

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

6/28/2022

Abstract: We compare behavior of two person teams with individuals in indefinitely repeated prisoner dilemma games with perfect monitoring. Team discussions are used to understand the rationale underlying these choices, and how these choices come about. There are three main findings: (1) Teams learned to cooperate faster than individuals, and cooperation was more stable for teams. (2) Strategies identified from team dialogues differ from those identified by the Strategy Frequency Estimation Method. This reflects the improvisational nature of teams' decision making. (3) Increasing cooperation was primarily driven by teams unilaterally cooperating in the hope of inducing their opponent to cooperate.

Key words: Infinitely repeated prisoner dilemma games, team decision making, analysis of team discussions

JEL codes: C72, C92, D01

*We have benefitted from discussions with Hal Arkes, Tim Cason, Russel Fazio, Guillaume Fréchette, and Dan Levin. Zeina Aboushaar, Jinsoo Bae, Joe Ballard, Puja Bhattacharya, Caleb Bray, Tobias Brevik, Jonathan Cooper, Adam Edwards, Molly Higgins, Mir Adnan Mahmood, Arthur Nelson, Kirby Nielsen, Luke Ortolani, Isabelle Wilcox, and Basil Wojdula provided valuable research assistance. Helpful comments were received from participants in seminars at UT Dallas, Chapman University, University of East Anglia Workshop on Behavioral Game Theory, UTS conference on experimental economics, Tokyo University, Central Michigan University, 2019 Southern Economic Associations meetings and the 2019 North American ESA meetings. Support was partially provided by NSF grants SES 1630288, SES 1630194, SES 2018704 and SES 2018690. Opinions, findings, conclusions or recommendations offered here are those of the authors and do not necessarily reflect the views of the National Science Foundation.

This paper reports an indefinitely repeated prisoner’s dilemma (IRPD) game experiment where the decision making “agents” were either freely interacting two-person teams or individuals. The primary goals of the paper are twofold: First, to compare the behavior of teams and individuals in IRPD games with perfect monitoring where mutual cooperation is both consistent with equilibrium and risk dominant. Many strategic interactions in economics involve groups of individuals. If the behavior of groups differs substantially from individuals, conclusions based solely on the decisions of individuals will not be broadly applicable. Second, to use team dialogues to understand what motivates agents’ choices and how these come about. Economists have become increasingly interested in using process data (e.g. response times, eye tracking, and fMRI) to understand decision making.¹ Analyzing the content of team dialogues provides direct insights into the thought processes underlying agents’ choices that cannot be easily obtained using other methods. This makes it possible to directly observe *how* and *why* strategies came about. Given that the broad patterns of play were similar for teams and individuals, team discussions presumably offer insights into the motivation underlying individuals’ choices as well.

Comparing teams and individuals, cooperation was rare for both in early supergames, with teams cooperating somewhat less than individuals. Cooperation increased significantly faster *across* supergames for teams than individuals, so that in later supergames teams cooperated at significantly higher rates. *Within* supergames, play was significantly more stable for teams than individuals. Play by teams was also more stable *between* supergames; once a team switched to a cooperative strategy, they rarely switched back.

Teams rarely discussed strategies as a game theorist would. Rather than pre-specifying a full, state contingent plan, strategies were typically incompletely specified, with teams improvising in response to unanticipated choices by their opponents. Despite this, we can almost always identify a strategy that corresponds to their initial plan for making choices within a supergame.

The modal strategy identified from team dialogues in early supergames was Always Defect (AD). Over half of all teams who started out with AD switched to more cooperative strategies, primarily Grim Trigger and its lenient variants, so that in later supergames a majority of teams employed cooperative strategies. Returning to AD, even temporarily, was rare.

Comparing the distribution of strategies identified from team dialogues with estimates based on the Strategy Frequency Estimation Method (SFEM; Dal Bó and Fréchette, 2011), SFEM

¹ See Fehr and Glimcher (2013) and Cooper, Krajbich, and Noussair (2019) for recent surveys.

puts more weight on variants of Tit-for-Tat (TFT, STFT, etc.) than variants of Grim Trigger, while the coding based on team dialogues does the opposite. This reflects the improvisational nature of teams' strategies. Teams did not try to anticipate all possible histories of play. Instead, they used simple plans covering the first few stage games (e.g. "My plan is choose C [cooperate] first for a few times, and if the other team keeps choosing D [defect], we will switch to D.") and then adjusted as needed. Choice patterns that appeared like Tit-for-Tat often reflected improvised reactions to their opponent, attempting to coordinate on mutual cooperation within a supergame, rather than a pre-conceived strategy.

Analysis of team dialogues makes it possible to identify the rationales underlying teams' adoption of strategies. Their frequent initial use of AD was primarily motivated by fear of their opponents defecting and a desire for the safety that AD provides. Teams typically had long conversations prior to first trying a cooperative strategy, with discussions stretching across multiple stage games and even across supergames. Realizing that *mutual* cooperation pays more in the long run, they often tried cooperating for a few stage games to see if the other team would go along. They discussed this in terms of "leading by example," hoping that their actions would send a message when direct communication was not possible. As one team noted, "this is all so hard without communication" "I know if we could just send them like one sentence we'd have it made." Teams varied in the specifics of how they conceptualized and executed this approach, making possible the identification of different strategies, but ultimately these were slight variations on the same basic theme.

Team dialogues also explain the stability of team play relative to individuals. When teammates disagreed about whether to stick with the status quo or make a change, the status quo usually won (87%). Teammates provided a check on switching that did not exist for individuals.

The comparison of teams and individuals raises a question whether behavior differs due to the presence of a teammate per se, or because of joint decision making and the associated communication between teammates. To distinguish between these two possibilities, a silent partners treatment was implemented. Subjects played the game in fixed pairs, like the team treatment, but one teammate was solely responsible for making decisions with no input from their silent partner. Behavior in the silent partner treatment differs little from the individual treatment, indicating that the effect of team play was largely due to a combination of bilateral communication and joint decision making.

The design and procedures used here parallel those employed in Kagel and McGee (2016) to compare individual and team play in finitely repeated prisoner's dilemma (FRPD) games. This makes possible a clean comparison of the differences between teams and individuals in IRPD and FRPD games. There are strong similarities between the two with inexperienced teams less cooperative than individuals in early stages of early supergames, only to become more cooperative with experience in early stage games. Likewise, play by teams was more stable than individuals. The differences between teams and individuals likely reflect underlying differences, present in many settings, between how teams and individuals make choices.

It is inherently interesting to understand *how* and *why* teams chose strategies, but this understanding also has implications beyond the experiment reported here. Teams approached IRPD games with a combination of simplicity and sophistication. Rather than fixing a detailed plan in advance, they typically start with a simple, incomplete, initial plan followed by improvisation in response to the behavior of their opponent. At the same time, leading by example showed a sophisticated ability to anticipate and manipulate other agents' behavior. The simple, flexible approach taken by teams is portable. Rather than only applying to the environment studied here, it is an approach that seems easily applied to other repeated games. The sophistication exhibited by teams and their tendency to improvise has important implications for how learning in IRPD games should be modeled, an issue discussed at length in the conclusion.

In terms of technique, the approach employed here for coding team dialogues is a departure from how economists have previously coded communication. Most previous coding exercises have been done at a granular level, but this would have missed vital content from conversations that extended over multiple stage games or even supergames. Identifying what strategies were being used (and why) required synthesizing the content of long-running conversations, in conjunction with choices, rather than focusing on what was said at any single point in time. These procedures are likely to be applicable to other experimental settings as well.

The outline of the paper is as follows. Section I briefly reviews past research with IRPD games. Section II reports the experimental design and procedures, and Section III lays out hypotheses for the data based on the relevant theory and the preceding literature. Section IV compares differences between teams and individuals for IRPD games. Section V describes the main features of the team dialogues, explains how strategies are identified from team discussions, investigates how strategies evolve with experience, compares the strategies identified from team

dialogues with the distribution of strategies estimated by SFEM, and examines the rationales underlying teams' choices. Section VI contrasts the results reported here with results from previously reported FRPD games. Section VII reports results from the silent partner treatment. Section VIII summarizes the results and discusses the implications of what has been reported.

I. Prior Research: There have been numerous experimental studies of IRPD games. Dal Bó and Fréchette (2018) provide an extensive survey of the experimental literature. This paper departs from the existing literature along two important dimensions. First, most of the economics research on IRPD games involves individuals making decisions. The focus here is on teams, specifically on the processes and motivation for their choices. Cason and Mui (2019) is a notable exception to the focus on individuals, comparing play in three person teams with individuals. Most of their analysis considers games with imperfect monitoring, with data primarily based on subjects' choices over a menu with predefined strategies (e.g., tit-for-tat or grim trigger). In a control treatment with perfect monitoring (like this experiment), *both* teams and individuals achieved very high cooperation rates (80 - 90%). These high cooperation rates reflect payoff matrices and continuation probabilities designed to achieve reasonable cooperation rates with imperfect monitoring, but leave little room for distinguishing between individuals and teams with perfect monitoring.² Teams in Cason and Mui could discuss their strategy choices. The paper describes few details of their formal coding of these discussions, which only play a minor role in the paper, but they do make use of samples of team discussions to illustrate features of team play. Identifying strategies from chat is a non-issue due to their use of direct elicitation.³ The focus of Cason and Mui, compared to the present paper, is quite different. More than anything else, the present paper is about what can be learned from the team chats. Cason and Mui is largely concerned with what strategies teams chose compared to individuals.

The second important difference between this paper and the existing literature is the use of team dialogues to identify what strategies are used and why. The most common method for identifying strategies is the Strategy Frequency Estimation Method (SFEM) introduced in Dal Bó and Fréchette (2011). This technique uses maximum likelihood estimation of a mixture model to

² They have a substantially smaller basin of attraction for Always Defect (SizeBAD): .05 compared to .26 in our design. See Section III for a definition and discussion of SizeBAD in relation to cooperation rates.

³ They report relationships between appeals to game theoretic reasoning and choice of AD and between concerns about the effect of noisy implementation of chosen actions and the use of lenient strategies.

estimate the distribution of strategies across the population. There are also a small number of papers that use direct elicitation, asking agents to specify a full strategy, either directly or by making a choice from a menu, at the beginning of each supergame (Romero and Rosokha, 2018, 2019; Dal Bó and Fréchette, 2019; Cason and Mui, 2019).

Each method of identifying strategies in IRPD games has its strengths and weaknesses, and which is best depends on the purpose of the exercise. Econometric approaches like SFEM are unintrusive, but only indirectly measure what strategies were used based on a predefined set of strategies. Such methods are best for estimating strategies employed when the researchers run a standard experiment with individuals making direct choices. This makes sense when identification of individual agents' strategies is not the primary purpose of the paper. Direct elicitation, by its very nature, perfectly identifies what strategies were used, but is the most intrusive method. Direct elicitation often limits subjects to a pre-specified list of strategies. Even in cases where any strategy can be constructed, the exercise implicitly directs subjects towards completely specified strategies rather than the incomplete strategies identified in our experiment.⁴ Direct elicitation is the best approach if the goal is to unambiguously identify strategies. Analysis of team chat lies somewhere between econometric methods and direct elicitation. Making decisions in a team obviously affects how choices are made, but discussions between teammates are a natural part of this process and recording the chats is non-intrusive. Identification of strategies from team chat is more direct than econometric methods, but does involve judgment by the coders, making it less accurate than direct elicitation. The true advantage of analyzing team dialogues is the ability to understand how and why decisions were made. If the goal is to better understand the nature of strategies and the rationale underlying agents' strategies, chat analysis is likely to be the best approach.

