

**Using Team Discussions to Understand Behavior
in Indefinitely Repeated Prisoner's Dilemma Games***

David J. Cooper
University of Iowa and University of East Anglia

John H. Kagel
Ohio State University

Online Appendix

Appendix A: Instructions

Individuals

This is an experiment in the economics of market decision making. The instructions are simple, and if you follow them carefully, you may earn a CONSIDERABLE AMOUNT OF MONEY which will be PAID TO YOU IN CASH at the end of the experiment.

You will be paired with another person choosing between options A and B in the payoff table shown below, where the number listed in the top left hand corner of each box is your payoff and the payoff for the person you are paired with in the lower right hand corner.

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

One round decisions work as follows: Think of yourself as the row player in terms of the payoff table. If you choose A and the other person chooses A your payoff will be 105 and the other person's payoff will be 105. If you choose A and the other person chooses B your payoff will be 5 and the other person's payoff will be 175. If you choose B and the other person chooses A your payoff will be 175 and the other person's payoff will be 5. Finally, if you both choose B, both of you will get a payoff of 75.

When making your choice, you will not know the choice of the person you have been paired with since choices are made simultaneously. After everyone has made their choices, the computer will report back your choice and the choice of the person you have been paired with for that round, along with your payoffs and the other person's payoffs.

All payoffs are denominated in an Experimental Currency Unit (ECU). ECUs will be converted into dollars at a rate to be described below.

Blocks of Rounds/Matches

1. In each match you will be repeatedly paired with the same person in the room. During each match, you will be asked to make decisions over a sequence of rounds using the payoff table just described.
2. The number of rounds in a match, is randomly determined as follows:

After each round, there is a 90% probability that the match will continue for one more round. Specifically, after each round, whether the match continues for

another round will be determined by a random number between 1 and 100 generated by the computer. If the number is lower than or equal to 90 the match will continue for at least another round, otherwise the match ends. For example, if you are in round 2, the probability that there will be a third round is 90% and if you are in round 12, the probability that there will be a tenth round is also 90%. That is, at any point in a match, the probability that the match will continue is 90%.

Earnings in each round of a match depend strictly on your choice and the choice of the person you have been paired with.

3. Once a match ends, you will be randomly paired with another person for a new match. You will not be able to identify the person you've interacted with in previous or future matches.
4. The experiment will end after 12-15 matches have been completed. In each and every match you are paired with a different person. Also, the payoffs for the different choices in each round are always the same.

To make a choice use your mouse to click on the relevant radio button to the right of the table and click the "Send Choice" button. Once you hit the Send Choice button your possible payoffs will be highlighted in red. (*Your choice will not be recorded until you hit the Send Choice button!*) If you want to change your choice, you must do so before hitting the send choice button.

In all rounds you will have up 1 minute to make your choice. (There is a countdown clock on your computer showing you how much time you have left. If you have not made your choice by then you will be prompted do so. Please do not plan on using the full 1 minute to decide what to do if you do not need to as the round proceeds once everyone has made their choices.

Once your choice has been sent, you will move to a waiting screen. (This will be a blank screen with no payoff table.) Once everyone in the room has made their choices, the round will end. When everyone has made their choices, you will see what you chose for that round, what the person you were paired with chose for that round, and your earnings for that round. One of the four possible outcomes shown here.

The outcome screen will be in place for up to 15 seconds. But please plan to click the OK button once you are ready to move on to the next round as this will speed things up.

To summarize: Following each round in a match there is a 90% chance of another round for that match and a 10% chance you will move on to another match with another person.

Payoffs

You earnings for today will be based on the sum total of your earnings in *each and every* round you participate in. ECUs will be converted into dollars at the exchange rate of \$1 = 250 ECU.
Are there any questions?

Please answer the following questions to make sure you understand the structure of the experiment. When you are done, raise your hand and one of us will be around to check your answers.

1. The person you are matched with will stay the same within each match - True / False (circle one).
2. The person you have been matched with will change between matches - True / False
3. If you choose A...
 ...and the other person chooses A, what will your payoff be? _____
 ...and the other person chooses B, what will your payoff be? _____
4. If you choose B...
 ...and the other person chooses A, what will your payoff be? _____
 ...and the other person chooses B, what will your payoff be? _____
5. There is an _____% chance a match will continue for another round.
 This probability is based on how many rounds the match has lasted so far? True/False

Teams

This is an experiment in the economics of market decision making. The instructions are simple, and if you follow them carefully, you may earn a CONSIDERABLE AMOUNT OF MONEY which will be PAID TO YOU IN CASH at the end of the experiment.

Teams

In this experiment, you will make decisions in teams. At the beginning of the experiment, you will be paired with another subject and this subject will be your partner for the entire experiment.

The One-Round Decision

You and your partner will be paired with another team choosing between options A and B in the payoff table shown below, where the number listed in the top left hand corner of each box is your payoff and the payoff for the team you are paired with in the lower right hand corner.

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

One round decisions work as follows: Think of yourself as the row player in terms of the payoff table. If your team chooses A and the other team chooses A your payoff will be 105 and the other team's payoff will be 105. If you choose A and the other team chooses B your payoff will be 5 and the other team's payoff will be 175. If you choose B and the other team chooses A your payoff will be 175 and the other team's payoff will be 5. Finally, if you both choose B, both teams will get a payoff of 75.

When making your choice, you will not know the choice of the other team since choices are made simultaneously. After all teams have made their choices, the computer will report back your choice and the choice of the team you have been paired with for that round, along with your payoffs and the other team's payoffs.

