

Learning, Teaching, and Turn Taking in the Repeated Assignment Game*

Timothy N. Cason^a, Sau-Him Paul Lau^b, and Vai-Lam Mui^c

^aDepartment of Economics, Krannert School of Management, Purdue University, 403 W. State St., West Lafayette, IN 47907-2056, U.S.A. (corresponding author: cason@purdue.edu; +1 765 494 1737 (phone); +1 765 494 9658 (fax))

^bSchool of Economics and Finance, University of Hong Kong, Pokfulam Road, Hong Kong

^cDepartment of Economics, Monash University, P.O. Box 11E, Clayton, Victoria 3800, Australia

August 2012

Abstract

History-dependent strategies are often used to support cooperation in repeated game models. Using the indefinitely repeated common-pool resource assignment game and a perfect stranger experimental design, this paper reports novel evidence that players who have successfully used an efficiency-enhancing turn-taking strategy will teach other players in subsequent supergames to adopt this strategy. We find that subjects engage in turn taking frequently in both the Low Conflict and the High Conflict treatments. Prior experience with turn taking significantly increases turn taking in both treatments. Moreover, successful turn taking often involves fast learning, and individuals with turn taking experience are more likely to be teachers than inexperienced individuals. The comparative statics results show that teaching in such an environment also responds to incentives, since teaching is empirically more frequent in the Low Conflict treatment with higher benefits and lower costs.

JEL Classification: C73, C91

Key words: Learning, Teaching, Assignment Game, Laboratory Experiment, Repeated Games, Turn Taking, Common-Pool Resources

* We are grateful to the Research Grants Council of Hong Kong (project HKU7223/04H) for financial support. We also thank Kyle Hyndman, two anonymous referees, and an associate editor, as well as conference and seminar participants at Monash, Purdue, Université Paris 1, the Australia New Zealand Workshop on Experimental Economics and the Economic Science Association for valuable comments. We retain responsibility for any errors. Ishita Chatterjee and Julian Chan provided valuable research assistance. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

1. Introduction

Repeated game models have been widely used by economists to study how repeated interaction may enhance cooperation. Like many theoretical models of strategic interactions, however, equilibrium analysis of repeated games only shows when cooperation can be supported as equilibrium, but is silent about how such equilibrium may arise.

One usual justification of equilibrium analysis of repeated interaction is that if players play a game sufficiently often, then they may *learn* to adopt equilibrium play. The importance of learning in generating equilibrium play in games has been actively studied by scholars (Fudenberg and Levine 1998; Camerer 2003, chapter 6). Furthermore, as suggested by some researchers (Fudenberg and Levine 1998, chapter 8; Camerer et al. 2002), some sophisticated players may anticipate that others will learn from their experience, and may incur short-term costs to *teach* others to learn to reach an equilibrium that improves their payoffs in the long run.

While teaching can potentially be important in affecting behavior in repeated games, only recently have researchers begun empirically investigating the role of teaching in repeated games. Most of this small, emerging literature focuses on how repetition affects the adoption of a particular Nash equilibrium in the one-shot game (Terracol and Vaksman 2009; Hyndman et al. 2009, 2012). Most applications of repeated game models, however, consider how players can use history-dependent strategies to support cooperation (see, for example, Mailath and Samuelson (2006) for a survey). This raises the natural further question that needs to be studied in teaching and repeated games: If players have successfully used a history-dependent repeated game strategy to cooperate in a supergame, will they teach another player in a subsequent supergame to adopt this efficiency-enhancing repeated game strategy? An affirmative answer to this question will provide support for emphasizing such history-dependent strategies in economic

applications of repeated game models in environments with ample teaching opportunities.

Using a simple 2x2 assignment game that has been influential in the study of common-pool resources (CPR) (Ostrom et al. 1994), this paper presents novel evidence that teaching is important in promoting the adoption of efficiency-enhancing history-dependent strategies. We also show that teaching responds to incentives: teaching has higher benefits and lower costs in the Low Conflict treatment in our experiment, and empirically teaching is more frequent in the Low Conflict treatment than in the High Conflict treatment.