There is an important line of research in the social psychology literature concerned with differences between individuals and teams in PD games. These experiments are quite different in structure from those commonly employed in economics. They typically involve a single supergame where agents are told they will be paired with the same opponent for between t and $t+n$ stage games, with the actual stopping point occurring at an unknown point in that interval. The

⁴ Romero and Rosokha (2018, 2019) allowed subjects to construct strategies in an almost unlimited fashion. In the latter paper, subjects could also revise their strategies midgame. In both cases, strategies are complete plans for how the game is to be played. Even if this plan can be revised, having subjects specify complete plans presumably has an effect on how subjects conceive of strategies. Dal Bó and Fréchette (2019) included a phase in which subjects need not follow their chosen strategies. The same concern about framing applies.

key finding from this literature is that teams are less cooperative than individuals (referred to as the “discontinuity effect”). This is attributed to greater fear and greed on the part of groups than individuals (see Wildschut et al., 2003 and Wildschut and Insko, 2007 for surveys).

Beyond IRPD games, there is a growing literature in experimental economics comparing how individuals and teams behave in games (e.g. Cooper and Kagel, 2005; Kocher and Sutter, 2005; Feri, Irlenbusch, and Sutter, 2010; Maciejovsky, Sutter, Budescu, and Bernau, 2013; Casari, Zhang, and Jackson, 2016). The typical finding is that teams are more “rational” than individuals; specifically, teams are more likely to use a theoretically optimal strategy, and are faster to learn how to maximize their payoffs. There is no optimal strategy in IRPD games, but, given that cooperative strategies earned higher payoffs than AD, finding that teams switched to cooperative strategies faster than individuals is consistent with this literature.

II. Experimental Design and Procedures: Throughout this paper “agent” is used as a generic term for either individuals or two person teams. Agents played a simultaneous move, indefinitely repeated prisoner’s dilemma (IRPD) game with perfect monitoring using the stage game payoffs reported in Table 1 (own payoffs are in red, opponent’s payoffs are in blue).

Table 1: Stage Game Payoffs

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

The continuation probability was $\delta = 0.90$, yielding an expected supergame length of ten stage games. Following each supergame, agents were randomly re-matched with the restriction that no two agents were matched in consecutive supergames. The instructions stressed that at the end of each stage game “there is a 90% chance of another round [stage game] for that match [supergame] and a 10% chance you will move on to another match with another team/individual.” After the last stage game within a supergame, agents were notified that their match had ended and that they would be starting another match with a different (randomly chosen) agent. Judging from the team chats, agents had no difficulty telling when they were starting a new supergame with a different opponent.

Using a between subjects design, all agents in a session were either individuals or teams with no mixing between the two. There were six individual agent sessions with 14 – 18 subjects in each session (104 total subjects). Sessions lasted 90 minutes, conducting as many supergames as possible within the allotted time. Four sessions had 13 supergames, the other two had 12. Individuals had up to one minute to make their choice, but this was never a binding constraint.

In the team treatment, subjects were randomly matched at the beginning of the session. Teammates remained the same throughout the session. Instructions were essentially the same as for individual sessions (see online Appendix A), except that teams were told to “make decisions jointly.” Each teammate could enter a choice, with the team’s choice implemented if choices agreed and did not change for five seconds. In the first two stage games of each supergame, teammates had two minutes to discuss their choices and reach an agreement. This was reduced to 40 seconds for all subsequent stage games. The default options, in case teammates could not agree, are reported in the instructions (see online Appendix A). Teammates reached an agreement on what action to take in the overwhelming majority of cases ($> 99\%$). To facilitate reaching agreement, teammates could send messages back and forth. They were told 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, not use profanity, and to not identify themselves. The instructions stressed that other teams could not see their messages.

There were six team sessions with between eight and twelve teams in each session for a total of 58 teams (116 subjects). The first two sessions had only six and seven supergames respectively; these were scheduled to last two hours as the amount of time the team chats would take was not anticipated. Subsequently, session time was increased to two and a half hours, along with modest reductions in the time teammates could discuss their choices.⁵ The remaining four team sessions had between nine and twelve supergames.

The six individual and six team sessions were paired, using matching seeds to generate the random length of the supergames (see Appendix B for a session list). This guaranteed parallelism between the two treatments in terms of the length of supergames. The main reason for variability

⁵ In the first two sessions, teams had two and a half minutes to discuss their choices in the first two stage games, with one minute after that. Subject feedback indicated that shorter times were more than adequate to reach agreement, so teams in the remaining sessions were given two minutes to discuss their choices in the first two stage games, with 40 seconds after that.

in the number of supergames across sessions, beyond the timing issue mentioned above, is that the different seeds led to different lengths of supergames.

Payoffs were denominated in experimental currency units (ECUs) which were converted into dollars at the rate of \$1 = 250 ECUs. Payoffs were summed over all stage games of all supergames, converted to dollars, and paid in cash at the end of an experimental session. Each member of a team received the team's full payoff. Earnings averaged \$44.98 per subject, with a typical session lasting 90 minutes for individuals and 135-150 minutes for teams. Subjects were recruited from the Ohio State University undergraduate population using ORSEE (Greiner, 2015). The software was programmed using zTree (Fischbacher, 2007).

III. Predictions: With sufficiently patient agents, IRPD games with perfect monitoring have an infinite set of subgame perfect equilibrium outcomes.⁶ Cooperative outcomes are consistent with equilibrium, but non-cooperative play is an equilibrium as well. A number of criteria have been explored to better predict when cooperation is more likely, the most popular of which reduces the game to a normal form game with only grim trigger (Grim) and always defect (AD) available as strategies. SizeBAD (Dal Bó and Fréchette, 2011) is defined as the size of the basin of attraction for AD. SizeBad measures the dynamic stability of cooperation. If SizeBad < 1, mutual play of Grim is a subgame perfect equilibrium, and, if SizeBad < 0.50, mutual play of Grim is risk dominant (Harsanyi and Selten, 1988). Empirically, cooperation rates are a decreasing function of SizeBad (Dal Bó and Fréchette, 2011 and 2018).⁷ The combination of stage game payoffs and continuation probability used here yields SizeBAD = .26; cooperation is risk dominant, but not so strong that we'd expect universal cooperation.

As noted previously, the experimental literature finds that teams are more likely than individuals to use a theoretically optimal strategy in a game, and are faster to learn how to maximize their payoffs. The former is of no help in making predictions, since the benefits of specific strategies depend on what equilibrium is played, but the latter implies that mutual cooperation rates will increase faster for teams. The discontinuity effect documented in the psychology literature, as described above, suggests that cooperation rates will initially be lower for teams.

⁶ Friedman (1971), Aumann and Shapley (1994), and Fudenberg and Maskin (1986).

⁷ Also see Lugovskyy, Puzzello, and Walker (2018) and Blonski, Ockenfels, and Spagnolo (2011).

Finally, previous experiments find a strong relationship between response times and cooperation in one-shot games (Rand, Greene, and Nowak, 2012), with more cooperative individuals having slower response times. This is attributed to cooperative play requiring greater deliberation, and remains true when individuals are prompted to respond either more or less rapidly. Team play forces slower decisions due to the need to discuss the team’s choices. As such, teams might be expected to be more cooperative. This finding refers to the level of cooperation, and does not deal with whether cooperation will grow more rapidly for teams versus individuals.

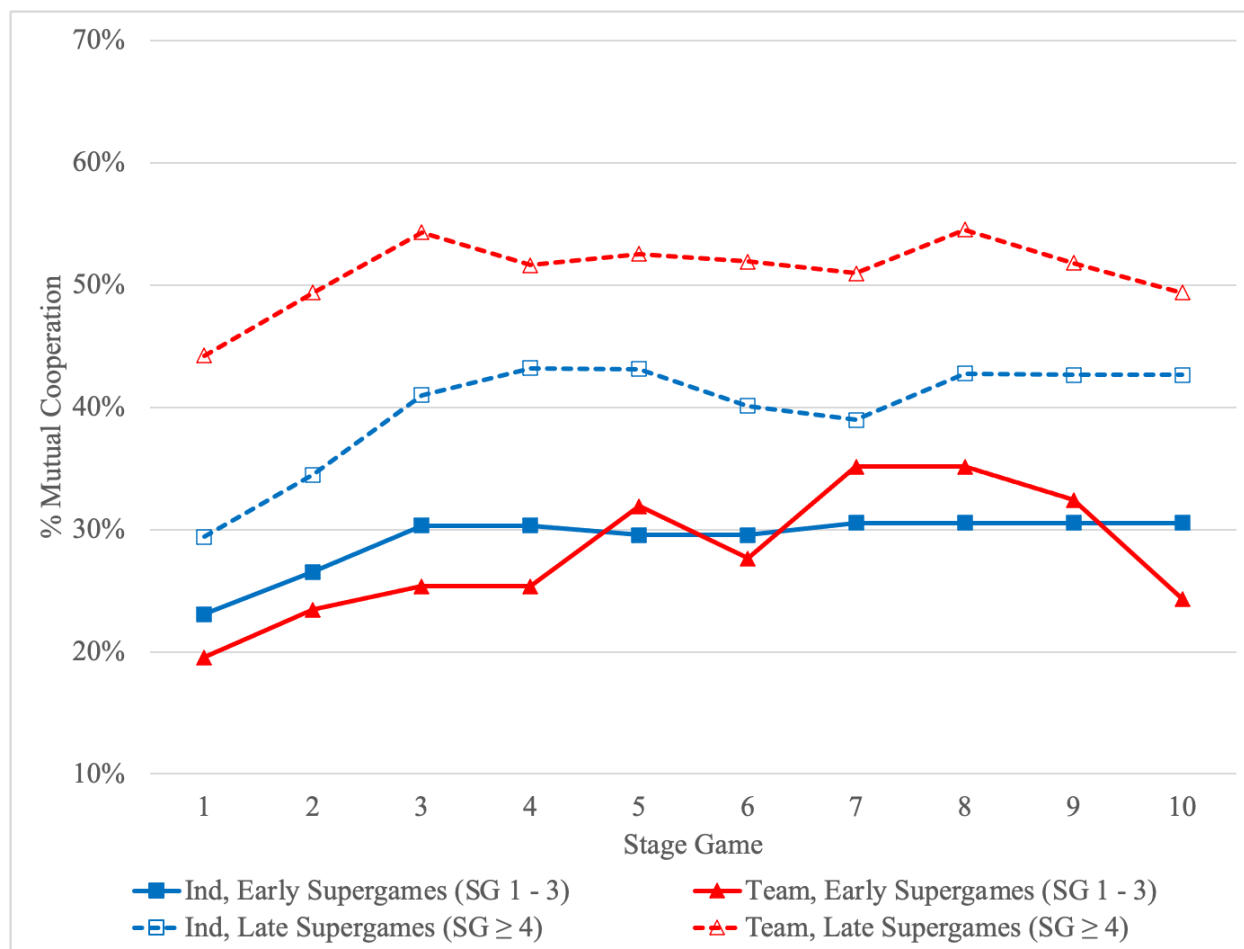
IV. Experimental Results, Individuals vs. Teams: Holding the seed fixed, the number of supergames completed was always less in the team treatment. To maintain parallelism, the dataset is restricted to data prior to and including the *final common supergame*, defined as the last supergame played by *both* individuals and teams using the *same* seed. For example, if teams played nine supergames with a given seed and individuals played twelve with the same seed, in both cases data used is from the first nine supergames.

The notation SG x refers to the x^{th} supergame in an experimental session (i.e. SG1 for the first supergame, SG2 for the second supergame, etc.) and St x refers to the x^{th} stage game within a supergame (i.e. St1 for the first stage game, St2 for the second stage game, etc.). Unless stated otherwise, statements about statistical significance comparing teams and individuals are based on Wilcoxon matched-pairs signed-rank tests, where observations are session averages paired by seed class. Likewise, statements about the statistical significance of changes over time within a session (e.g. does mutual cooperation differ between the first and last supergame) are also based on Wilcoxon matched-pairs signed-rank tests. These are weak tests that are biased in favor of Type II errors. Regressions yielding similar results are reported in online Appendix C.

