, to go over and reconcile any obvious differences in interpretation, and then coded the remaining sessions on their own. Experience suggests that the PIs keep on top of this – checking the coding as it goes along to catch any obvious gaps.

MLEs. Substantially more complicated cases will be discussed below. Examples of chats related to each of these five strategies are reported below. Spelling and grammatical mistakes are not corrected for. Since choices in the experiment were labeled generically as A and B, these will be replaced with C and D for clarity.

Always Defect: See the examples reported above supporting this.

Grim: "My plan is choose C first for a few times, and if the other team keeps choosing D, we will switch to D." A team willing to show some patience before defecting. There were also a number of cases where teams planned to start with Grim, but after mutual cooperation was established, to unilaterally defect at some point. This will be referred to as Grim with counting, and as shown below, is based on the "gambler's fallacy" - that the super game was likely to end somewhere around the tenth round, so they should defect to take advantage of their opponent or to protect against them doing the same. Both of these alternatives to the standard Grim strategy – defecting after a single stage game where their opponent defected, are combined when comparing against Grim for the MLEs.

TFT: "So this is the dove's and hawk's game. … C first then play tit for tat?" And in a later super-game from the same team: "anyway, I heard the best why to play the hawks and doves game is tit for tat with some room for forgiveness" The chats show that specifying some variation of TFT are typically, but not always, based on what one of the subjects had learned in one of their classes. This is not surprising given the ubiquitous use of PD games in a large number of disciplines. We will have more to say below about strategies teams with obvious knowledge of PD games used.

Suspicious TFT: We adopt a generalized version of STFT as there is considerable variability in teams exhibiting the two key common characteristics of STFT: (i) choosing D to begin with out of fear of getting the sucker payoff, followed by (ii) efforts for mutual cooperation if their opponent plays C early on. An example of the simplest of these cases went as follows: "I will choose D this match …if they choose C we can catch on" – which led to mutual cooperation as their opponent was playing Grim with forgiveness.[18] An example of a substantially more complicated case is as follows: beginning with the first stage game for the super-game in question (own choice listed first): DD, DD, CD, CD, DC, CD, CC thereafter. What the coders characterized as "change" early on, typically identifies cases like this.

Always Cooperate: This was never even considered in the team chats. However, when Grim meets Grim, this will look like Always Cooperate. Likewise we found no discussions corresponding to win stay, lose shift, or choices consistent with same.[19]

---

[18] Note that if their opponent's strategy was Always Defect, STFT would not be identified other than through the team chats.
[19] WSLS is inconsistent with the long sequences of DD observed when teams failed to cooperate early on.

Table 4
Frequencies of Team Strategies: Chats versus Maximum Likelihood Estimates
(bootstrap standard errors of the mean in parentheses for MLE)

Early (SGs 1-3)	Always Defect	Grim	TFT	STFT	Always Cooperate	WSLS
Chats[a]	56	30	8	7	0	0
ML	42*** (3.7)	27*** (2.8)	9 (10.7)	18*** (4.7)	3 (10.0)	0
Late (SGs 5-7)						
Chats[a]	24	53	12	11	0	0
ML	23*** (3.3)	51*** (9.0)	11 (16.5)	3 (6.5)	12 (10.8)	0

[a] Percentage of chats not coded: early (7%) and late (2%). *, **, *** significant at the 10%, 5% and 1% levels respectively.

Table 4 reports the frequency with which these strategies were played for teams. First based on the ML estimates. Second based on the team chats. Data are reported for the first three super-games combined, followed by super-games 5-7.[20] The MLE and chat estimates are broadly consistent with each other and the results reported earlier – low cooperation rates for teams in early super-games, with increasing cooperation in later super-games.[21] As already noted Always Cooperate is positive in the ML estimates but zero in the chats. However the large standard error relative to the estimated value in the MLE estimates is such that one cannot reject a null hypothesis of zero, consistent with the zero frequency reported in the chats.