We choose to use the assignment game to study teaching of history-dependent strategies because “taking turns” is an efficient way to cooperate in this game, and turn taking is important in facilitating cooperation in the management of CPRs (Janssen and Ostrom 2006). In addition, the one-shot version of the assignment game has a unique dominant strategy Nash equilibrium, and cooperation in the repeated assignment game requires that players take different actions in every period. These features make it easier to identify whether a player is incurring a short-term payoff cost to invest in actions that can plausibly be interpreted as teaching in the repeated game.

Table 1: The (one-shot) CPR assignment game

1\2	Good Spot (Tough)	Bad Spot (Soft)
Good Spot (Tough)	$(0.5h, 0.5h)$	(h, l)
Bad Spot (Soft)	(l, h)	$(0.5l, 0.5l)$

Table 1 illustrates the CPR assignment game. In this game, two fishermen independently decide to go to one of two fishing spots. The good spot has a value of h fish, and the bad spot has a value of l fish, where $h > l > 0$. If they choose different spots, each fisherman will obtain the respective value of the spot. If they choose the same spot, they will split the value of the spot. We

shall focus on the case of $h > 2l$, so that the good spot is much better than the bad spot.

In this game, the two asymmetric outcomes—(Tough, Soft) and (Soft, Tough)—maximize the sum of the players' payoffs, where Tough and Soft denote choosing Good Spot and Bad Spot, respectively. The players would like to coordinate on playing one of these two asymmetric outcomes. Conflict, however, is present in this game because a player prefers the asymmetric outcome in which she plays Tough to the other asymmetric outcome in which she plays Soft. Furthermore, when $h > 2l$, Tough is the dominant strategy for each player. Hence, (Tough, Tough) is the unique equilibrium in the one-shot assignment game.

One might expect, however, that if this game is played repeatedly, players may take turns to play Tough. For example, Berkes (1992) reports that fishermen in Turkey use turn taking to allocate fishing spots. A turn taking strategy is efficiency-enhancing, because it enables the players to obtain an average payoff higher than the payoff in the unique Nash equilibrium (Tough, Tough) in the one-shot game.

Besides turn taking, however, another focal subgame perfect equilibrium of this repeated game is for each player to play the stage-game dominant strategy Tough every period regardless of the past history. This is the benchmark case of non-cooperation, in which repetition does not enable the players to do any better than one-shot interaction.

While turn taking can ensure that the players will get the maximum total harvest, $h + l$, a player who is supposed to take the “bad turn” and play Soft may be tempted to play Tough to obtain the higher payoff of $0.5h$ instead of the low payoff of l . Thus, for a fixed amount of total resources ($h + l$) available, larger differences between the attractiveness of the good and the bad spot increase difficulties in coordinating on turn taking. This implies that the ratio h/l —what we refer to as “the degree of conflict”—can be important in affecting behavior in this environment.

As we shall explain later, an increase in the degree of conflict increases the costs and lowers the benefits of teaching. Our experiment includes a Low Conflict treatment and a High Conflict treatment to evaluate how differences in the degree of conflict affect teaching. In the experiment, each subject plays the same indefinitely repeated assignment game seven times, but she plays with a different opponent in each supergame. This perfect stranger design allows subjects to teach new individuals across supergames, while eliminating the possibility that they will be playing a “repeated game of repeated games.”

We find that turn taking occurs frequently in both treatments, but is more common in the Low Conflict treatment. In both treatments, successful turn taking often involves fast learning: The modal pattern among turn taking pairs involves one player choosing Soft in the first period, followed by the other choosing Soft in the second period. Subjects who have experienced successful turn taking in previous supergames typically use simple alternation between Soft and Tough to teach others to adopt an efficient turn taking equilibrium described below. In addition, when a subject who has experience in successful turn taking plays with a subject who has never adopted turn taking, the former is more likely to be the one who plays Soft in the first period and in later odd periods. We also find that experienced individuals are more likely to be teachers than inexperienced individuals. Furthermore, teaching occurs more frequently in the Low conflict treatment. While a lower degree of conflict promotes turn taking when both subjects are inexperienced in turn taking, this difference disappears when subjects are experienced in turn taking. This suggests that once subjects learn about the benefits of turn taking, experience is more important than the degree of conflict in explaining turn taking adoption.