Throughout this section, an observation is defined as a single play of a stage game. There are three possible outcomes for each observation: mutual cooperation (CC), mutual defection (DD), and mixed (CD). Mutual cooperation is the primary measure of cooperation used throughout the paper. Stressing mutual outcomes rather than individual agents’ choices is largely a matter of convenience; individual cooperation and mutual cooperation are highly correlated, and it is redundant to describe results for both. Analysis reported in Appendix C shows that the main

conclusions are not affected by using mutual cooperation as the measure of cooperation rather than cooperation by individual agents.⁸

Figure 1a: Mutual Cooperation Across Stage Games



Notes: Figure 1a is based on 9,498 observations.

Figure 1a reports the rate of mutual cooperation over the first ten stage games (St1 – St10). The data are broken down by individuals or teams, and early (SGs 1 – 3) or late (SG ≥ 4) supergames. As another way of seeing how the data changes with experience, Figure 1b shows the average mutual cooperation rate in St1 of SGs 1 – 8; after this point in time, more than half of all agents had dropped out of the sample due to sessions ending. For parallel figures based on individual cooperation rates, along with a brief discussion, see Appendix C.

⁸ In 89% of all stage games, agents either mutually cooperated or mutually defected.

Looking at Figure 1a, notice that the frequency of mutual cooperation changed little after the first few stage games.⁹ Differences between treatments were largely driven by what happened in St1. Given this, the following analysis focuses on mutual cooperation *in St1*. This measure has the advantage of not being affected by the differing length of supergames and is highly correlated with mutual cooperation in later stage games ($\rho = 0.69$). See Appendix C for results showing that our conclusions are not affected by using data from all stage games.

Figure 1b: Mutual Cooperation in Stage Game 1 Across Supergames

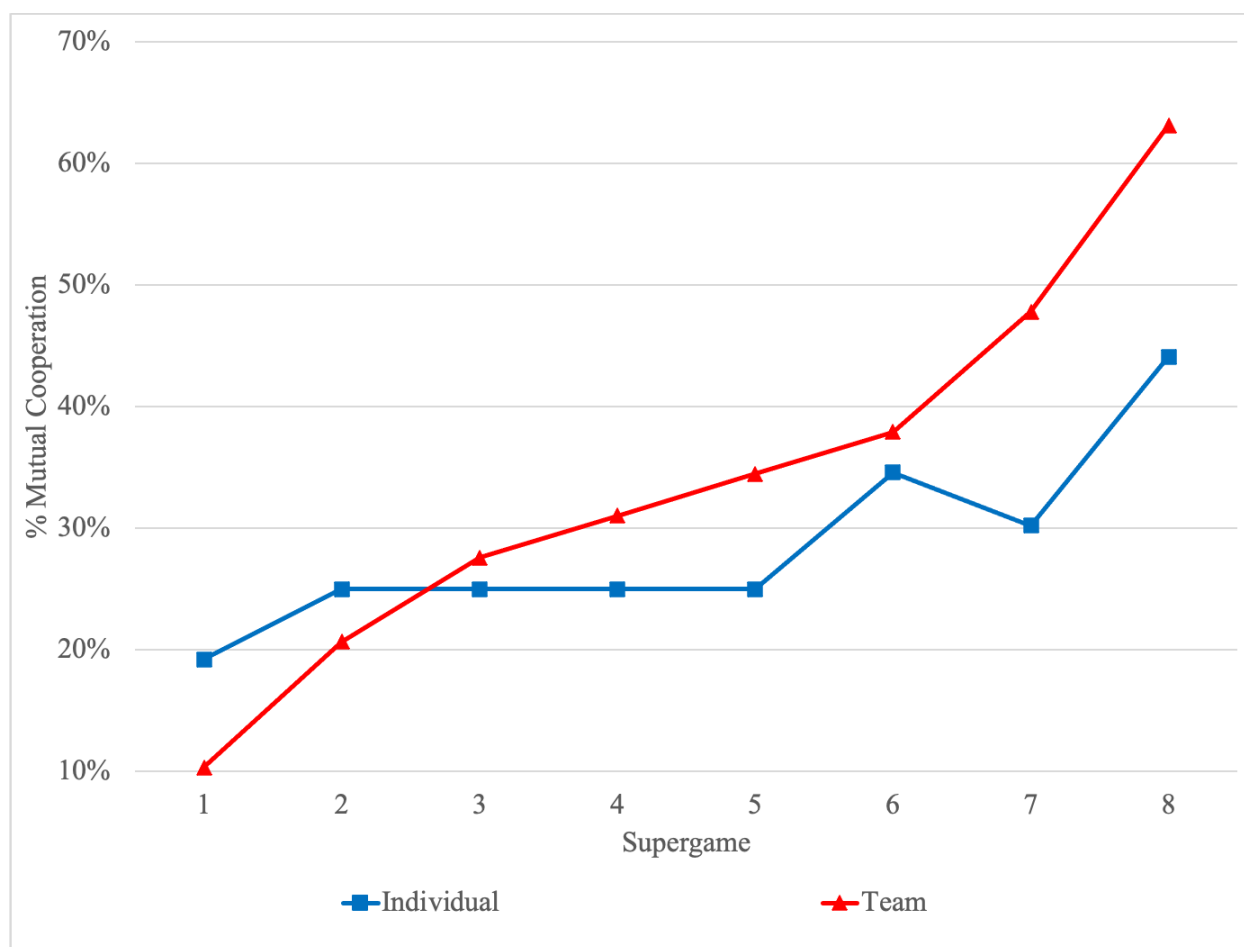


Figure 1b is based on 723 observations.

Mutual cooperation in St1 was *lower* for teams than individuals in SG1 (10.3% vs. 19.2%), but with experience, teams overtook individuals: by the final common supergame, mutual cooperation in St1 was *higher* for teams (55.2% vs. 36.5%). Comparing SG1 and the last common

⁹ The early increase in mutual cooperation is not statistically significant. A probit regression using data from all stage games, including controls for the treatment, supergame, and seed class, the dummy for stage games greater than or equal to 3 has a positive parameter estimate, but is not significant (est. = .035; s.e. = .028; $p = .201$).

supergame, mutual cooperation in St1 increased significantly for teams ($n = 6$; $z = 2.20$; $p = 0.028$) but not for individuals ($n = 6$; $z = 1.05$; $p = 0.292$), with the increase significantly larger for teams compared to individuals ($n = 12$; $z = 1.78$; $p = 0.075$). Agents who cooperated in St1 earned more than those who did not.¹⁰ The faster increase in mutual cooperation on the part of teams is consistent with past findings in other settings (cited in Section I) that teams learn to maximize payoffs more rapidly than individuals.

Observation 1: Mutual cooperation increased faster with experience for teams than for individuals.

Table 2: Number of Switches per Supergame

		Individual	Team
All Observations	Average	1.43	0.85
	# Obs	344	194
Mutual Cooperation (CC)	Average	0.93	0.69
	# Obs	100	70
Mutual Defection (DD)	Average	1.48	0.44
	# Obs	79	52
Mixed (CD)	Average	1.71	1.31
	# Obs	165	72

While *average* levels of cooperation were stable across stage games, this hides a fair amount of switching between outcomes (mutual cooperation, mutual defection, or mixed) within *individual* supergames. A “switch” occurs when the outcome for the current stage game differs from the outcome in the previous stage game within a given supergame.¹¹ Table 2 reports the average number of switches per supergame, excluding very short supergames with only one or two stage games (top row). Results are also reported separately for each possible outcome in St1 (Rows 2 - 4 respectively).

The number of switches was significantly lower for teams than individuals across all observations ($n = 6$; $z = 1.99$; $p = 0.046$). Teams had fewer switches than individuals for all initial outcomes, but the difference was largest following mutual defection in St1. The difference was

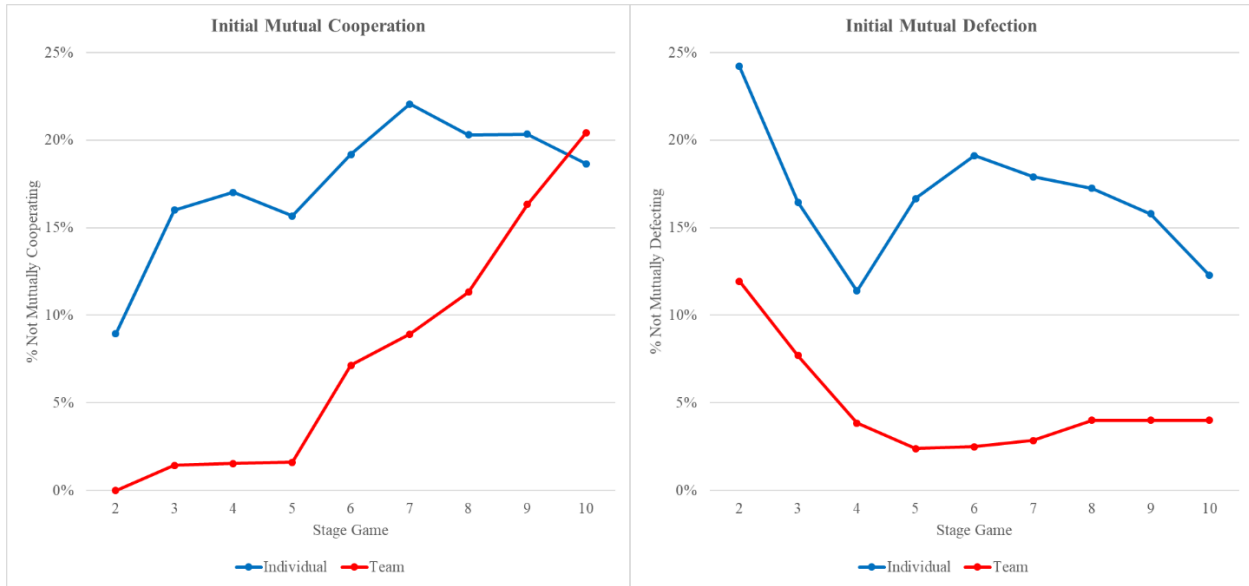
¹⁰ The difference in average payoffs per stage game between those agents who cooperate in St1 versus those who don't was 90 vs. 85 ECUs for individuals and 91 vs. 83 for teams.

¹¹ For example, suppose a pair of agents have the outcomes C/C C/C C/D D/D D/D and C/C in a supergame lasting for six stage games. There are three switches, in St3, St4, and St 6.

slightly smaller in late supergames ($SG \geq 4$), but the average number of switches remained significantly higher for individuals (1.26 vs 0.76; $n = 6$; $z = 1.99$; $p = 0.046$).

The likelihood of a switch was higher for individuals following either mutual cooperation (3.9% vs. 2.5%) or mutual defection (5.0% vs. 1.7%). These small differences in the probability of switching had substantial cumulative effects relative to outcomes in St1. This is illustrated in Figure 2 which shows the fraction of observations where the outcome *differed* from the *initial* outcome in St1, distinguishing between starting with mutual cooperation (left panel) and mutual defection (right panel). For example, consider *pairs* of opponents that mutually cooperated in St1. In the team treatment, 0% switched to a *different* outcome in St2, compared to 8.9% for individuals. For St3, only 1.4% of teams are no longer mutually cooperating, far lower than the 16.0% figure for individuals. Greater stability is a double edged sword; teams are better at sustaining mutual cooperation than individuals, but worse at escaping from mutual defection.

Figure 2: Stability, Likelihood of Not Playing Initial Outcome



Note: The left panel is based on 995 observations, and the right panel includes 1272 observations.

The greater stability of team play applied *between* supergames as well. Classify an agent as having made a “switch between supergames” if their action in St1 of the current supergame differs from St1 in the previous supergame. The proportion of switches between supergames was 23% for individuals versus 14% for teams, which is a weakly significant difference ($n = 6$; $z = 1.78$; $p = 0.075$). Aggregating across common supergames, 69% of teams switched St1 actions between supergames only once, compared with 49% for individuals. As reported below in Table

4, analysis of the team chat shows that teams tend to switch from Always Defect to cooperative strategies, but rarely switch back, leading to greater stability of team play between supergames.