As noted, one advantage of the chats is that one can observe changes in strategies over time within a given super-game, along with relatively complex strategies.[22] Table 5 tracks team strategies over time across super-games 1-7. Strategies are aggregated into two categories – Always Defect and Cooperate – where the latter aggregates Grim, TFT, and generalized STFT. Around 20% of teams always played Defect over the seven super-games, never trying one of the cooperative strategies. Just over 36% consistently choose one or another of the cooperative strategies, doing so except for an occasional super game in which they reverted to Defect. (The

[20] Games 4, 5, and 6 for Team 1 that had only 6 super-games.
[21] For the SFEM estimates, the distribution of strategies changes significantly over time ($\chi^2 = 329.8$; d.f. = 9; $p < .001$.
[22] This is not to imply that all of the advantages are one sided in favor of the chats. In particular generalized STFT is subject to differences between coders, and there are no chats for individuals. Coding is also quite time and resource intensive.

number of teams doing so for one, or at most two, super-games is reported in parentheses below the raw data counts.) The third category shows the percentage of teams (44.6%) that started with Always Defect, only to switch and consistently play one of the cooperative strategies.[23] (No team switched from consistently cooperating to consistently defecting). What backsliding there is, typically resulted from frustration with efforts to cooperate. For example, one team who typically cooperated using Grim with three rounds of forgiveness (Grim3), switched to Always Defect after several super-games in which their opponents failed to "get it", returning to Grim3 after a single super-game.

Looking at the chats for those switching from Defect to Cooperate, the general pattern involved one or more earlier super-games where their opponent chose C in the first stage game and they played D, only to play DD after that, plus one or more longer super-games with DD. For example, in one case, in three of their first four super-games their opponent started with C, while they chose D, only to get stuck in DD for the remainder of each super-game, with games lasting between 4 and 22 rounds. At which point they switched to Grim in super-game 5.[24] This is consistent with the regressions in Table 1 that capture these two effects through - "length of last super-game" and their opponent cooperating in stage one of the previous super-game.

Table 5: Changes in Individual Team Strategies Over Time

Strategy	Frequency	Raw Data[a]
Always Defect	19.6%	11
Cooperate[b]	36.7%	20 (4)
Defect to Cooperate	44.6%	25 (7)

[a] Numbers in parentheses represent deviations from strategy in question: typically a single SG.
[b] Includes Grim, TFT, STFT and GSTFT. Two teams could not be reliably coded.

There are also a number of more unorthodox strategies uncovered in the team chats. As noted teams often played Grim but with two or three rounds of forgiveness (Grim2 and 3).[25] This has been reported in games with a random probability that intended actions would be

[23] Again, the number of teams reverting to defect for one, or at most two, super-games are reported in parentheses.
[24] "do you want to try choosing C first this time … bc it's a new team" "THAT'S SO CRAZY I was thinking the same thing" "and if they choose C we get 105and we could keep doing that … if not we go back to choosing D"
[25] 16 out of the 36 teams that used Grim for two or more super-games used Grim2 or Grim3 one or more times.

replaced with the opposite action (Fudenberg, Rand, and Dreber, 2012; Aoyagi, Bhaskar, and Fréchette, 2019).[26] The present results show that Grim with forgiveness is present even with no noise in their opponent's behavior, and at reasonably high levels: Repeating the MLE estimates accounting for Grim2 and 3, for the last three super-games, the combined frequency is 19%, equal to that of Grim1.[27] This is eminently sensible since teams were not permitted to communicate with each other, nor did they have any opportunity to publicly commit to Grim. Also identified is the strategy called "Grim with counting", where a team planned to play Grim, but once cooperation was established, chose to unilaterally defect after a number of stage games. Eleven teams did this one or more times, defecting on average between stage-games 9-10. No team did this consistently, almost always reverting to Grim, or Grim with forgiveness after a couple of times. Team chats indicate these defections were rooted in the gambler's fallacy:

> Round 7: "when do you want to go D?" "I say round 10 since they can last this long"
> Round 8: "10?" "9" "cool with me case thats the average amount"
> Round 9: The team in question defected in round 9.
> Round 10: "okay we gotta stick with b now lets hope it ends soon" "true"

Teams locked into mutual defection for a number of rounds would, occasionally, make an effort to restore, or establish, cooperation. For example, take the following team that started with Grim2 establishing mutual cooperation with a STFT player in round 2, resulting in mutual cooperation through round 9. In round 10 STFT reverted to D, resulting in mutual defection (DD) through round 17. At which point STFT reverted unilaterally to C, re-establishing mutual cooperation for another nine rounds. While patterns of this sort were not common, tending to occur with some experience in longer super-games, they did result in re-establishing cooperation in four out of the eight cases identified.

Table 6 compares SEMF estimates for strategies employed for both teams and individuals. The SEMF estimates have incorporated Grim2 and Grim3 as these are straightforward strategies identified in the chats, as opposed to some of the other strategies that are too complicated, and fragmented, to easily characterize for estimation (e.g., generalized STFT, Grim with counting). Results for Grim 2 and 3 have been pooled, while still estimating Grim1 (standard Grim).