All payoffs are denominated in an Experimental Currency Unit (ECU). ECUs will be converted into dollars at a rate to be described below.

Blocks of Rounds/Matches

1. In each match your team will be repeatedly matched with another team in the room. During each match, your team will be asked to make decisions over a sequence of rounds using the payoff table just described.

2. The number of rounds in a match, is randomly determined as follows:

After each round, there is a 90% probability that the match will continue for at least another round. Specifically, after each round, whether the match continues for another round will be determined by a random number between 1 and 100 generated by the computer. If the number is lower than or equal to 90 the match will continue for at least another round, otherwise it will end. For example, if you are in round 2, the probability that there will be a third round is 90% and if you are in round 12, the probability that there will be a tenth round is also 90%. That is, at any point in a match, the probability that the match will continue is 90%.

Earnings in each round of a match depend strictly on your teams choice and the choice of the team you have been paired with.

3. Once a match ends, your team will be randomly paired with another team for a new match. You will not be able to identify the team you've interacted with in previous or future matches.
4. The experiment will end after 12-15 matches have been completed. In each and every match your partner remains the same but the team you are paired with will be a new team. Also, the payoffs for the different choices in each round are always the same.

Talking to Your Partner and Making a Decision

As a team you will make decisions jointly. That is, the two of you must decide together what choices to make with your payoffs depending on these choices. To facilitate this, there will be a chat box on your screen to send messages back and forth to each other. Although we will record these messages, only you and your partner will see them. You should use this chat box to discuss your strategy and come to an agreement regarding what choice to make. (*E points to where the chat box is and where messages show up both sent and received.*)

To make a choice use your mouse to click on the relevant radio button to the right of the table and click the "Send Choice" button. Once you hit the Send Choice button your possible payoffs will be highlighted in red. (*Your choice will not be recorded until you hit the Send Choice button!*) Once this is done your partner will see your choice in the grey radio button to the left of the payoff table. (*E points to this on slide 1*). Similarly, once your partner has chosen you will see his/her choice. Once you and your partner have agreed on a choice it will become binding after 5 seconds. If you want to change your choice, you must do so before the choice becomes final.

In the first and second round of each match you will have 2 minutes to discuss what to do with your partner and to coordinate your choices. After that you will have 40 seconds to do the same. There is a clock on your choice screen telling you how much time is left for that round.

If time expires and you and your partner have not coordinated your choices, the computer looks to see if one of you has made a choice and use that for your team's choice. And if you have both made choices it will randomly pick one of these as your team's choice. If neither you nor your partner has made a choice before time expires, the computer will automatically send the choice that your team made in the last round. *Please don't plan on using any of these options in making your choices as a round ends after all teams have made their choices, so these options are just designed to deal with "sleepy" teammates.*

Once your choice has been sent, you will move to a waiting screen where you will be able to continue chatting. (This will be a blank screen with no payoff table but with the chat box open.) Once all teams have made their choices, the round will end. When all teams have made their choices, you will see what your team chose for that round, what the team you are paired with chose for that round, and your earnings for that round. One of the four possible outcomes shown here.

The outcome screen will be in place for up to 10 seconds. But please plan to click OK button once you are ready to move on to the next round as this will speed things up.

Note, in sending messages back and forth between you and your teammate we request you follow two simple rules: (1) Be civil to each other and do not use profanity and (2) Do not identify yourself. The communication channel is intended for you to use to discuss and coordinate your choices and should be used that way.

To summarize:

1. Your teammate will remain the same throughout the experiment (same in all matches).
2. Following each round in a match 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.

Payoffs

Your earnings for today will be based on the sum total of your team's earnings in *each and every* round you participate in. You will each get your team's earnings (they will *not* be split between you). ECUs will be converted into dollars at the exchange rate of \$1 = 250 ECU. Are there any questions?

Please answer the following questions to make sure you understand the structure of the experiment. When you are done, raise your hand and one of us will be around to check your answers.

1. Your partner will be the same for the entire session... True / False (circle one)
2. The team you are be matched with will stay the same within each match - True / False
3. The team you are matched with will change at the start of a new match. - True / False
4. If your team chooses A...

 ...and the other team chooses A, what will your payoff be? _____
 ...and the other team chooses B, what will your payoff be? _____
5. If you and your partner choose B...

 ...and the other team chooses A, what will your payoff be? _____
 ...and the other team chooses B, what will your payoff be? _____
6. There is an ____% chance a match will continue for another round.
 This probability is based on how many rounds the match has lasted so far? True/False

Appendix B: Session List

<u>Date</u>	<u>Treatment</u>	<u># Stage Games</u>	<u># Supergames</u>
11/29/16	Team	66	7
12/1/16	Team	52	6
2/14/17	Individual	178	13
2/17/17	Individual	152	13
3/2/17	Team	131	9
3/24/17	Team	127	12
3/30/17	Team	127	12
3/30/17	Individual	152	13
9/20/17	Team	131	10
9/20/17	Individual	170	13
10/1/18	Individual	84	12
10/2/18	Individual	99	13
11/11/21	Silent Partner	112	11
11/15/21	Silent Partner	142	9
11/17/21	Silent Partner	136	10
11/23/21	Silent Partner	112	11

Note: The number of stage games reported is the total number of stage games across all supergames. Only data from common supergames is used in our analysis.

Appendix C: Regression Analysis and Robustness Checks

Regressions: The non-parametric tests reported in Section 4 are useful but conservative and somewhat limited since there are no controls for either the varying length of supergames and sessions or agents' prior experience. The regressions in Tables C1 and C2 correct for this.