Observation 2: Play was more stable for teams than individuals, both within supergames and between supergames.

The non-parametric tests reported above are useful but conservative and somewhat limited since there are no controls for either the varying length of supergames, or agents' prior experience, which are known to affect cooperation rates. Appendix C reports regressions that control for these issues. The regressions confirm Observations 1 and 2, with statistical significance at the 5% level or better in both cases. Observations 1 and 2 continue to hold with a variety of changes in the specification (e.g. using all stage games rather than St1, using individual cooperation rather than mutual cooperation).

V. Analysis of Team Discussions and Strategies: The main innovation of this paper is the use of teams' discussions to understand the thought processes and motivations underlying teams' "strategies": What strategies teams used, why did they use them, and how (and why) did they change with experience?

For all discussions, quotation marks separate the different team members' messages, with choices "A" and "B" changed to "C" and "D" to aid the reader. Spelling and grammatical errors are not corrected.

5.1 Features of Team Discussions: Several common features of the team discussions dominate the analysis and interpretation of the data: (1) once a strategy was adopted, teams often continued to use it without restating it, although it is clear from their choices that they were following the same strategy, (2) ideas were often developed across multiple stage games, with choices frequently based on discussions in earlier stage games as well as earlier supergames, and (3) teams frequently did not specify a complete set of state contingent actions, improvising instead. The following examples illustrate these three features.

- 1) Continuing to use a strategy without restating it: This team first used a cooperative strategy in SG8. They had started all prior supergames with D but had multiple discussions about possibly using a lenient Grim strategy. Their play from the beginning of SG8 was consistent with Grim with leniency for 3 stage games (Grim3), with a brief reiteration of the logic ("D?" "Let's try C a couple more times to see if they catch on"). For the subsequent four supergames (SG9 –

SG12), play was 100% consistent with Grim3 although they never discussed their strategy again. It was coded as such.

- 2) Using a strategy based on an earlier extended discussions: This team switched to Grim at the beginning of SG2. They had discussed this switch across multiple stage games in SG1 (e.g. “think C until opponent choose D is a good strategy to start each match?” “OK”) but did not talk about their strategy in SG2. The relevant discussions had all taken place earlier.
- 3) The improvisational nature of team decision making: The sequence of actions this team experienced was as follows, with own choice listed first: D/D, D/D, C/D, C/D, D/C, C/D, C/C, followed by continued mutual cooperation. They started out not trusting the other team, choosing D out of fear of getting the “sucker payoff.” In St2, they discussed trying to achieve mutual cooperation and in St3 switched to C:

St2: “I’m wondering if we can coordinate with them and pick C. Then we would both get 105 unless they choose D again, which they probably will... Lets do D again” “yes that’s what I’m thinking...” “One of our teams needs to get smart and always choose C so we all get the most lol” “true! we need to coordinate.”

St3: “what about now...” “If one of us makes the jump to C, there is a chance that the other team will keep doing D. Is this a test of our selfishness? lol” “ok!” “Shall we try C?”

After trying to achieve mutual cooperation in St3 and St4, they gave up in St5.

St5: “Well that other team is selfish lol lets go with D. We tried” “i think no one will choose C... let’s choose D” “Yeah. If we all take a chance and choose C, we will all get more!”

In St5, the other team switched to C just as they returned to D. In St6, they scrambled to recover from their poor timing, choosing C while the other team chose D.

St 6: “gosh dang it lol ... they picked C ... Back to C?” “let’s do it lol.”

St7: “well ... C again?” “why other team just not coordinate with us... D maybe?” “just read our minds, jeez” “lol ... I think we should stick with C one more time.”

At this point, mutual cooperation was finally achieved. There was no over-arching plan behind this team’s choices, so it does not fit neatly into the paradigm of picking a strategy at the beginning of a supergame, in the standard game theoretic sense. Rather they changed their mind about their plan in St2 and improvised after that.

It is worth noting that there is no evidence of *intentional* mixing in the team discussions as economists would define it. This contrasts with findings from Breitmoser (2015) and Romero and Rosokha (2021).

5.2: Coding Team Strategies: To quantify the content of team discussions, two coding exercises were conducted. The first was a game-by-game coding of the dialogues in which the unit of coding was a stage game (specifically, a dialogue started after teammates learned the outcome of the previous stage game). The two authors went separately through the team chats from a randomly selected session and developed categories to be coded. These categories included discussing what action (C or D) to choose, reasons for choosing either an action or strategy, and generic questions about game play (e.g. asking a teammate about the continuation probability).

Coders were also asked to identify which of six simple strategies agents used: Always Defect (AD), Always Cooperate (AC), Grim Trigger (Grim), Tit-for-Tat (TFT), Suspicious Tit-for-Tat (STFT), and Win Stay, Lose Shift (WSLS).¹² Based on a meta study of IRPD experiments, Dal Bó and Fréchette (2018) found that these generally accounted for more than three-quarters of the strategies used. Teams often discussed versions of Grim Trigger with leniency that permitted two or three defections before triggering punishment (Grim2 and Grim3), coded as Lenient Grim.¹³ Complex versions of STFT (labeled “Complex STFT”)¹⁴ and Grim with Counting were added to the strategy set at the suggestion of the coders. Complex STFT refers to teams that started with D, but quickly switched to C in an attempt to achieve cooperation (possibly in response to cooperation by the other team).¹⁵ Grim with Counting describes cases where a team started with Grim, but at some preset point, usually close to St10, unilaterally defected on the grounds that the supergame was likely to end (i.e., these teams suffered from the gambler’s fallacy). Grim with Counting had received little attention in the existing literature but was previously reported in Romero and Rosokha (2018). The dialogue below is an example from a team that unilaterally defected in St9:

St7: “when do you want to go D?” “I say 10 since they can last this long”

¹² STFT is the same as TFT, except it starts with defect rather than cooperate. WSLS starts with cooperate, and then cooperates if and only if both players used the same action in the previous stage game.

¹³ Grim2 and Grim3 were difficult to distinguish, hence we combined them into a single category. We also added more complex versions of TFT (e.g. tit for two tats; two tits for two tats, etc.) to the coding scheme, paralleling the addition of more complex versions of Grim, but no examples were observed.

¹⁴ STFT and Complex STFT are not identical. For instance, the following sequence: D/C C/D C/C is not consistent with STFT which would have had D/C in St3.

¹⁵ Complex STFT often led to mutual cooperation. This might require multiple plays of C/D before abandoning C or settling down to mutual cooperation. For example, this team’s dialogue prior to St2 stated the underlying strategy: “Should we switch to C ... I feel like yes” “C?” “they’ll go C if we do a couple times” “ok” “why not haha ... worth a shot ... maybe end up with 5 once but if we get 5 twice we will switch back to D.” A second example of Complex STFT did not involve responding to initial cooperation by the other team. Play began D/D, and then one team unilaterally switched to C. Prior to St1, this team briefly stated their strategy: “D first then ride C” “I’d say so.” They reiterate this strategy again before flipping to C in St2: “i think we might have to stick to D this round” “C for 2 times. If they don’t switch, then we go back to D” “ok.” The two teams achieved mutual cooperation in St3 and subsequently.

St8: “10?” “9” “cool with me case thats the average amount”

After defecting in St9 ...

St10: “okay we gotta stick with D now lets hope it ends soon” “true”

Two graduate student RAs coded the dialogues independently. The coders were provided with copies of the coding categories along with explanations of each category.¹⁶ The RAs were not told about any hypotheses the authors had, and conversations between the RAs and the co-authors were limited to clarifications of the coding scheme, rather than suggestions about how any specific discussion should be coded. The coders were instructed that their task was to quantify the content of messages rather than interpreting the messages. The coders were free to code multiple categories for a single stage game. The coders were encouraged to add categories if they felt the researchers had missed something; complex STFT and Grim with Counting were both suggested by the coders. Appendix D reports the full set of coding categories, frequencies for each category, and Cohen’s kappa for agreement rates between the two coders, which are generally quite high.

A second coding exercise was conducted to identify strategies from the team dialogues at the *supergame* level. As the name implies, the unit of observation for this exercise was an entire supergame. In extending the game-by-game coding to the supergame level, the simplest case is when a team explicitly stated the strategy underlying their choices. These coders were told to “use the stage game coding as a guide” in determining what strategy to code but were not to be bound by it.¹⁷ If they felt the stage game coding had misidentified the strategy, they should code what they felt was the correct strategy. The instructions to coders stressed that once a strategy was identified, “They don’t have to keep saying [their strategy] *as long as their choices correspond to the strategy they had been using.*” Coding for a team only changed if they explicitly stated a new strategy, or their actions deviated from their previously stated strategy. When a new strategy was *not* explicitly stated at the time of the change, the coders were instructed to take a holistic approach to determine what strategy was being used. This included looking at the discussions surrounding the point where a change took place along with the team’s choices. As noted previously, strategies

¹⁶ For instance, Table A1 lists “Myopia” as a category. The coders’ instructions added the following explanation: Myopia has a number of possible characteristics all of which should lead to classification under myopia. There is no need to distinguish between the characteristics. (a) Focused on getting 175 in current round with no consideration of longer run implications/impact on other teams’ choices in subsequent rounds. (b) Focusing on total earnings in terms of getting 75 each round. No consideration of tradeoffs from choosing C vs D (considering the future). (c) Short sighted – when deciding to defect no consideration of tradeoffs/ possible negative consequences of choosing D.

¹⁷ The coding instructions stated, “You will find that occasionally the previous coders got it wrong. So in general anchor off their coding but if it’s obviously off, feel free to change the coding.”

were often developed over a number of stage games and the clearest statement of a team's strategy was often prior to the time the team started using it. In short, coders used a combination of all the available evidence along with their best judgment to determine the strategy employed.

If a team's strategy changed midway through a long supergame, the supergame was coded based on the *initial* strategy for that supergame. For example, if a team started a supergame with AD but decided in St15 to start playing Lenient Grim, the supergame would be coded AD. If no further discussion of strategies occurred and play in subsequent supergames was consistent with Lenient Grim, these supergames were coded as Lenient Grim.

This extension of the coding to the supergame level was done by two economic graduate student RAs. The coders were given detailed examples of each strategy as well as instructions on how to conduct the coding. They coded the strategies independently and were asked to meet (without the researchers present) to reconcile any differences in their coding. Agreement between the coders was high before they met ($\kappa = 0.79$).¹⁸ The combination of what teams said and did usually made it clear what strategy a team was using.¹⁹ The coders were given the option of leaving a supergame uncoded if they could not identify, or agree, on the strategy, which happened occasionally even after reconciliation.²⁰

When choosing the list of strategies to be coded, we drew heavily on the relevant literature in game theory and experimental economics. A reviewer suggested that this may have unwittingly biased the results. To check this, we employed two teams of undergraduate RAs to develop their own lists of strategies with minimal direction. Their lists were short and included strategies comparable to AD and Grim. They tended to ignore subtleties that are important in understanding behavior; for example, both teams bundled Grim, Grim with Leniency and Grim with Counting into a single strategy. Details are reported in Appendix E.

5.3 Teams' Strategies: Table 3 shows the distribution of strategies identified from the team chats at the supergame level. The number of observations in Table 3 declines in later supergames due to

¹⁸ The most common disagreements were cases where one coder coded a supergame while the other left it uncoded, and where the coders agreed that some version of Grim was being played but disagreed on which version.

¹⁹ For example, the combination of what teams said and did made AD easy to recognize. In all 183 cases where a team was coded for AD, the team chose D in the first stage game.

²⁰ There were cases where the two coders did not successfully reconcile their coding. In all such cases, one of the two coders picked a strategy while the other left the supergame uncoded. The analysis that follows uses the strategy from the coder who coded that supergame.

shorter sessions ending. Frequencies are not reported for Always Cooperate or Win-Stay-Lose-Shift as these strategies were never observed.