[26] P = 1/8 in most cases.
[27] An interesting side effect of introducing Grim2 and 3 is that the MLE estimate for Always cooperate goes to 0 along with a sizable jump in the estimated frequency of TFT from 12% to 27% over super-games 4-7.

Table 6 Table
Team versus Individual Strategy Frequencies: Maximum Likelihood Estimates with Grim2 and 3 (bootstrap standard errors of the mean in parentheses)

Early (SGs 1-3)	Always Defect	Grim	Grim2 and 3.	TFT	STFT	Always Cooperate	WSLS
Individuals	30*** (6.6)	11 (7.1)	6 (4.0)	30*** (8.5)	22** (9.0)	1 (2.3)	1
Teams	44*** (10.4)	2 (6.3)	2 (4.0)	32*** (12.5)	20** (8.1)	0 (0.1)	0
Late (last 3)							
Individuals	15*** (5.5)	10** (5.2)	24*** (8.7)	33*** (7.8)	16*** (5.2)	2 (4.4)	0
Teams	19* (10.0)	19 (18.6)	19** (9.6)	27* (16.2)	15** (5.9)	0 (0.0)	0

Always Defect is substantially lower for individuals to begin with (30% vs 44%), consistent the probit estimates that teams are less cooperative to begin with. In both cases Always Defect is halved by super-games 4-7, with the differences between teams and individuals essentially eliminated. Grim 2 and 3 are estimated with high frequency over the last 3 super-games, with negligible frequency to begin with, in both cases, so that it takes some experience for these to develop. One can compare the results here, with Grim 2 and 3, to those reported earlier for teams (Table 2), and in Table 1A in the appendix for individuals. In both cases adding Grim 2 and 3 has a strong impact on estimates for some of the other strategies: Always Cooperate is at 20% (12%) for individuals in super-games 4-7 without Grim 2 and 3 and 2% (0%) with then included. This is much more in line with the team chats, where we found no team planning to always cooperate. TFT also increases substantially, from 19% (11%) for individuals (teams) to over super-games 4-7 to 33% (27%) for individuals (teams) with Grim 2 and 3 compared to without.

As a final note, we report strategies for the 10 teams whose discussions showed some familiarity with the prisoner's dilemma game. Of these, 5 started with Always Defect, 4 with Grim and 1 with STFT. Three of the 5 starting with Always Defect continued to do so throughout the session, with the other 2 switching to Grim.[28] Those choosing Always Defect had, apparently, never covered repeated play games in their classes, or missed that day's class:

> 206: we should choose D …haha yeah its called the prisoner's dilemma from nash equilibrium … You ever learn about that in econ?

On a more serious note, initial cooperation rates for these teams is only slightly higher than for teams as a whole.

Summary: Comparing strategies from the team chats to the SFEM estimates, the two are broadly consistent. This provides support for the strategy estimation method. Always Defect decreases over time in favor of one or more of the cooperative strategies – Grim, TFT, and generalized STFT. The team chats also reveal a number of variations not captured by the strategy estimation method. These include Grim with forgiveness for early stage game defections, Grim with counting (what stage game to defect in), a variety of patterns where teams defect in the first stage game, after which they try to achieve joint cooperation if their opponent cooperates early on, and no evidence for Always Cooperate. The SFEM estimates show that teams are much more likely to play Always Defect in early super-games compared to individuals. Students with some background in PD games from their studies, are only slightly more likely to play one or another of cooperative strategy early on, compared to students who reveal no formal educational background regarding PD games.

IV. 3 Finite versus Infinitely Repeated PD Games:

Figure 2, reported earlier, showed the typical decline in cooperation rates from the first to the last stage game within supergames, with effectively zero cooperation at the end. In contrast cooperation rates are relatively across stage games for IRPD games, as the super-game end point is randomly determined. While neither result is surprising, it is important to confirm it.

Eyeballing the data in Figure 2, suggests that stage 1 cooperation rates are the same, or higher, under FR compared to IR games. Table 7 reports these differences for the first stage game in supergame 1 and for the last common supergame.[29] For both teams and individuals,

---

[28] The Generalized STFT team provided out favorite quote: "i think we should choose b everytime" "Game theory says we should choose A" "because theortically we could make the most that way" "But this is Trump's america lol and I was thinking B too" "i think the other team will do the same thats why"

[29] Three of the 5 FRPD games with teams had 7 super-games, with 9 and 10 super-games, respectively, in the other two. All five of the FRPD games with individuals had 10 super-games. So the last common super-game means ….

stage 1 cooperation rates are higher in supergame 1 for FRPD games ($p < 0.10$ for individuals). Stage 1 cooperation rates were also higher for FR games in the last common super-game, but the differences are smaller, with none statistically significant. The positive stage 1 cooperation rates for FRPD games is not new in the literature, and can be rationalized in terms of standard theory given a small percentage of "crazy" types who always cooperate, or who are conditional cooperators (Kreps et al., 1982; Reny, 1992).