Table C1: Regression Analysis, Mutual Cooperation in Stage Game 1

	(1)	(2)
Team, IRPD	0.104 (0.060)	-0.104 (0.078)
Team, FRPD	0.032 (0.093)	-0.221 (0.089)
Silent Partner	0.057 (0.058)	0.009 (0.113)
Lagged # Stage Games	0.002 (0.001)	0.002 (0.001)
Experienced Defection St1, Previous Supergame	-0.242 (0.037)	-0.220 (0.032)
Supergame * IRPD		0.039 (0.015)
Supergame * Team, IRPD		0.038 (0.016)
Supergame * Team, FRPD		0.073 (0.024)
Supergame * Silent Partner		0.009 (0.014)
Log-Likelihood	-692.15	-683.00
Observations	1,233	1,233

Table C1 reexamines Observation 1, that mutual cooperation in St1 increased faster with experience for teams than for individuals. The dependent variable is whether the outcome for the *first* stage game of a supergame (St1) is mutual cooperation. This is a binary variable, so a probit model is used. Marginal effects are reported. There is one observation per supergame, with standard errors clustered at the session level. The dataset includes the FRPD data from Kagel and Magee (2016). This allows us to confirm that Observations 1 and 2 also hold for FRPD games as claimed in Observation 8. Data from the silent partners treatment is also included, making it possible to confirm Observation 9. Both regressions include dummies for the supergame and seed

class, the length of the previous supergame,¹ and agents' experience with defection in St1 of the *previous* supergame.²

Beyond these standard controls, Model 1 has dummies for the team treatments in the FRPD and IRPD treatments along with a dummy for the silent partners treatment.³ Team play had a weak positive effect on mutual cooperation in St1 of the IRPD games. The surprise is that any effect is detected since teams started out less cooperative than individuals, becoming more cooperative over time. Model 2 accounts for these dynamic effects, adding interaction terms between the three treatment dummies and the supergame. The interaction between the team treatment for the IRPD games and supergame is positive and significant, confirming Observation 1 while controlling for a number of potential confounds: mutual cooperation in St1 increased significantly faster for teams than individuals. Observation 1 extends to the FRPD games; the interaction between the team treatment and supergame is positive and significant for FRPD games. The corresponding interaction term for the silent partners treatment is small and does not approach statistical significance.

Table C2 confirms Observation 2 that play was more stable for teams than individuals, and shows that this result also holds for the FRPD games. The dependent variable is the number of switches (as defined in the text) that took place within a supergame. A tobit model is used since the number of switches is constrained to be non-negative. There is one observation per supergame, and all supergames are used regardless of length. Standard errors are corrected for clustering at the session level. All regressions include controls for the outcome in St1 with mutual cooperation as the omitted category, the length of the previous supergame, whether the two agents experienced defection in the first stage game of the previous supergame, and the length of the current supergame,⁴ as well as dummies for the supergame and seed class.

Model 1 uses data from all supergames, while Models 2 and 3 use data from the early (SG 1 – 3) and late supergames ($SG \geq 4$) respectively. The variable of greatest interest is “Team, IRPD” which captures the difference between individuals and teams in the IRPD games. This is negative and significant, confirming that play by teams in the IRPD games was stabler than for individuals

¹ For SG1 this is set equal to 10, the expected supergame length.

² This is averaged across the two agents in a pair, and the mean value is used for SG1.

³ The FRPD treatment is treated as a seed class, so a dummy for the FRPD games is not included. The team treatment dummies measure the difference from the corresponding individual treatments (IRPD or FRPD).

⁴ This is interacted with a dummy for the IRPD games, since all FRPD supergames are the same length.

after controlling for the outcome in St1. This is an important point since teams were less likely to start with a Mixed Outcome than individuals (37% vs. 51%) and supergames that start with the Mixed Outcome had more total switches (see Table 2). These disparities do *not* explain the difference between teams and individuals.

Table C2: Tobit Models: Number of Outcome Switches

	(1)	(2)	(3)
Supergames	All	Early (SG 1 – 3)	Late (SG ≥ 4)
Team, IRPD	-0.390 (0.129)	-0.530 (0.181)	-0.314 (0.144)
Team, FRPD	-0.163 (0.085)	-0.316 (0.136)	-0.094 (0.099)
Silent Partner	0.351 (0.286)	0.233 (0.321)	0.401 (0.405)
FRPD Mutual Defection in St1	-1.531 (0.095)	-1.399 (0.132)	-1.667 (0.127)
FRPD Mixed in St1	-0.394 (0.096)	-0.189 (0.172)	-0.502 (0.107)
IRPD Mutual Defection in St1	0.054 (0.205)	-0.065 (0.335)	0.039 (0.236)
IRPD Mixed in St1	0.576 (0.111)	0.222 (0.213)	0.750 (0.147)
Lagged # Stage Games	-0.014 (0.005)	-0.038 (0.022)	-0.014 (0.006)
Experienced Defection St1, Previous Supergame	-0.163 (0.126)	-0.097 (0.304)	-0.127 (0.122)
Number of Stage Games	0.069 (0.011)	0.091 (0.021)	0.060 (0.013)
Log-Likelihood	-2169.41	-785.04	-1374.64
Observations	1,233	444	789

The coefficient for “Team * FRPD” in Model 1 captures the difference in stability between teams and individuals in the FRPD games. The estimate is negative and weakly significant; as for IRPD games, team play is stabler than play by individuals. The parameter for the silent partner treatment is positive but does not approach statistical significance. Once again, there is little difference between play in the individual and silent partner treatments.