Table 3: Strategy Frequencies from Team Chats

SG	Always Defect	Grim Trigger	Lenient Grim	Grim w/ Counting	TFT	STFT	Complex STFT	Uncoded	#Obs
1	55.17%	22.41%	3.45%	3.45%	6.90%	1.72%	6.90%	0.00%	58
2	53.45%	12.07%	8.62%	8.62%	8.62%	3.45%	3.45%	1.72%	58
3	48.28%	22.41%	10.34%	1.72%	3.45%	5.17%	6.90%	1.72%	58
4	43.10%	24.14%	20.69%	3.45%	1.72%	1.72%	3.45%	1.72%	58
5	36.21%	29.31%	12.07%	8.62%	5.17%	5.17%	3.45%	0.00%	58
6	34.48%	25.86%	18.97%	5.17%	5.17%	6.90%	3.45%	0.00%	58
7	17.39%	36.96%	17.39%	4.35%	8.70%	4.35%	6.52%	4.35%	46
8	7.89%	39.47%	23.68%	7.89%	10.53%	5.26%	5.26%	0.00%	38
9	7.89%	42.11%	18.42%	13.16%	7.89%	5.26%	2.63%	2.63%	38
10	16.67%	30.00%	20.00%	6.67%	6.67%	13.33%	3.33%	3.33%	30
11	20.00%	25.00%	20.00%	5.00%	5.00%	20.00%	5.00%	0.00%	20
12	15.00%	25.00%	25.00%	5.00%	10.00%	10.00%	10.00%	0.00%	20
Total	33.89%	27.04%	15.19%	5.93%	6.30%	5.56%	4.81%	1.30%	540

Initially, Always Defect (AD) was easily the modal choice, averaging 52% across the first three supergames. Variants of Grim Trigger (Grim Trigger, Lenient Grim, and Grim with Counting) were also common, totaling 31% of the observations in SGs 1 – 3. The weight on AD decreased continuously across supergames, shifting primarily to variants of Grim Trigger. By SGs 5 – 7, the last point before substantial dropouts due to sessions ending, the frequency of AD was 32%, while variants of Grim Trigger combined for 51% of observations.²¹ Although never common, Grim with Counting was always present. Variants of tit-for-tat (TFT, STFT, and Complex STFT) were less common than variants of Grim Trigger and did not change frequency with experience (16% in SGs 1 – 3 and SGs 5 – 7). The movement away from AD towards cooperative strategies parallels the increased cooperation rates in the choice data.

Because the coding in Table 3 was done at the team level, rather than the population level, it is possible to identify how strategies for individual teams changed with experience. We classify strategies into two broad categories: AD versus potentially cooperative strategies (variants of Grim

²¹ For the one session with only six supergames, data from SGs 4 – 6 was used.

Trigger, TFT, or STFT).²² Table 4 reports how frequently individual teams switched between these two categories across all supergames, broken down by which category the team was in for the first supergame (SG1).

Table 4: Number of Switches Between Categories of Strategies for Individual Teams

Switches	SG1: Always Defect	SG1: Cooperative	Row Total
0	10 31.3%	16 61.5%	26 44.8%
1	16 50.0%	2 7.7%	18 31.0%
2	2 6.3%	6 23.1%	8 13.8%
3	4 12.5%	1 3.9%	5 8.6%
4	0 0.0%	1 3.9%	1 1.7%

Ten teams played AD for all supergames, and sixteen always played one of the cooperative strategies (see the 0-switch row of Table 4). Half of the teams who started out playing AD (50%) switched to a cooperative strategy and never switched back (the 1-switch row). Collectively, a clear majority of teams (76%) had zero or one switch. Movement was generally away from AD towards cooperative strategies; 63% of the teams (20/32) coded for AD in SG1 used a cooperative strategy in their final supergame, but only 12% of the teams (3/26) that started with a cooperative strategy in SG1 chose AD in their final supergame. Changes in strategy were largely a one-way street from AD to potentially cooperative strategies.

Observation 3: At the supergame level, AD was initially the most frequently used strategy for teams. With experience, cooperative strategies, primarily variants of Grim Trigger, gained weight at the expense of AD. There wasn't much back and forth between AD and cooperative strategies. For the most part, teams either picked AD or one of the potentially cooperative strategies to begin with and stuck with it, or there was one-way movement from AD to a cooperative strategy.

Table 5 compares strategies identified from team dialogues with the distribution of strategies estimated by the Strategy Frequency Estimation Method (SFEM, Dal Bó and Fréchette, 2011). SFEM is, by far, the most commonly used technique for inferring the strategies underlying

²² In the few cases with missing codes, the coding from the previous supergame was employed. There were no cases with more than one consecutive uncoded supergame.

choices in IRPD games.²³ SFEM models individuals as playing finite automata, capturing common strategies such as Grim Trigger or Tit-for-Tat. Using a pre-specified set of strategies, the model calculates the likelihood of each agent's observed actions subject to some probability distribution over strategies. The weights on strategies are then fit by maximum likelihood estimation. SFEM is a mixture model; it estimates the probability distribution of strategies across the entire population rather than assigning specific strategies to specific agents. This implies that it cannot identify when individual agents have changed strategies, making an exercise like Table 4 impossible. To ease identification, SFEM is generally estimated on a block of supergames rather than a single supergame. The estimates in Table 5 follow this approach, estimating the model separately for the early (SGs 1 – 3) and late (SGs 5 – 7) supergames.²⁴ Appendix F provides an extended discussion of how we fit SFEM.

Table 5: Strategy Frequencies: Chat Coding vs. SFEM

Time Period	Method	Always Defect	Grim Trigger	Lenient Grim	Counting	TFT	Complex TFT/STFT	STFT
SGs 1 – 3	Coding	52.3%	19.0%	7.5%	4.6%	6.3%	3.4%	5.7%
	SFEM	43.7%	2.7%	0.0%	0.0%	28.7%	5.0%	19.9%
SGs 5 – 7	Coding	31.6%	28.7%	16.7%	5.7%	5.7%	5.7%	4.6%
	SFEM	19.4%	19.2%	19.0%	0.0%	27.3%	0.0%	15.0%

Note: Lenient Grim combines Grim2 and Grim3. Complex TFT combines TF2T, TF3T, and 2TF2T from SFEM and Complex STFT from the coding. For the session with only six supergames, data from SG4 – SG6 are used. Data is included from 348 supergames (coding) and 4,278 choices (SFEM)

Both approaches, SFEM and coding from the chats, capture the broad movement away from AD to more cooperative strategies, with neither identifying much use of more complex versions of TFT or STFT. That said, there are substantial differences between the distribution of strategies identified by the two approaches. Categorizing contingent strategies as variants of Grim Trigger (Grim Trigger, Lenient Grim, and Grim with Counting) or variants of TFT (TFT, STFT, and Complex TFT/STFT), coding based on the chats always puts less weight on variants of TFT and more on variants of Grim Trigger than SFEM. This is especially pronounced in early supergames (SG 1 – 3), where the coding from team chats assigns 16% of the population to variants of TFT versus 31% to variants of Grim Trigger, compared to 54% and 3% for SFEM. The

²³ See Dal Bó and Fréchette (2018) for a summary of existing papers that use SFEM.

²⁴ Data from SGs 4 – 6 were used for the one team session with only six supergames.

difference between methods remains large in late supergames (SG 5 – 7): The coding assigns 16% and 51% to variants of TFT and Grim Trigger respectively, as opposed to 42% and 38% for SFEM.

To understand the broad differences between SFEM and the coding, consider the types of histories that make it possible to distinguish between TFT and Grim Trigger. A sequence like the following must be observed to identify TFT: C/C C/D D/D D/C. There needs to be initial cooperation that falls apart, followed by an attempt to re-establish cooperation. This rather complex set of events occurred infrequently, giving teams little reason to discuss this option before it happened. Teams generally used incomplete strategies, improvising when faced with an unanticipated outcome. SFEM identifies teams' strategies solely based on what actions they take, with no means of distinguishing a premeditated choice from improvisation.

Observation 4: Compared to SFEM, the coding from team chats identified a relatively higher proportion of variants of Grim Trigger and relatively fewer variants of TFT.

Table 6: SFEM Estimates, Individuals vs. Teams

	<i>Individual</i>		<i>Team</i>	
	SGs 1 – 3	SGs 5 – 7	SGs 1 – 3	SGs 5 - 7
AD	29.78%	15.80%	43.67%	19.38%
AC	1.11%	1.11%	0.00%	0.00%
Grim	12.16%	8.53%	2.69%	19.20%
Lenient Grim	4.88%	6.32%	0.00%	18.99%
TFT	23.32%	25.02%	28.74%	27.43%
STFT	20.66%	18.41%	19.89%	14.99%
Complex TFT	7.05%	24.81%	5.01%	0.00%
WSLS	1.05%	0.00%	0.00%	0.00%
# Observations	3,624	4,178	1,938	2,340

Comparing strategies between individuals and teams must rely on SFEM rather than the coding. Table 6 compares the distributions of strategies estimated for the individual and team data, subdivided between early (SGs 1 – 3) and late (SGs 5 – 7) supergames. To simplify the table, we have combined some strategies: “lenient grim” includes both Grim2 and Grim3, and “complex TFT” includes TF2T, TF3T, and 2TF2T. Table B1 in Appendix F includes estimates for the component strategies in these categories, as well as standard errors and the noise parameter.

The estimated distribution of strategies changes over time for both individuals ($\chi^2 = 392.96$; d.f. = 12; $p < .001$) and teams ($\chi^2 = 331.31$; d.f. = 12; $p < .001$).²⁵ In both cases there was a shift away from AD towards more cooperative strategies, but which cooperative strategies gained weight differed: Variants of TFT (TFT, STFT, and Complex TFT) gained 17 percentage points for individuals but lost 11 percentage points for teams, while versions of Grim lost 2 percentage points for individuals and gained 36 points for teams.

Focusing on the late supergames (SGs 5 – 7), when agents have had a chance to learn and the strategies employed have settled down, the estimated parameters are significantly different between individuals and teams ($\chi^2 = 517.44$; d.f. = 12; $p < .001$). The primary difference is that teams are estimated to put much more weight on variants of Grim Trigger (38% vs. 15%) and much less on variants of Tit-for-Tat (42% vs. 68%). Almost a quarter of individuals are identified as using complex versions of TFT (Complex TFT), while this is assigned no weight for teams.

Given the coding results, it seems likely that SFEM overestimates the frequency of variants of TFT for individuals just as these are overestimated for teams. Subject to this caveat, what might the differing estimates between individuals and teams indicate? Play by individuals was inherently less stable than team play. This shows up in the higher frequency of switching reported in Table 2 and is paralleled by higher estimated error rates for individuals in SFEM. The team dialogues indicate that responses to unexpected changes in the opposing players strategy were not specified in advance, but instead were improvised. Individuals faced more unexpected changes than teams because individual data was noisier, and hence gave agents more opportunities to improvise. We suspect that SFEM puts more weight on variants of TFT for individuals compared to teams because of more frequent improvisation (due to greater instability) by individuals relative to teams.

Observation 5: Based on SFEM, individuals use variants of TFT more and variants of Grim less than teams.

5.4 Coding the Rationale for Switching to Cooperation: Given that switching to cooperation generally involved extended discussions between teammates, the game-by-game coding isn't terribly useful for understanding *why* teams switched from AD to a potentially cooperative strategy. A separate coding at the *team* level was developed to examine teams' rationale for *unilaterally* switching to cooperation. Two cases were considered: (1) switching from an initial

²⁵ Test statistics are log-likelihood ratio tests.

choice of D to an initial choice of C *between* supergames and (2) switching to C following mutual defection *within* the same supergame. There were 41 switches to cooperation between supergames and 25 within supergames, with 39 of the 58 teams having at least one switch.