Table 7: Stage 1 Cooperation Rates

	Individuals			Teams		
	Finite	Infinite	Diff	Finite	Infinite	Diff
Super-Game 1	61.5%	47.1%	14.4%*	65.4%	59.6%	5.8%
Last Common Super-Game	65.4%	59.6%	5.8%	68.6%	65.5%	3.1%

That mutual cooperation rates are higher in stage 1 for the FRPD games is contrary to Dal Bó (2005) who compared FRPD with IRPD games with the same stage game payoffs. There, after several super-games, stage 1 cooperation rates were consistently lower for FR games. However, the continuation value for these games was much lower than the one employed here - .1 and .3 versus .9 here. This has two effects: First, any unraveling from the last stage game will necessarily take longer here, given the limited backward induction teams and individuals exhibit in FRPD games.[30] Second, the larger the number of FR stage games the greater the incentive for fully rational types to imitate the "crazies" early on. Whether the present results would hold up over longer number of super-games remains to be determined.

Summary: Comparing FRPD with IRPD games with the same stage game payoffs and the same expected number of stage games, FR games show the typical reduction in cooperation rates from early to later stage games. In contrast, cooperation rates for IRPD games remain stable across stage games as would be expected. This holds for both teams and individuals. Average cooperation rates are the same or higher in stage 1 games in both the first and last common

[30] There is a large status quo bias (no change) in response to defection by opponents in later stage games. In addition, 70% of the agents who do defect earlier do not account for their opponent doing the same, with the remainder going no further than two steps of backward induction (see Kagel and McGee, 2016, for details).

super-game for FR compared to IR games. These results differ from Dal Bó (2005), who showed that with experience, stage 1 cooperation rates were lower for FR games. This difference in results can be attributed to the much higher expected number of stage games in each super-game here compared to Dal Bó.

V. Summary and Conclusions

This paper reports results from IRPD games comparing individuals to two person teams. Consistent with previous results from IRPD games that satisfy risk dominance and a BAD of less than .50, cooperation rates are high and increasing over time in both cases. Further, teams are less cooperative to begin with, with a substantially higher rate of Always Defect over the first three super-games than for individuals, consistent with the "discontinuity effect" reported in the psychology literature for PD games (Wildschut et al., 2003; Wildschut and Insko, 2007). However, after these initial super-games, cooperation rates are the same, or higher, for teams compared to individuals. The higher rate of defection in early super-games for teams can be attributed to seeking safety from the possibility of the sucker payoff if they cooperate, as well as myopia, with teams focusing on the fact that in any given stage game payoffs for defection are higher regardless of what their opponent does. This essentially ignores the repeated nature of each super-game.

Analysis of the team chats provide an alternative basis for understanding strategies compared to the strategy estimation method (Dal Bó and Fréchette, 2011; Furdenberg et al., 2012) or having agents choose from a set of predefined strategies (Cason and Mui, 2018; Dal Bó and Fréchette, 2018b). The chats show that absent some sort of commitment device, teams will change strategies within a super-game, while identifying a number of strategies not identified in previous studies. The latter include Grim with counting, planning to defect in later stage games one cooperation is established, and the complete absence of any discussion of Always Cooperating. Many teams move from Always Defect to one of the cooperative strategies with experience, with no consistent changes in the opposite direction. Comparing strategies identified from the team chats to the strategy estimation method, outcomes are remarkably similar at an aggregate level. This is important verification of the strategy estimation method since most PD experiments involve individuals, while teams are expensive, will typically involve fewer super-games, and coding of team discussions is time intensive. Further, once identified, new strategies can be easily incorporated into the strategy estimation method to allow for greater forgiveness than strict Grim or TFT (see Fudenberg et al. 2012).

Comparing stage 1 cooperation rates in both early and later super-games, they are the same or higher in FR games for both teams and individuals. In contrast, with experience Dal Bó (2005) found lower stage 1 cooperation rates for FR games. This difference can be attributed to the substantially higher continuation value employed here relative to Dal Bó. There is a clear need for more comparisons of this sort with a larger number of super-games.

This study does not fit easily into arguments that teams are more rational than individuals since, given the folk theorem, IRPD games have multiple equilibria with both low and high cooperation rates. If cooperation rates are the measure of rationality when they are supported by risk dominance or the size of the basin of attraction for always defect, then individuals are smarter to begin with, but teams learn faster.