Comparing Models 2 and 3, “Team, IRPD” is smaller in the late supergames, but still easily significant. The difference between teams and individuals shrank with experience in the IRPD games, but did not disappear. The difference also shrank with experience for FRPD games. Unlike

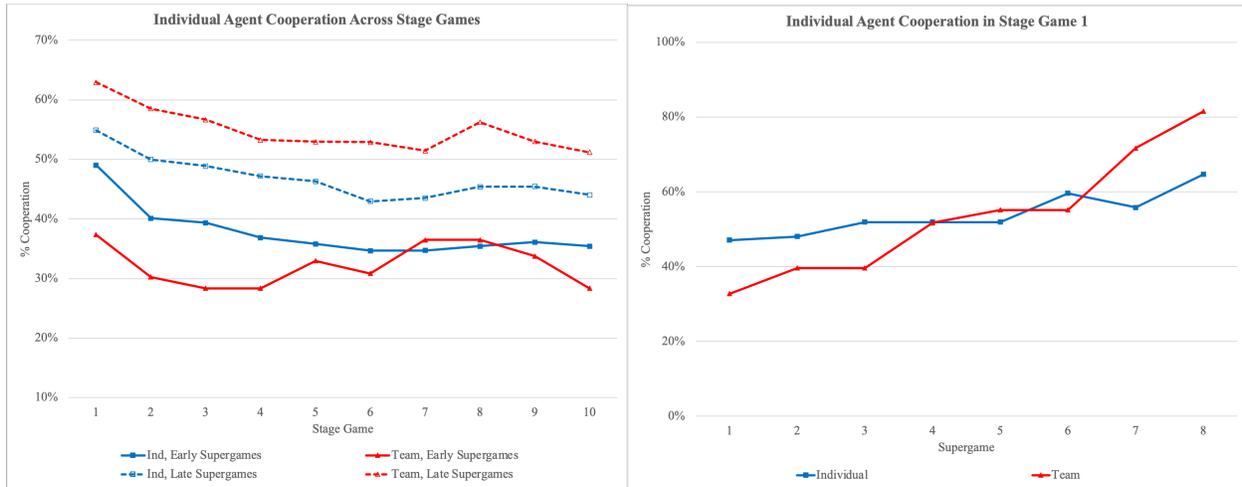
the IRPD games, the difference is sufficiently weak in late supergames that it is nowhere close to statistical significance.

Table C2 addresses stability within supergames, but we have also run probit regressions looking at stability between supergames: if an agent started one supergame with cooperation (defection), how likely were they to start the next supergame with defection (cooperation)? We answer this question via regressions that parallel those in Table C2. The unit of observation is an individual agent, and the dependent variable is a dummy for whether the agent switched their action, cooperate or defect, between St1 of the previous supergame and St1 of the current supergame. In a model without an interaction between the team treatment and supergame (like Model 1 in Table C2), the estimated marginal effect for the team treatment is negative and statistically significant (est. = -0.083; s.e. = 0.020; $p < .001$). When the interaction term is added, as in Model 2 in Table C2, the dummy for the team treatment is still negative but no longer significant (est. = -0.040; s.e. = 0.036; $p = 0.288$), but the interaction term is negative and significant (est = -0.009; s.e. 0.003; $p = 0.012$). These regressions are consistent with our claim that team play was significantly more stable across supergames, with the difference becoming larger with experience. This does not hold in the FRPD games. The dummy for the team treatment is *positive* and significant for the FRPD games (est = 0.279; s.e. 0.116; $p = 0.006$) and the interaction term is negative and significant (est = -0.056; s.e. 0.019; $p = 0.004$). Teams are initially *less* stable between supergames than individuals in the FRPD games, but this flips with experience. There are no significant differences between the individual and silent partners treatments in the IRPD games.⁵

Robustness: The analysis of the data underlying Observations 1 and 2 is based on mutual cooperation in St1 by pairs of agents playing an IRPD game against each other. This is not the only metric we could have used; natural alternatives include using the cooperation rate by individual agents or using data from all stage games. The body of the paper explains why mutual cooperation in St1 is the best metric in our opinion, but also notes that the choice of metric is not terribly important as the obvious metrics are all highly correlated. The purpose of this appendix is to document that Observations 1 and 2 are robust to the use of different metrics.

⁵ Neither the dummy for the silent partner treatment (est = 0.024; s.e. 0.067; $p = 0.714$) nor the interaction term (est = -0.012; s.e. 0.016; $p = 0.460$) is significant.

Figure C1: Cooperation by Individual Agents in IRPD Games



Recall that Observation 1 states, “Mutual cooperation increased faster with experience for teams than for individuals.” Figure 1 showed data supporting this conclusion based on mutual cooperation in St1. Figure C1 is the parallel figure using individual agents’ cooperation rates rather than mutual cooperation rates. The same patterns seen in mutual cooperation are readily apparent. Individual agents’ cooperation rates were initially higher for individuals than teams, but this flipped with experience.

The regressions in Table C3 provide more formal evidence that Observation 1 is robust to how cooperation is measured. The first column replicates Model 2 from Table C1. The critical variable for Observation 1 is the interaction term “Supergame * Team, IRPD.” In Model 1, this term is positive and significant, indicating that mutual cooperation increased faster for teams than individuals in the IRPD games.