The authors read through a sample of dialogues when teams switched to cooperation and identified common rationales for the change.²⁶ The point was not to identify what strategy was used, but rather *why* a switch took place. Two graduate student RAs then coded all of the dialogues where a switch to cooperation occurred.²⁷ The instructions the RAs received stressed the need to look not just at the stage game when a change took place, but also the surrounding dialogues, given the extended nature of team discussions. The coders each went through the dialogues separately, and then met to reconcile their codings. The discussion below is based on their reconciled codings. Appendix D (Table D2) reports descriptions and frequencies of these coding categories, along with Cohen's kappa for agreement rates between the two coders prior to reconciliation. For the most part the coders agreed reasonably well on the coding.

5.5 Understanding the Rationale for Changes in Teams' Choices: Define a "substantive" discussion as a dialogue for a single stage game with teams discussing what choices to make either in the current stage game or the future (15% of stage games). This eliminates dialogues that were unrelated to the experiment. By far the most common topic was what action to take for the current stage game (84% of substantive discussions). This occurred in 13% of the stage games as there was little to discuss given that mutual cooperation and mutual defection were quite stable and the interface made it easy to coordinate choices. When there was a genuine need to discuss what action to take, teammates typically did so. In their initial interaction (SG1, St1), 97% of teams discussed what action to take. Discussions of what action to use were also common when changing initial actions from the previous supergame (81%), changing actions within a supergame (71%), or responding to a change in their opponent's action (47%).

Discussions of strategies (as opposed to actions) were rarer, occurring in only 21% of substantive discussions. Once teams adopted a strategy, they felt little need to continue discussing their strategy if they weren't changing it. Discussions of strategies were surprisingly rare at key moments. A clear majority of teams discussed strategies before the first stage game of the first

²⁶ The sample was drawn from the first two team sessions. We used two sessions rather than one to get a reasonably large sample for generating categories.

²⁷ One of these RAs also did the supergame level coding; due to availability, the other coder was new.

supergame (68%), but relatively few did when switching initial actions at the *start* of a new supergame (33%), or when making a unilateral deviation from mutual cooperation or mutual defection *within* a supergame (29%). This speaks to the point made previously, that discussions about changing strategies were often extended affairs. When implementing a change in strategy, teammates frequently relied on earlier discussions rather than a discussion at the point in time when the change actually occurred. It follows that we need to use a team’s extended conversations when coding their strategies rather than taking a more granular approach.

Teams’ discussions provide insight into the motivation behind their initial choices. In SG1, 59% of teams were coded as using AD. The most common reason given for this in St1 (43%) was not trusting the other team (e.g. “D. You know they’ll pick D, so let’s get 75 instead of 5.”). This was twice as frequent as either of the next two most common reasons for choosing AD – myopia (focusing on the short run benefits of D without recognizing any possible longer-term repercussions) at 19% and discussing the impact of choosing D on future play at 21%.²⁸ For teams that chose cooperative strategies (variants of Grim or TFT) in SG1, the most common reason given in St1 was mutual gains from cooperation (24%). The relationship between what teams discussed and choices in St1 was quite strong as can be seen in the probit regression below; the dependent variable is a dummy for choosing C in St1 of SG1 (with p-values in parentheses).²⁹

$$C = -0.753 \text{ Distrust} - 0.420 \text{ Myopia} + 0.816 \text{ Mutual Benefits} - 0.029 \text{ Discuss Future} + 0.135 \text{ Confusion} \\ (0.006) \quad (0.006) \quad (0.002) \quad (0.778) \quad (0.331)$$

Observation 6: Distrust of the other team, resulting in fear of getting the “sucker” payoff, was the most common reason coded for justifying an initial choice of AD.

Switching to a cooperative strategy first involved realizing that the team could do better through mutual cooperation. While this may seem obvious, it was an “aha” moment for many teams, who then had to figure out *how* to coordinate on mutual cooperation.³⁰ The most common approach was an attempt to lead by example and/or to signal an intent to cooperate (“Leading”). Leading was coded for 53% of teams when unilaterally switching to cooperation, being especially

²⁸ It may seem strange that teams choose D while discussing the future negative repercussions of this choice. A team could recognize the negative effect of initially choosing D on future cooperation and still feel that the benefits of protecting themselves against the sucker payoff justified choosing D.

²⁹ P-values are based on robust standard errors. Confusion captures teams that discussed the rules of the game, largely because one of the teammates was confused.

³⁰ Prior to switching, 34% of teams were coded for explicitly discussing the benefits of mutual cooperation and 79% of teams were coded as explicitly discussing the need to adopt a new strategy.

common (76%) when teams switched to cooperation within a supergame. By cooperating, they hoped their opponent would view them as willing to cooperate and be willing to follow their example. The dialogues are full of examples like the following: “i wonder if choosing C once will make them willing to switch” “I guess we could try?” “it might be worth it” “here we go lol.” Teams often viewed leading by example as a way of sending a message in an environment where direct communication was not possible, as one team noted: “this is all so hard without communication” “I know if we could just send them like one sentence we'd have it made.”³¹

Leading is a good strategy (in the non-technical sense) for IRPD games with relatively high continuation probabilities. Because there are, in expectation, many stage games within a supergame, it costs little to take the sucker payoff for a few stage games, compared to the high potential payoff from mutual cooperation in later stage games. Indeed, some teams explicitly discussed the tradeoff between the losses from switching to C and potential gains from achieving mutual cooperation:

“What if we start with C and do it the whole time” “then we would get 5 points when they choose D” “i think they will start with D, but they are good people so will switch to C” “I don’t think they will” “we would make up the first [stage game] loss quickly” “alright let’s do C this round but if they choose D i think that’s where they’ll stay.”

Leading was often associated with the lenient grim strategy, as teams understood they needed to give their opponent a chance to catch on. For example, the team just quoted had their opponent choose D in St1 after which the teammate who wanted to cooperate said, “we have to do it twice to see if they change.” Their teammate reluctantly agreed, and they achieved mutual cooperation in the next stage game.

A second common rationale for trying cooperation (“learning”) was to determine whether the opposing team was willing to cooperate, rather than trying to influence them per se. This was coded for 24% of teams unilaterally switching to cooperation, less than half as frequent as leading. The following brief exchange illustrates learning:

³¹ There is a related literature on leading by example in public goods literature (for a summary see Cooper and Hamman, 2021). The mechanism that makes leading by example successful in public goods game is reciprocity; leaders contribute to the public good, anticipating that the other group members are conditional cooperators and will reciprocate by contributing themselves even though this is *not* an equilibrium strategy (Gächter, Nosenzo, Renner, and Sefton, 2012; Arbak and Villeval, 2013). This differs from teams’ discussions of leading by example in IRPD games, which were typically framed in terms of sending a message to coordinate on mutual cooperation. Teams were solving a problem of equilibrium selection rather than relying on the kindness of others.

“we can risk it to see what kinda team they are and press C or go safe and press D again”
“press C” “kk.”

Leading and learning were similar rationales for switching to cooperation, primarily distinguished by whether the intent was to influence the other team (leading) or to determine whether the other team were willing to cooperate (learning). It was uncommon to combine the two with only 8% of teams coded for both leading and learning.

Only 8% of teams explicitly discussed prior play as a reason for switching to cooperation. This is surprising given that teams were more likely to start cooperating if they had recently experienced an opponent who initially cooperated, or following longer supergames with long stretches of mutual defection.³² The history of prior play presumably prompted teams to think about switching strategies but was not explicitly discussed before switching.

Observation 7: The most common approach teams took when trying to coordinate on mutual cooperation was leading by example, signaling their willingness to cooperate.

Finally, the team dialogues provide clues as to why team choices were so stable compared to individuals' choices. Consider cases where the status quo was either mutual cooperation or mutual defection in the previous stage game. If a team only discussed switching from the status quo, they went through with the switch 54% of the time. But when they discussed switching to cooperation *and* the status quo of defection, they switched only 13% of the time. Inertia favored the status quo, consistent with “pluralistic ignorance” noted in the psychology literature (Prentice and Miller, 1996). This holds that even when a member of a group privately rejects an opinion or practice, they tend to believe that other members of the group accept it, making it much easier to abide by an established convention than to change it.

VI. Finite versus Infinitely Repeated Prisoner's Dilemma Games: This subsection compares the data here to results from Kagel and McGee (KM; 2016) who ran a series of finitely repeated prisoner's dilemma (FRPD) games using the same stage game payoffs, the same subject population, and the same procedures as those employed here. The number of stage games, ten, matched the expected number of stage games in the IRPD games.³³

³² Two-thirds of teams that switched to cooperating in St1 experienced an opponent cooperating in St1 of the previous supergame, compared to only one-third two supergames ago. The average length of the supergame preceding the switch was 17.5 stage games, as opposed to 10.8 for two supergames before the switch.

³³ The dataset included five individual sessions and five team sessions. All sessions had at least seven supergames.

The results for the FRPD games parallel those reported for IRPD games as mutual cooperation was initially lower for teams than individuals (12.0% vs 46.2% in the first stage game of SG1), but became greater with experience (52.0% vs. 38.5% in the final common supergame). As with the IRPD games, mutual cooperation increased faster with experience for teams than individuals ($n = 10$; $z = 1.786$; $p = 0.074$).

Although one might expect less mutual cooperation in St1 for FRPD games than IRPD games, mutual cooperation was *more* common in St1 of SG1 for the FRPD games for individuals (46.2% vs. 19.2%) and teams (12.0% vs. 10.3%). This difference was initially significant for individuals ($n = 11$; $z = 2.64$, $p < 0.01$), but not significant by the final common supergame ($n = 11$; $z = 0.367$; $p = 0.714$).³⁴ For teams, the difference in St1 cooperation rates between FRPD and IRPD games was not initially significant ($n = 11$; $z = 0.299$; $p = 0.765$). Although the gap grew with experience, it never became statistically significant.³⁵

For individuals, mutual cooperation in St1 decreased from 41.0% to 31.7% between early (SGs 1 – 3) and late FRPD supergames but increased from 23.1% to 29.4% for IRPD games. Dal Bó (2005) reports similar reductions in St1 cooperation rates for individuals in FRPD games, albeit in much shorter games, commenting that “the effect of the shadow of the future increases with experience.” However, this did *not* carry over to teams, as mutual cooperation in St1 increased substantially with experience for *both* FRPD games (from 26.7% to 51.0%) and IRPD games (from 19.5% to 43.3%).³⁶ Teams rapidly learned to cooperate in St1 for FRPD games even though the incentives for starting with C were worse for teams than individuals in early supergames; cooperating in St1 increased expected supergame payoffs by 92 ECUs for individuals vs. 22 ECUs for teams.

Like the IRPD games, play was more stable for teams than individuals in the FRPD games. The difference in the number of switches looks small in the raw data (1.43 for individuals vs 1.36 for teams), but this is an artifact of differences in the initial conditions. Regressions reported in Table C2 that control for differences in initial conditions between individuals and teams find

³⁴ To compare apples with apples, these comparisons are based on the first seven supergames in all cases. Dal Bó (2005) reports lower St1 cooperation rates with experience in FRPD games (for individuals) than in parallel IRPD games, but this likely reflects differences in the number of stage games: ten in KM versus two or four in Dal Bó. Embry et al. (2018) report that St1 cooperation rates (for individuals) are increasing in the length of FRPD games.

³⁵ The difference between FRPD and IRPD games was at its maximum in SG6 (56% vs. 38%), but was still not statistically significant ($n = 11$; $z = 0.739$; $p = 0.460$).

³⁶ This increase was significant for teams in both FRPD ($n = 5$; $z = 1.761$; $p = 0.078$) and IRPD ($n = 6$; $z = 2.201$; $p = 0.028$) games.

significantly fewer switches with teams (est. = -0.165; s.e. = 0.085; $p = 0.051$). In short, the main differences between teams and individuals in the IRPD games reported here were also present in the FRPD games reported in KM, suggesting that the pattern of results is relatively broad.