Our study focuses on teaching history-dependent repeated game strategies across different matches of supergames, and differs from the recent work of Terracol and Vaksman

(2009), and Hyndman et al. (2009, 2012), who focus on how finite repetition affects the adoption of a particular Nash equilibrium of the stage game. Hyndman et al. (2012) consider repeated play in two games with unique pure strategy Nash equilibria, and they observe some subjects teaching others to play the Nash equilibrium when they choose Nash equilibrium actions that are not best responses to their own reported beliefs. Terracol and Vaksman (2009) and Hyndman et al. (2009) also elicit beliefs from subjects, and both present evidence that subjects are teaching by not best responding to their own reported beliefs. Terracol and Vaksman (2009) find that in their asymmetric game, the players who have more to gain from teaching others to play a preferred equilibrium are more likely to teach. Hyndman et al. (2009) manipulate the costs and benefits of teaching, and they find that teaching is more likely in their low cost-high benefit treatment.¹

In these studies it is natural to use elicited beliefs to measure teaching, because either multiple equilibria exist in the stage games (Terracol and Vaksman 2009, and Hyndman et al. 2009), or there is no dominant strategy in the stage game (Hyndman et al. 2012). The use of elicited beliefs raises potential concerns, however, such as incentives for hedging and measurement noise because beliefs are deliberately not strongly-incentivized (Hyndman et al. 2009). Besides looking at the different questions of teaching history-dependent strategies and teaching across different supergames, our study complements this earlier work by providing more direct evidence of teaching. The assignment game has a unique stage-game dominant strategy equilibrium. Hence, subjects' choice of the dominated action Soft and in alternation with Tough, is a strong indication that they are incurring costs to teach the other to take turns.²

¹ Duersch et al. (2010) study how subjects learn to play against computers that are programmed to follow one of a number of standard learning algorithms. They find that teaching occurs frequently and that all learning algorithms are subject to exploitation with the notable exception of imitation.

² Unlike the contributions discussed above but like our paper, Camerer et al. (2002) consider teaching in repeated game strategies. Their main concern, however, is on teaching by a player who faces a sequence of different players in a finitely repeated trust game.

Other researchers have also documented turn taking behavior for various repeated games in the laboratory, ranging from the game of chicken (Bornstein et al. 1997); a route choice game in traffic management that is similar to the assignment game (Helbing et al. 2005); an entry game with incomplete information (Kaplan and Ruffle 2012); an inter-group contest game (Leibbrandt and Sääksvuori 2012); to other 2x2 games with two efficient asymmetric outcomes (Prisbrey 1992, Bednar et al. 2012). Recently, biologists Harcourt et al. (2010) present experimental evidence that pairs of stickleback fish use turn taking to solve coordination and conflict problems. None of these studies, however, focus on teaching.

Our study also contributes to a growing literature on the experimental study of indefinitely repeated games. While the laboratory offers a useful environment in which one can implement a probabilistic termination design to assess the effects of indefinite repetition (Roth and Murnighan 1978), as suggested by Duffy and Ochs (2009), only recently has an emerging experimental literature exploited this possibility to identify the empirical conditions under which indefinite repetition facilitates cooperation. Not surprisingly, given the prominence of the prisoner's dilemma (PD), the majority of this literature focuses on the indefinitely repeated PD (see Blonski et al. (2011) for a detailed review).³ An important difference between the indefinitely repeated PD and the indefinitely repeated assignment game is that in the former, cooperation requires the players take the same action (both cooperate) in every period, while in the latter, cooperation requires that players take different actions (one plays Soft while the other plays Tough) in every period. Besides its empirical significance, this feature of the indefinitely repeated assignment game makes it particularly useful to study teaching, as we can exploit this

³ Notable exceptions include, for example, Engle-Warnick and Slonim (2006a,b) and Duffy et al. (forthcoming), who study cooperation in the indefinitely repeated trust game.

asymmetry to better identify who is a teacher and a learner. To our knowledge, this paper is the first contribution that focuses on the importance of teaching in indefinitely repeated games.