Models 2 – 4 offer parallel specifications using different metrics for cooperation. Model 2 also uses data from St1, but the measure of cooperation is cooperation by an individual agent rather than mutual cooperation by a pair of agents. The specification is otherwise identical to Model 1. Model 3 uses data from *all* stage games rather than just St1. The measure of cooperation is mutual cooperation by a pair of agents. The only change to the specification is the addition of a control for the stage game.⁶ Model 4 is the same as Model 3 except the measure of cooperation is cooperation by an individual agent rather than mutual cooperation by a pair of

⁶ This is interacted with a dummy for the type of game, IRPD or FRPD. The changes across stage games are obviously quite different for the two types of games.

agents. We have omitted the control variables (length of previous supergame, lagged defection in St1) from Table C1 for the sake of brevity.

Table C3: Observation 1, Robustness Checks

	(1)	(2)	(3)	(4)
Dependent Variable	Mutual Cooperation	Cooperation	Mutual Cooperation	Cooperation
Stage Games	St1	St1	All	All
Team, IRPD	-0.104 (0.078)	-0.169 (0.101)	-0.055 (0.088)	-0.121 (0.094)
Team, FRPD	-0.221 (0.089)	-0.126 (0.094)	-0.100 (0.077)	-0.130 (0.082)
Silent Partner	0.009 (0.113)	0.018 (0.114)	-0.060 (0.089)	-0.019 (0.103)
Supergame * IRPD	0.039 (0.015)	0.014 (0.015)	0.052 (0.013)	0.048 (0.014)
Supergame * Team, IRPD	0.038 (0.016)	0.037 (0.022)	0.021 (0.012)	0.029 (0.014)
Supergame * Team, FRPD	0.073 (0.024)	0.044 (0.023)	0.030 (0.012)	0.035 (0.013)
Supergame * Silent Partner	0.009 (0.014)	0.009 (0.012)	0.008 (0.012)	0.007 (0.014)
Log-Likelihood	-683.00	-1585.34	-7890.20	-16867.03
Observations	1,233	2,473	13,614	27,298

Notes: All models include controls for the lagged # stage games and experiencing defection in St1 of the previous stage game. Models 3 and 4 include controls for the current stage game.

The main takeaway from Table C3 is that the parameter estimate for “Supergame * Team, IRPD” is always positive and significant. It is worth noting that the parameter for “Team, IRPD” is always negative, but only significant in one of the four regressions. There is an initial discontinuity effect (teams cooperate less than individuals) in the data, but it is not especially strong. Overall, Observation 1 does *not* depend on the details of how cooperation is measured.

Observation 2 states, “Play was more stable for teams than individuals, both within supergames and between supergames.” Stability is defined at the level of outcomes for a pair of agents playing an IRPD game. A “switch” occurs when the mutual outcome (Mutual Cooperation, Mutual Defection, or Mixed) for the current stage game differs from the outcome in the previous stage game within a given supergame. Alternatively, stability can be defined at the level of choices by an individual agent, with a switch occurring whenever an agent changes between C and D. Table C4 reproduces Table 2 from the data, except using switches in individual agent choices as

the measure of stability rather than switches in mutual outcomes. Note that the data is still broken down by the initial *mutual* outcome for the supergame and data is again only included from supergames that lasted at least three stage games.

Table C4: Number of Switches per Supergame

		Individual	Team
Mutual Cooperation (CC)	Average	0.59	0.37
	# Obs	200	140
Mutual Defection (DD)	Average	1.21	0.26
	# Obs	158	104
Mixed (CD)	Average	1.43	0.86
	# Obs	330	144
All Observations	Average	1.14	0.52
	# Obs	688	388

The conclusion from Table C4 match those from Table 2. Play was less stable for individuals than for teams. The level of stability varied depending on the initial outcome for the supergame, but there were always more switches for individuals than teams regardless of the initial mutual outcome.

Table C5 provides formal evidence that Observation 2 does not depend on how stability is measured. Model 1 in Table C5 reproduces Model 1 from Table C2. The key variable is “Team, IRPD.” The negative estimate for this variable indicates that play was stabler for teams than for individuals. Model 2 replicates Model 1 with a different dependent variable. Rather defining a switch as a change in the mutual outcome for a pair of agents, a switch is defined as a change between C and D for an individual agent. The specification is otherwise unchanged from Model 1. In particular, the dataset includes *all* common supergames regardless of length. The control variables (initial outcome, length of previous supergame, lagged defection in St1, length of supergame) from Table C1 are omitted in the interest of brevity.

The main takeaway from Table C5 is the lack of qualitative differences between Models 1 and 2. Specifically, the number of individual switches was significantly lower for teams in the IRPD games. Once again, this finding also held for the FRPD games (and was actually somewhat stronger). Regardless of how switching is measured, there was never a significant difference

between the individual and silent partner treatments. To summarize, Observation 2 does not depend on what measure of stability is used.

Table C5: Tobit Models: Observation 1, Robustness Checks

Dependent Variable	(1)	(2)
	Switches in Mutual Outcome Pair of Agents	Switches in Cooperation Individual Agents
Team, IRPD	-0.390 (0.129)	-0.413 (0.119)
Team, FRPD	-0.163 (0.085)	-0.210 (0.081)
Silent Partner	0.351 (0.286)	0.272 (0.249)
Log-Likelihood	-2169.41	-4105.85
Observations	1,233	2,473