Team chats in FRPD games were coded at the stage game level by two graduate students, using essentially the same procedures described above for IRPD games. Like the results reported above for IRPD games, teams that chose D in St1 of SG1 did so primarily out of safety considerations (91.7%; 22 of 24 teams choosing D). Like the IRPD games, 70.6% (12 out of 17 teams choosing C) did so to elicit mutual cooperation with its higher payoff. Lenient Grim was rarely practiced, likely due to the finite number of stage games. In several cases teams managed to generate cooperation by choosing D in St1, switching to C in St2 if the other team chose C, and continuing with C in St3, hoping to achieve cooperation, similar to the generalized STFT strategy described above.

VII. Silent Partner Treatment: The comparison of teams and individuals raises a question whether behavior differs due to the presence of a teammate per se, possibly due to being responsible for their teammate's earnings, or because of communication and joint decision making between teammates. To distinguish between these two hypotheses, a silent partners treatment was implemented. Like the team treatment, subjects were assigned to fixed two-subject teams at the beginning of the experiment, with payoffs the same for both teammates as a consequence of their choices. One member of the team was randomly chosen for the role of Decider and the other was assigned the role of Silent Partner. These roles were fixed for the duration of the experiment. All decisions for the team were made by the Decider. The Decider knew they were responsible for their Silent Partner's payoffs. The Silent Partner observed the Decider's choices and outcomes, but there was no communication between the two.³⁷

Four silent partner sessions were conducted. The same procedures were employed as in the main sessions, and the same seeds were used as in the final four team sessions. All sessions ran for at least as long as the matched team session, with the data reported on below restricted to common supergames. Redoing the regressions reported in Appendix C with only data from the

³⁷ To make it less obvious which subjects were Deciders, silent partners were encouraged to make choices indicating what "they would have [done] if they were the Decider." These played no role in determining the team's choice and Deciders could *not* observe their silent partner's choices. Most silent partners made choices (88% of all stage games), but the correlation with the Deciders' choices was not especially high ($\rho = 0.47$, subject to making a choice).

four seeds used in the silent partner treatment, the differences between individuals and teams reported in Observations 1 and 2 remain statistically significant.

The silent partner treatment was designed to test how much of the difference between teams and individuals was due to having a partner per se, rather than joint decision making and communication. This might matter because decision makers might become more risk averse if their decision affects the risks borne by a passive second party (as in Bolton, Ockenfels, and Stauf, 2015), making them less willing to initially take the risk of cooperating. Additionally, if Deciders exert greater cognitive effort on choosing a strategy because of other-regarding preferences (their effort now benefits their Silent Partner as well as themselves), they should learn to cooperate faster in the silent partners treatment than in the individual treatment. The greater stability of play by teams grows from the process of joint decision making – when the teammates disagree, there is a strong tendency to stick with the status quo. Merely having a partner should not have the same effect.

Limiting the dataset to the four common seeds, initial mutual cooperation rates and the changes in cooperation rates over time differed little between individuals and the silent partner treatment. In the first stage game of SG1, the mutual cooperation rate for the silent partners treatment was 25%, compared with 26% for individuals and 16% for teams. Turning to the final common supergame, mutual cooperation in the first stage game was 56% for the silent partners treatment, an increase of 31 percentage points. The analogous mutual cooperation rates were 44% and 68% for individuals and teams, representing increases of 18 and 53 percentage points. The direction of these effects for the silent partner treatment relative to individuals were the same as in the team treatment, but the magnitudes were far smaller.

Table C1 reports probit regressions estimating the size of these effects, controlling for the seed, supergame, length of the previous supergame, and experience with cooperation by opponents. None of the effects of the silent partners treatment, relative to individuals, were statistically significant.

The average number of switches per supergame was 1.89 for the silent partners treatment. This is higher than the average number of switches in the individual treatment (1.29), rather than lower as in the team treatment (0.87). That said, the estimated difference between the silent partners treatment and the individual treatment is not significant (see Table C2). Once again, play in the silent partner treatment differed little from play by individuals.

The similarities between the silent partners treatment and the individual treatment are consistent with the idea that the differences between individual and team decision making were primarily due to joint decision making and the associated communication between teammates. Having a partner per se could potentially have an effect through either risk or social preferences, but this was not the case.

The effects of joint decision making and team communication are inherently difficult to disentangle since it is not possible to coordinate choices without some type of communication between teammates. To the extent that this has been studied, full bilateral communication is a necessary condition for the strong performance of teams relative to individuals. Neither unilateral communication nor communication limited to an exchange of proposed actions with no further explanation is sufficient to replicate the effect of team play with free form communication (Cooper and Kagel, 2016; Arad, Grubiak, and Penczynski, 2021).

VIII. Discussion and Conclusions: There were two motivations for the experiment reported here: (1) To compare the behavior of individuals and freely interacting two-person teams in IRPD games with perfect monitoring. (2) To use the dialogues between teammates to understand teams' underlying behavioral processes. Arguably, the greatest value of studying teams is use of their conversations to understand how and why decisions were made. Economists have become increasingly interested in process data (e.g. fMRI, eye tracking, reaction times) to understand decision making, including analysis of team chat. Team discussions are a natural part of the decision making process and provide direct insights into how and why decisions come about.

Teams were less cooperative than individuals in early supergames, but cooperation rates increased more rapidly for teams, resulting in significantly more cooperation than individuals in late supergames. Additionally, team play was more stable both *within* supergames and between supergames. This implies that both mutual cooperation and mutual defection were more persistent within supergames for teams.

Analysis of teams' dialogues provided direct insights into the strategies teams used and the rationale underlying their choices. Discussions that explicitly laid out strategies in the game theoretic sense were rare. Rather than specifying plans for all possible histories, strategies typically developed over time and were improvisational in nature. Further, it was not uncommon for teams to switch strategies midway through a supergame. Teams' choices were not arbitrary, since the dialogues showed that they generally thought carefully about how to play the game, but their

thought processes were far less structured than modelers typically assume. Given the broad similarity of team and individual behavior, it seems reasonable to assume that similar processes underlie individual choices.³⁸

Related to the latter point, results from Romero and Rosokha (2019) are consistent with individuals improvising within supergames. Their experiment used individuals as agents, long IRPD games ($\delta = .98$), and direct elicitation of strategies. There were frequent changes within supergames (2.37 per supergame) when subjects could costlessly change their strategy within a supergame, consistent with the improvisation observed in the team chats.³⁹

Strategies identified from the team dialogues show that teams moved steadily across supergames from always defecting (AD) to cooperative strategies with little backsliding. Comparing strategies identified from the team dialogues with the distribution of strategies estimated by SFEM, coding assigns more weight to variants of Grim and less to variants of TFT. Rather than deciding in advance how to handle unlikely contingencies, teams adopted simpler strategies like Grim and its variants and then adjusted on the fly to unexpected circumstances.

The most common reason teams gave for choosing to defect in early supergames was distrust of their opponent. The most frequent approach used when switching to a cooperative strategy was leading by example – cooperating for a few initial stage games to signal their willingness to cooperate, hoping that their opponent would catch on. Another common approach was learning about their opponent, where teams initially cooperated in an attempt to find out if their opponent was willing to cooperate. These two approaches, leading and learning, are distinguished by the motivation for cooperating, either trying to influence their opponent's behavior or testing for their responsiveness to cooperation.

The analysis of team dialogues has important implications for how cooperation emerges and suggests new ways of modeling this process. There have been relatively few learning models used to study IRPD games (see Dal Bó and Fréchette, 2011 and Romero and Rosokha, 2019 for exceptions), but the strong dynamics in the data point to the importance of such models. Teams

³⁸ Play by individuals and teams clearly is not identical. That said, it is qualitatively similar with cooperation rates starting low and growing with experience. Likewise, we argue that the **underlying processes are likely to be qualitatively similar as well**. For example, individuals, like teams, no doubt rely largely on improvisation as opposed to well thought out state contingent plans.

³⁹ Improvisation is somewhat different than simply changing strategies midway through a supergame. Teams in our experiments generally did not have fully specified strategies, and filled in the details in response to the actions of their opponents. Because Rosokha and Romero used direct elicitation, agents always had a complete strategy.

that led by example were consciously trying to affect their opponent's choices.⁴⁰ This suggests that learning models for IRPD games need to incorporate strategic sophistication since trying to anticipate and influence the decision process of other players is the essence of strategic sophistication.⁴¹ Teams often improvised, exploring new strategies. Models that treat teams as choosing a strategy at the beginning of each supergame and sticking with it miss much of what teams actually did. They observed what happened, thought about it, possibly engaged in some experimentation, and often adjusted their strategy. A good model of learning in IRPD games would capture this continuous process of experimentation and learning.

Teams started out cooperating less than individuals consistent with the “discontinuity effect” reported in the psychology literature for PD games. That literature attributes the lower cooperation rates for teams to greater distrust and support for self-serving choices. The data here are consistent with this explanation, although bounded rationality (myopia) also played an important role in initial decisions to defect. What the psychology literature fails to identify, due to typically implementing a single supergame in an experimental session, is that with experience the cooperation rate for teams surpasses that of individuals.

Finding that teams' choices were more stable than individuals' choices was not expected. Two explanations for this suggest themselves. First, this might reflect teams being more rational than individuals. In line with this, Proto, Rustichini and Sofianos (2019), who investigated repeated games where individuals were stratified into higher and lower cognitive ability cohorts, reported that higher cognitive ability individuals were more stable in their choices (specifically, less likely to deviate following mutual cooperation in an IRPD game). A second possibility is that team choices were more stable because team decision making was inherently biased in favor of the status quo. The dialogues support this as teams that discussed both switching strategies and the status quo generally stuck with the status quo. Switching choices was difficult for teams because unanimity was necessary to make a change, consistent with the pluralistic ignorance literature in psychology (Prentice and Miller, 1996).

⁴⁰ This relates to strategic teaching as documented by Hyndman, Terracol, and Vaksman (2009) in coordination games and Hyndman, Ozbay, Schotter, and Ehrblatt (2012) in normal form games with a unique Nash equilibrium which is on the Pareto frontier.

⁴¹ See the SEWA model of Camerer, Ho, and Chong (2002) for an example of how sophistication can be added to a model of learning.

In previous work (Cooper and Kagel, 2005) comparing strategic play by individuals and teams, we have used the truth wins (TW) model of Lorge and Solomon (1955). This was developed to model decision making by teams for logic problems that have a demonstrably correct solution, with the reference point being that a team solves the logic problem if any of its members, working independently, would have solved the problem. The TW model does *not* apply to IRPD games as there is no singular optimal strategy (i.e. a demonstrably correct solution). Instead, the optimal strategy varies depending on the strategies adopted by other players. However, playing C in St1 leads to higher average payoffs than playing D (see fn 10). Based on this, define play of C in St1 as the “empirically optimal” action. Under a version of the TW model using this alternative definition of optimality, team play is initially below the TW benchmark as teams initially cooperate less than individuals. Over time, as teams become more cooperative, team play just catches up with the lower range of the 90% confidence interval for this TW benchmark. The relatively fast growth of cooperation in the team data is not captured by the TW model, consistent with our observation that the TW model is not well suited for IRPD games. Our interpretation of the data is that the comparison between teams and individuals reflects a combination of two factors. Initially, what dominates is the discontinuity effect (i.e. teams are more sensitive to the risk of the other side defecting, and are therefore less willing to cooperate than individuals). However, teams are faster than individuals to learn how to build cooperation through tactics like leading by example. The TW model is designed to capture the adoption of eureka insights, such as realizing that leading by example is a good approach, but is poorly suited for factors like matters of judgment (e.g, risk preferences) or the lack of an unambiguously optimal strategy in IRPD games. See Appendix G for a fuller discussion of the TW model.

The results reported here are based on a single specific environment, with relatively long supergames and perfect monitoring. Both of these features likely affected the findings. The popularity of leading by example was no doubt due in part to the length of the supergames. With a continuation probability of .90, it was cheap to send the opposing team a message by choosing C for one or two stage games, compared to the potential benefits of inducing them to cooperate. Likewise, the lack of noise made it less important for teams to think about what to do under various contingencies. This may change with imperfect monitoring as teams respond to the inherent uncertainty in their opponents’ actions. The point will not be just to see how teams’ choices change

as the structure of the game varies, but to use their discussions to understand the process underlying their choices.