References

Aoyagi, Masaki, V. Bhaskar and Guillaume R. Fréchette. 2019. "The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private" *American Economic Journal: Microeconomics.*

Aumann, Robert J., and Lloyd S. Shapley. 1994. "Long-term Competition – a Game Theoretic

Analysis." Essays in Game Theory in Honor of Michel Maschler, edited by N. Megiddo, Springer, New York, 1994, pp. 1-15

Blonski, Matthias, and Giancarlo Spagnolo. 2004. "Prisoners' other Dilemma." http://www.gianca.org/Index.htm

Cason, Timothy N. and Vai-Lam Mui. 2018. "Individual versus Group Choices of Repeated Game Strategies: A Strategy Method Approach." *Purdue University Working Paper*

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

Dal Bó, Pedro. 2005 "Cooperation under the shadow of the future: experimental evidence from infinitely repeated games." *American Economics Review*, 95(5):1591–1604.

Dal Bó, Pedro and Guillaume R. Fréchette. 2011. "The evolution of cooperation in infinitely repeated games" *American Economic Review*, 101(1), 411-429.

Dal Bó, Pedro and Guillaume R. Fréchette. 2018. "On the determinants of cooperation in infinitely repeated games: A survey". *Journal of Economic Literature*, 56 (1), 60-114.

Dal Bó, Pedro and Guillaume R. Fréchette. 2018. "Strategy Choice In The Infinitely Repeated Prisoners' Dilemma." *Working Paper*.

Friedman, James W. 1971. "A Noncooperative Equilibrium for Supergames." *Review of Economic Studies*, 38(113): 1-12.

Fudenberg, Drew, and Eric Maskin. 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica,* 54(3): 533-556.

Fudenberg, Drew, David G. Rand and Anna Dreber. 2012. "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World" *American Economic Review, Vol.102*

Kagel, John H. and Peter McGee. 2016. "Team versus Individual Play in Finitely Repeated Prisoner Dilemma Games." *American Economic Journal: Microeconomics*, 8 (2), 253-276.

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.

Kugler, T., Kausel, E. E. and Kocher M. G. (2012). Interactive decision making in groups. Are groups more rational than individuals? *Wiley Interdisciplinary Reviews: Cognitive Science, 3,* 471-482.

Lugovskyy, Volodymyr, Daniela Puzzello and James Walker. 2018. "On Cooperation in Finitely and Indefinitely Repeated Prisoner's Dilemma Games" *Working Paper*.

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

Romero, Julian and Yaroslav Rosokha. 2018. "Constructing strategies in the indefinitely repeated prisoner's dilemma game." *European Economics Review, 2018*, vol. 104, issue C, 185-219.

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

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

Table 1A

Team versus Individual Strategy Frequencies: Maximum Likelihood Estimates with and without Grim 2 and 3

(bootstrap standard errors of the mean in parentheses)

Early (SGs 1-3)	Always Defect	Grim	Grim 2 and 3	TFT	STFT	Always Cooperate	WSLS
Individuals	24*** (3.8)	28*** (3.1)		18** (8.8)	24*** (5.8)	5 (5.5)	1
Individuals with Grim 2 and 3	30*** (6.6)	11 (7.1)	6 (4.0)	30*** (8.5)	22** (9.0)	1 (2.3)	1
Late (last 3)							
Individuals	17*** (3.2)	33*** (5.4)		19** (9.7)	10 (6.3)	20*** (7.2)	0
Individuals with Grim 2 and 3	15*** (5.5)	10** (9.0)	24*** (8.7)	33*** (7.8)	16*** (5.2)	2 (4.4)	0
Early (SGs 1-3)	Always Defect	Grim	Grim 2 and 3	TFT	STFT	Always cooperate	WSLS
Teams	42***	27***		9	18***	3	0

	(3.7)	(2.8)		(10.7)	(4.7)	(10.0)	
Teams with Grim 2 and 3	44*** (10.4)	2 (6.3)	2 (4.0)	32*** (12.5)	20** (8.1)	0 (0.1)	0
Late (last 3)							
Teams (last 3)	23*** (3.3)	51*** (9.0)		11 (16.5)	3 (6.5)	12 (10.8)	0
Teams with Grim 2 and 3	19* (10.0)	19 (8.6)	19** (9.6)	27* (16.2)	15** (5.9)	0 (0.0)	0

Numbers are percentages. *, **, *** significant at the 10%, 5% and 1% levels respectively.