Appendix D: Summary of Coding Categories

Table D1: Game-by-Game Coding: Summary of Coding Categories

- 1) Current Action (12.9%; $\kappa = 0.999$)
 - a. C discussed (61.2%; $\kappa = 0.914$)
 - b. D discussed (57.9%; $\kappa = 0.902$)
- 2) Strategy (3.1%; $\kappa = 0.981$)
 - a. Always Defect (32.1%)
 - b. Always Cooperate (0.8%)
 - c. Grim Trigger (28.2%)
 - d. Grim 2 and Grim 3 (11.1%)
 - e. Grim w/ Counting (4.7%)
 - f. TFT (3.6%)
 - g. TFT variations (e.g. TF2T, TF3T, 2TF2T, 2TFT) (0.0%)
 - h. Suspicious TFT (3.1%)
 - i. Win Stay, Lose Shift (0.0%)
 - j. Signaling (15.0%)
- 3) Discuss past play (1.3%; $\kappa = 0.987$)
- 4) Discuss future play (2.0%; $\kappa = 1.000$)
- 5) Myopia (0.7%; $\kappa = 0.681$)
- 6) Discuss possibility of mutual gains (1.0%; $\kappa = 0.856$)
- 7) Discuss distrust of opponent (1.0%; $\kappa = 0.808$)
- 8) Confusion (errors, gambler's fallacy) (0.5%; $\kappa = 0.555$)

Notes: Frequencies and Cohen's kappa are reported in parentheses. Frequencies for each category are over all observations, where the unit of observation is a team's conversation prior to choosing an action for a stage game. For sub-categories, frequencies are conditional on being coded for that category (e.g. percentage coded for "Always Defect" subject to being coded for Category 2 by at least one coder).

The team discussions showed that 10 of the 58 teams had some familiarity with PD games from one of their classes. An example of this follows.

St1: "yeah its called the prisoner's dilemma ...from nash equilibrium ... You ever learn about that in econ?"

This is not surprising as PD games are included in the curriculum for a variety of disciplines. But does speak to concerns that behavior in lab experiments using college students may be influenced by what they have learned in class. Of these ten teams, five started with AD, four started with Grim, and one started with STFT, little different from other teams. There is nothing obviously unique about teams who have been exposed to PD games in a class.

Table D2: Team Level Coding: Summary of Coding Categories

Category	All Obs.	Between SG	Within SG	κ
Exploring New Strategy	0.788	0.786	0.792	0.650
Benefits of Mutual Cooperation	0.341	0.369	0.292	0.682
History of Previous Play	0.083	0.024	0.188	0.306
Try to Learn Opponent's Type	0.242	0.310	0.125	0.502
Lead by Example / Signal Intent	0.530	0.405	0.750	0.635

Note: Frequencies for each category are over observations with a switch to cooperation as defined within the text, where the unit of observation is a team's entire conversation leading to the switch.

Appendix E: Unstructured Coding

Borrowing a term from the computer science literature, the method of coding teams' strategies described in Section 5.2 is "supervised." We chose the set of strategies to be coded, reflecting our knowledge of the relevant literatures in game theory and experimental economics. As noted in the text, subjects don't particularly think of strategies as game theorists do - an overarching plan that applies in all possible contingencies. This raises the concern that we may have unwittingly biased the results of the coding exercise by specifying the list of possible strategies.

To address this concern, we carried out an exploratory analysis of the dialogues that was "unsupervised", in the sense that we did *not* propose a list of possible strategies. We hired six undergraduate RAs, none of whom had taken a course in game theory or were familiar with our research, to categorize what strategies were used by teams. Their instructions defined a strategy as follows: "A team's strategy is defined at the level of a supergame. A strategy is a plan for how to make decisions for the supergame. It encompasses the entire supergame, not just one stage game within the supergame." To avoid biasing the RAs, we were careful to *not* give them specific examples of strategies.

Initially, all six RAs independently developed a list of strategies. They were also asked to describe the teams' motivation for choosing these strategies and to provide sample dialogues for each strategy. We then had the RAs meet in two groups of three to formulate unified lists of strategies. We subsequently showed them our list of strategies and asked them individually to compare their group's list with ours.

Table E.1 shows the strategy lists from the two groups (Blue and Yellow). The strategy names and material in quotations are directly from the RAs; further descriptions not in quotations are our summary based on materials provided by the RAs. We have modified the terminology used by the RAs to match what is used elsewhere in this paper (e.g. we substitute "supergame" for "match").

Several points stood out from the RAs' characterizations of strategies. First, the lists of strategies were short for both groups - four for one group, five for the other. This was less than the seven strategies included in Table 3, and the RAs did not identify the large number of slightly differentiated strategies included in most fitting exercises. One of the RAs gave the following explanation for having relatively few strategies when comparing our list of strategies with his

group’s list: “I think a few of your strategy categories could be simplified into to one inclusive category ... I think that lenient grim and grim trigger are of the same category and could thus be simplified into one ‘grim trigger’ category as the only difference between the two is how many defections occur before the team swaps to playing D consistently. They both fit the same category of strategy, just with slight variation in execution. Ultimately, my group ended up reducing any ‘grim trigger’ strategy to a TFT strategy since the two were so hard to tell apart ...” This underlines a point that became clear as we analyzed teams’ dialogues. Most strategies are variations on a theme. All the variants of Grim are closely related implementations of the same basic strategy, and even TFT is not so different in practice. In all cases, the basic rationale is to try cooperation for a while in the hope that one’s opponent will get the hint and also cooperate.