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*, 11(1), 1-43.
- Arad, Ayala, Kevin P. Grubiak, and Stefan P. Penczynski. 2021. "Does communicating within a team influence individuals' reasoning and decisions?" working paper.
- Arbak, Emrah and Marie Claire Villeval. 2013. "Voluntary Leadership: Motivation and Influence." *Social Choice and Welfare*, 40, 635-662.
- 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, 1-15
- Blonski, Matthias, Peter Ockenfels, and Giancarlo Spagnolo. 2011. "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence." *American Economic Journal: Microeconomics*, 3(3), 164-192.
- Breitmoser, Yves. 2005. "Cooperation, but No Reciprocity: Individual Strategies in the Repeated Prisoner's Dilemma." *American Economic Review*, 105(9), 2882–2910.
- Camerer, C, Ho, T, Chong, K. 2002. "Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games." *Journal of Economic Theory*, 104(1), 137-188.
- Casari, Marco, Jingjing Zhang, and Christine Jackson. 2016. "Same process, different outcomes: Group Performance in an Acquiring a Company Experiment." *Experimental Economics*, 19(4), 764-791.
- Cason, Timothy N. and Vai-Lam Mui. 2019. "Individual versus Group Choices of Repeated Game Strategies: A Strategy Method Approach." *Games and Economic Behavior*, 114, 128-145.
- Cooper, David J. and John R. Hamman. 2021. "Leadership and Delegation of Authority." in Marie Claire Villeval and Klaus Zimmerman (eds.), *The Handbook of Labor, Human Resources and Population Economics*, forthcoming, Berlin: Springer-Verlag.
- Cooper, David J. and Kagel, John H. 2005. "Are Two Heads Better Than One? Team versus Individual Play in Signaling Games." *American Economic Review*, 95(3), 477-509.
- Cooper, David J. and John H. Kagel. 2016. "A Failure to Communicate: An Experimental Investigation of the Effects of Advice on Strategic Play." *European Economic Review*, 82, 24-45.

- Cooper, David J., Ian Krajbich, and Charles Noussair. 2019. "Choice-Process Data in Experimental Economics." *Journal of the Economic Science Association*, 5, 1–13.
- Dal Bó, Pedro. 2005. "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games." *American Economic 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. 2019. "Strategy Choice In The Infinitely Repeated Prisoners' Dilemma." *American Economic Review*, 109 (11), 3929- 3952
- Davis, James H. 1992. "Some Compelling Intuitions about Group Consensus Decisions, Theoretical and Empirical Research, and Interpersonal Aggregation Phenomena: Selected Examples, 1950-1990." *Organizational Behavior and Human Decision Processes*, 52(1), 3-38.
- Feri, F., Irlenbusch, B., and Sutter, M. 2010. "Efficiency Gains from Team-Based Coordination—Large-Scale Experimental Evidence," *American Economic Review*. 100(4), 1892-1912.
- Embrey, Matthew, Guillaume Fréchette, and Sevgi Yuksel. 2018. "Cooperation in the Finitely Repeated Prisoner's Dilemma." *Quarterly Journal of Economics*, 133(1), 509-551.
- Fehr, Ernst and Paul M. Glimcher (eds.). 2013. *Neuroeconomics: Decision Making and the Brain*, London: Academic Press.
- Fischbacher, Urs. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments" *Experimental Economics* 10 (2), 171-178
- 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.
- Gächter, Simon, Daniele Nosenzo, Elke Renner, and Martin Sefton. 2012. "Who Makes A Good Leader? Cooperativeness, Optimism, and Leading-By-Example." *Economic Inquiry*, 50, 953-967.
- Greiner, Ben. 2015. "Subject pool recruitment procedures: organizing experiments with ORSEE", *Journal of the Economic Science Association* 1, 114 – 125.
- Harsanyi, John C. and Reinhard Selten. 1988. "A General Theory of Equilibrium Selection in Games." MIT Press Books, edition 1, vol. 1.

Hyndman, Kyle, Antoine Terracol, and Jonathan Vaksman. 2009. "Learning and Sophistication in Coordination Games." *Experimental Economics*, 12, 450 – 472.

Hyndman, Kyle, Erkut Ozbay, Andrew Schotter, and Wolf Ehrblatt. 2012 "Convergence: An Experimental Study of Teaching and Learning in Repeated Games." *Journal of the European Economic Association*, 10(3), 573–604.

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.

Kocher, M. G. and Sutter, M. 2005. "The Decision Maker Matters: Individual versus Group Behavior in Experimental Beauty-Contest Games," *Economic Journal* 115(500), 200-223.

Lorge, Irving and Herbert Solomon. 1955. "Two models of group behavior in the solution of eureka-type problems." *Psychometrika*, 20(2), pp. 139-48.

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

Maciejovsky, Boris, Matthias Sutter, David V. Budescu, and Patrick Bernau. 2013. "Teams Make You Smarter: How Exposure to Teams Improves Individual Decisions in Probability and Reasoning Tasks," *Management Science*, 59(6), 1255–1270.

Prentice, Debra, A and Miller, Dale, T. 1996. "Pluralistic Ignorance and the Perpetuation of Social Norms by Unwitting Actors," *Advances in Experimental Social Psychology*, 28, 161-209.

Proto, E., Rustichini, A. and Sofianos, A. 2019. "Intelligence, personality and Gains from Cooperation." *Journal of Political Economy*, 127 (3), 1351-1390.

Rand, David, Joshua Greene, and Martin Nowak. 2012 "Spontaneous giving and calculated greed." *Nature*, 489, 427–430.

Romero, Julian and Yaroslav Rosokha. 2018. "Constructing Strategies in the Indefinitely Repeated Prisoner's Dilemma Game." *European Economic Review*, 104, issue C, 185-219

Romero, Julian and Yaroslav Rosokha. 2019. "The Evolution of Cooperation: The Role of Costly Strategy Adjustments." *American Economic Journal: Microeconomics*, 11(1). 299-328.

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

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

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	<div>105<div>105</div></div>	<div>5<div>175</div></div>
B	<div>175<div>5</div></div>	<div>75<div>75</div></div>

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

You 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 teams 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

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

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,⁴² and agents' experience with defection in St1 of the *previous* supergame.⁴³

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.⁴⁴ 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,⁴⁵ 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

⁴² For SG1 this is set equal to 10, the expected supergame length.

⁴³ This is averaged across the two agents in a pair, and the mean value is used for SG1.

⁴⁴ 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).

⁴⁵ 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

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

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.

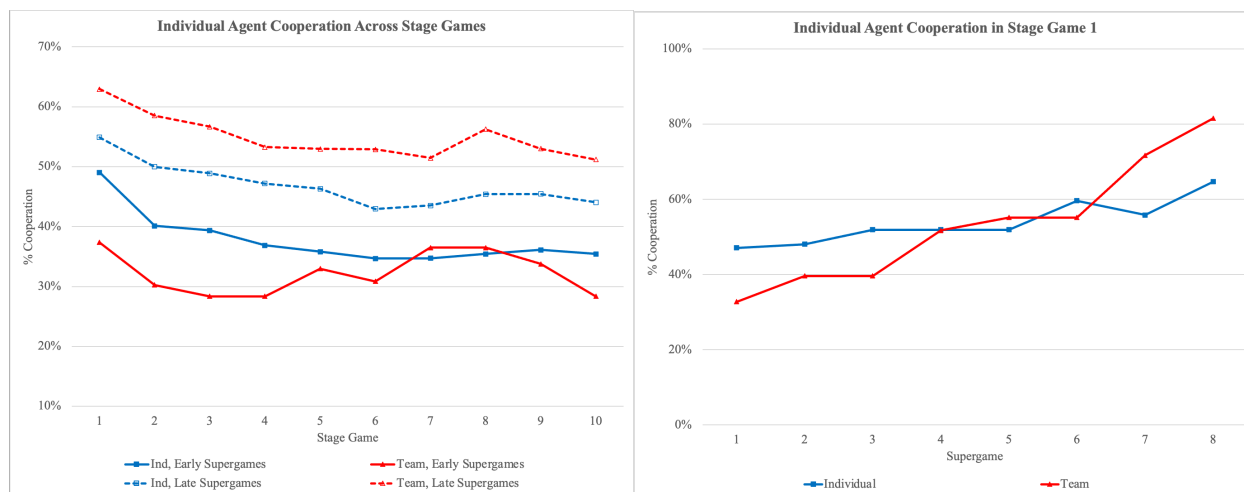
Table C2 addresses stability within supergames, but we have also run probit regressions looking at stability between supergames: if an agent started one supergame with cooperation (defection), how likely were they to start the next supergame with defection (cooperation)? We answer this question via regressions 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.⁴⁶

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 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

⁴⁶ 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.

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.

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.⁴⁷ 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 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: Three (***), two (**), and one (*) stars indicate significance at the 1%, 5%, and 10% level, respectively, using a two-tailed test. 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,

⁴⁷ 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.

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

	(1)	(2)
Dependent Variable	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

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

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 (3.2%; $\kappa = 0.981$)
 - a. Always Defect (34.8%)
 - b. Always Cooperate (0.6%)
 - c. Grim Trigger (27.8%)
 - d. Grim 2 and Grim 3 (10.1%)
 - e. Grim w/ Counting (4.0%)
 - f. TFT (2.8%)
 - g. TFT variations (e.g. TF2T, TF3T, 2TF2T, 2TFT) (0.0%)
 - h. Suspicious TFT (3.4%)
 - i. Win Stay, Lose Shift (0.0%)
 - j. Signaling (16.6%)
- 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).

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	0.780	0.800	0.650
Benefits of Mutual Cooperation	0.341	0.378	0.280	0.682
History of Previous Play	0.083	0.024	0.180	0.306
Try to Learn Opponent's Type	0.242	0.293	0.160	0.502
Lead by Example / Signal Intent	0.530	0.390	0.760	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 Fréchette's (2019) estimation of SFEM for IRPD games with perfect monitoring and $\delta = 0.90$.⁴⁸ 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.⁴⁹ 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.

⁴⁸ 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.

⁴⁹ 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%	15.80%	43.67%	19.38%
	(6.76%)	(5.37%)	(11.01%)	(10.25%)
AC	1.11%	1.11%	0.00%	0.00%
	(1.93%)	(3.74%)	(0.00%)	(0.00%)
Grim	12.16%	8.53%	2.69%	19.20%
	(7.59%)	(5.56%)	(5.26%)	(19.19%)
Grim 2	4.88%	6.32%	0.00%	12.43%
	(3.59%)	(5.99%)	(0.93%)	(7.63%)
Grim 3	0.00%	0.00%	0.00%	6.56%
	(1.07%)	(3.52%)	(0.44%)	(3.31%)
TFT	23.32%	25.02%	28.74%	27.43%
	(7.80%)	(7.39%)	(11.66%)	(15.44%)
STFT	20.66%	18.41%	19.89%	14.99%
	(8.27%)	(6.07%)	(8.03%)	(5.86%)
TF2T	4.36%	13.28%	0.00%	0.00%
	(2.43%)	(8.36%)	(3.26%)	(3.40%)
TF3T	0.00%	9.93%	0.00%	0.00%
	(2.66%)	(5.55%)	(0.81%)	(2.50%)
2TF2T	2.69%	1.60%	5.01%	0.00%
	(2.98%)	(1.65%)	(5.82%)	(3.40%)
WSLS	1.05%	0.00%	0.00%	0.00%
	-	-	-	-
Gamma	0.398	0.338	0.3137	0.278
	(0.048)	(0.024)	(0.025)	(0.025)
p(error)	7.50%	4.95%	3.96%	2.67%
Log-likelihood	-1098.10	-952.02	-380.59	-330.36

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.