Table E.1: Strategies for RA Groups

Team	Strategy Name	Description
Blue	Always Defect	"Always defect for the entire supergame." This is equivalent to AD.
	Hesitant Cooperation	"Defect the first stage game. If the opponent cooperated, signal cooperation with C for one or two stage games. Mirror opponent’s behavior." This category is roughly equivalent to STFT, although the team members’ comments make it clear that ‘mirroring’ refers to how the team will respond to their opponents’ actions after the two rounds of cooperation rather than a general policy of TFT.
	Trusting Cooperation	"Cooperate the first two stage games. Mirror opponent’s behavior." This category encompassed variants of Grim and TFT. Judging by comments from team members, the description is more specific than what they had in mind; they recognized that not all groups cooperated for exactly two rounds.
	Cooperate then Betray	"Variant of Trusting Cooperation. Cooperate, then defect after a number of stage games." This is equivalent to Grim w/ Counting.
Yellow	Strong C	"Participants believe that C is the best option and plan to consistently choose C." The description makes this sound like AC, but one of the team members’ individual descriptions makes it clearer what was meant: "They would start by selecting C and continue to select C for as long as their opponent also selected C. They would heavily favor C only switching if the other team constantly chose D." This category included all Grim variants, although they seem to have largely had lenient Grim in mind.
	Strong D	"Participants believe that D is the best option and plan to consistently choose D." This is equivalent to AD
	Backstabbing	"Participants start with C with the hopes of eventually changing to D, catching their opponents off guard, for a higher pay off." This is equivalent to Grim w/ Counting.
	Dependent on Opponent	"Participants start with C/D but plan to mirror their opponents actions." This category was a combination of STFT, TFT, and the TFT variants.
	Stubborn Opponent	"Participants are choosing D, but are willing to switch to C. However, they are hesitant to follow through because they fear the potential loss of profit." This roughly matches what we called generalized STFT.

Second, the strategies the groups described are easily matched with categories in our coding, albeit with less differentiation. Both groups clearly identified AD and Grim with Counting. Both groups accounted for variants of TFT, STFT, and Grim; what differed is how they grouped them together. The one exception is generalized STFT; one group identified this, while the other group regarded these cases as representing a mid-game change in strategy (see quote below).

Both groups' definitions of strategies focused on discussions in the initial stages of supergames, paying little attention to how teams responded to their opponent's behavior later in the supergames. Like us, the RAs recognized that teams were not operating with fixed strategies: "Usually in the rounds I saw where a team tried to switch to C mid-way through after defecting, this was not a strategy devised at the beginning but rather a sudden change in plans after realizing the benefits of cooperating part way through the round."

Finally, the rationales for strategies identified by the RAs differ little from what is described above. For example, one group gave the rationale for Always Defect as follows: "Usually involves a lack of trust in their opponent ...". For Trusting Cooperation, a strategy which broadly included all variants of Grim and TFT, they stated: "Belief that scoring a 5 on the first round would be offset by future cooperation over a long period, assuming that the opponent will cooperate."

Obviously, this was a speculative exercise. Based on the results, we don't recommend replacing a more structured coding with having the RAs come up with their own categories. The differences between strategies identified by game theorists may be subtle, arguably too subtle, but details such as how patient agents will be before punishing defection play an important role in determining whether mutual cooperation can be achieved. However, given how little direction we gave these RAs, and their lack of experience with game theory, it is surprising how close their lists of strategies came to ours. This exercise provides some confidence that our findings in the main text are not an artifact of our choice of strategies to include in the coding scheme.

Appendix F: SFEM Estimates and Comparison of Teams vs. Individuals

SFEM models individuals as playing finite automata, capturing common strategies such as Grim Trigger or Tit-for-Tat. Critically, the model includes an error component - every time an action is made the intended action is implemented with probability β and the other available actions with probability $1 - \beta$. The possibility of errors implies that any possible series of actions is played with positive probability by any finite automaton, making it possible to calculate a likelihood function. Using a pre-specified set of strategies, the model calculates the likelihood of each individual/team's observed actions subject to adoption of each possible strategy. The probability distribution over possible strategies is then used to generate a weighted average of the likelihoods of the available strategies. SFEM fits the weights on strategies and the noise parameter via maximum likelihood estimation; specifically the weighted average likelihood over strategies is maximized. SFEM is a mixture model. It does not assign specific strategies to specific individuals or teams. Rather, it estimates the probability distribution of strategies across the entire population.

A critical issue in working with SFEM is determining the set of strategies to include in the model. We use the set of strategies receiving positive weight in Aoyagi, Bhaskar, and Fréchet's (2019) estimation of SFEM for IRPD games with perfect monitoring and $\delta = 0.90$.⁷ Table B1 reports the distribution of strategies estimated by SFEM. The model is estimated separately for the individual and team data, subdivided between early (SGs 1 – 3) and late (SGs 5 – 7) supergames.⁸ Table 6 in the main text combines the estimated weights Grim2 and Grim3 into “lenient grim” and TF2T, TF3T, and 2TF2T into “complex tft.”

The SFEM estimates reflect the two main differences between teams and individuals. Observation 1 notes that mutual cooperation increased faster across supergames for teams than individuals. Underlying this, the estimated proportion of AD decreased with experience for both individuals and teams, but the decrease was almost twice as large for teams (24% vs. 14%). Observation 2, that behavior was more stable for teams, was reflected by the lower estimated error rates for teams, as this is the only way SFEM can capture this feature of the data.

⁷ Aoyagi et al. (2019) include 15 strategies based on achieving statistical significance in earlier papers. Four complex strategies (CDDD, Sum2, 2TFT, and SSum2) received 0% weight in their estimation for IRPD games with perfect monitoring and $\delta = 0.90$. None of these were detected in the coding exercise here, and are not included in our estimates.

⁸ All of the IRPD sessions except one ran for at least seven supergames. We use data from SGs 4 – 6 for this session as well as the matching individual session (the session using the same random seed).

Table F1: SFEM Estimates

	Individual		Team	
	SGs 1 – 3	SGs 5 – 7	SGs 1 – 3	SGs 5 - 7
AD	29.78%	19.30%	43.67%	24.76%
	(6.65%)	(5.71%)	(10.30%)	(9.65%)
AC	1.11%	5.85%	0.00%	0.00%
	(1.84%)	(4.39%)	(0.00%)	(0.00%)
Grim	12.16%	0.00%	2.69%	9.37%
	(7.10%)	(3.10%)	(5.80%)	(9.55%)
Grim 2	4.88%	7.81%	0.00%	0.00%
	(3.67%)	(5.50%)	(0.93%)	(2.49%)
Grim 3	0.00%	0.00%	0.00%	0.00%
	(1.03%)	(3.41%)	(0.43%)	(0.00%)
TFT	23.32%	36.98%	28.74%	30.53%
	(7.90%)	(7.00%)	(11.78%)	(15.68%)
STFT	20.66%	22.10%	19.89%	17.94%
	(8.39%)	(5.87%)	(7.99%)	(5.98%)
TF2T	4.36%	5.79%	0.00%	9.42%
	(2.54%)	(6.86%)	(3.01%)	(8.90%)
TF3T	0.00%	0.00%	0.00%	0.00%
	(2.87%)	(4.20%)	(0.90%)	(0.00%)
2TF2T	2.69%	2.17%	5.01%	7.98%
	(2.80%)	(1.13%)	(5.84%)	(9.80%)
WSLS	1.05%	0.00%	0.00%	0.00%
	-	-	-	-
Gamma	0.398	0.352	0.314	0.284
	(0.050)	(0.025)	(0.024)	(0.022)
p(error)	7.50%	5.52%	3.97%	2.87%
Log-likelihood	-1098.10	-1028.68	-380.59	-372.94

Appendix G: The Truth-Wins Model

The truth wins (TW) model of Lorge and Solomon (1955) applies to problems that have a demonstrably correct solution. The idea is that if any team member solves the problem, they will communicate the solution to their teammates who will grasp that the solution is correct and implement it. Making this more concrete, consider a problem that an individual has probability p of solving. For an n person team trying to solve this problem, the TW model predicts that the probability of a correct solution equals $1 - (1-p)^n$. There is an extensive psychology literature examining the TW model. The general finding is that teams usually fall below the TW baseline and only very rarely exceed it (Davis, 1992). The failure to meet the TW baseline is typically attributed to “process loss.” This accounts for both inefficiencies due to costly and imperfect communication as well as free-riding on the efforts of others. The evidence for games is mixed. In work with signaling games, Cooper and Kagel (2005) find that teams significantly outperform the TW baseline for some cases. However, this does not hold for all games. For example, Casari, Zhang, and Jackson (2016) provide an example in the takeover game where teams not only fail to beat the TW baseline, but actually do worse than individuals. Cooper and Kagel (2016), while replicating the result of Cooper and Kagel (2005), provide evidence that beating the TW baseline reflects positive synergies generated by teammates’ discussions rather than the mechanical explanation underlying the TW model. In an advice treatment where insights can only flow from an advisor to an advisee, advisors who have learned the optimal strategy frequently do not communicate this to their advisee, and advisees who receive correct advice often fail to adopt it. Both observations undermine the TW model. Behaviorally, the TW model is a useful benchmark but misses much of what drives team performance.

The TW model is not obviously germane to IRPD games as there is no correct solution (i.e. no optimal strategy). If we weaken the model to consider empirically optimal strategies, rather than theoretically optimal strategies, it is possible to put together a version of TW and calibrate it to the data. In particular, focus on play in St1. Playing C in St1, consistent with cooperative strategies such as Grim Trigger, TFT, and their variants, leads to higher average payoffs (see fn. 10). Thus, we can define play of C in St1 as the “empirically optimal” action. If we replace an optimal action with the empirically optimal action in the TW model, we can compare individual play, team play, and a variant TW benchmark (with the 90% confidence range) as shown in Figure G1. Initially, team play is below the TW benchmark as cooperation rates for teams are below

cooperation rates for individuals. Over time, as teams become more cooperative, team play just catches up with the lower range of the 90% confidence interval for the TW benchmark. This is consistent with the standard result in the psychology literature that teams rarely meet let alone beat the TW benchmark (Davis, 1992).



Importantly, the variant TW model misses a central feature of the data, faster growth of cooperative play by teams. It instead predicts that cooperative play grows no faster for teams than individuals. The relative fast growth of cooperation in the team data must reflect factors not captured by the TW model, consistent with our observation that the TW model is not germane for IRPD games. The literature comparing individual and team decisions often contrasts eureka and judgment type problems. Eureka problems feature a demonstrably correct solution. While there are important insights to be had in an IRPD game (e.g. leading by example works well), IRPD games do not have an optimal strategy per se as the best response changes as a function of the other player’s strategy. In a judgment problem, there is *not* a demonstrably correct solution and choosing an option depends on an individual agent’s tastes and/or beliefs. This is directly relevant for play of IRPD games. As dialogues in the initial games make clear, an important component of the problem facing agents is balancing the risk of receiving the sucker payoff against the long-

term benefits of achieving mutual cooperation. In other words, the choices facing agents in an IRPD game include aspects of both eureka *and* judgment problems, fitting cleaning into *neither* category. The TW model is meant for problems that fall cleanly into the eureka category. As such, it is to be expected that the TW model would fail to capture major features of our data.