In this new experimental literature on indefinitely repeated games, Dal Bó and Fréchette's (2011) study on the repeated PD is closest to ours, as both papers study how experience promotes cooperation.⁴ But there are also crucial differences. In Dal Bó and Fréchette (2011), subjects play a sequence of repeated PDs, with a random stranger matching protocol. They vary both the probability of continuation and the payoff from cooperation, and find that cooperation decreases with experience when cooperation cannot be supported as equilibrium, while cooperation increases with experience when cooperation can be supported as equilibrium. They do not, however, consider teaching in their analysis. Our study varies the degree of conflict in the payoffs but does not change the probability of continuation, and we only consider cases in which cooperation (in the form of turn taking) can be supported as equilibrium. Our finding that experience has a stronger effect on increasing turn taking in the Low Conflict treatment in our repeated assignment game, and Dal Bó and Fréchette's (2011) finding that experience can decrease or increase cooperation in the repeated PD in different treatments, provide mutually reinforcing support for the general message that the influence of experience on cooperation in indefinitely repeated games depends crucially on the primitives of the game. Importantly, by focusing on the question of teaching, we provide novel evidence that teaching behavior responds to incentives and is important in shaping how experience may affect cooperation differently. In the Low Conflict treatment players have a stronger incentive to teach, and teach more than in the High Conflict treatment. This can explain why we observe a higher incidence of turn taking in

⁴ Fudenberg et al. (2012) identify the repeated game strategies most commonly used by players in an indefinitely repeated PD when intended actions are implemented with a noise. Engle-Warnick and Slonim (2006b) employ a statistical approach to identify repeated game strategies in finitely and indefinitely repeated trust games.

the late matches in the Low Conflict treatment despite the fact that the two treatments offer the same scope for experience to affect behavior.

2. Experimental Design

To study whether learning and teaching are important in affecting the adoption of turn taking behavior in the repeated assignment game, we conducted 12 sessions at the University of Hong Kong, involving 192 human subjects. Subjects were students recruited through flyers and classroom announcements from the general student population, and each subject participated in only one session of this study. The majority (86%) of subjects had never participated in a previous economics experiment, and none participated in more than one session of this study.

As illustrated in Table 1, the assignment game is completely described by the two parameters, h and l . Now consider an alternative specification of the assignment game by defining the following two parameters: $\lambda = h + l$ and $\theta = h/l$. The parameter $\lambda = h + l$ is the total value (of fish) in the good and bad spots, which is the maximum surplus available to the two players when the players achieve the asymmetric outcome. The parameter $\theta = h/l$ is the ratio of the value of the good spot to the value of the bad spot, which reflects how the total surplus in an asymmetric efficient outcome is distributed, and can be interpreted as the *degree of conflict* of the game. Using the fact that $h = \frac{\theta\lambda}{1+\theta}$ and $l = \frac{\lambda}{1+\theta}$, the assignment game can also be represented using the two parameters λ and θ , as illustrated in Table 2.⁵

We conduct both the Low Conflict and the High Conflict treatments to evaluate how changes in the degree of conflict affect teaching. The games we implemented in the experiment

⁵ Note that for any probability p with which player 2 may play Tough, player 1 gets a higher payoff by playing Tough instead of Soft, and the difference in payoff between using these two different responses is given by $\frac{\lambda}{1+\theta}(\theta - 0.5) - 0.5p\lambda$, which is increasing in θ . Hence, other things being equal, an increase in the degree of conflict increases the gain from playing Tough instead of Soft.

are illustrated in the left and right columns of Table 3, respectively. These are experimental “francs” that were converted to Hong Kong dollars at a pre-determined exchange rate. Note that the High Conflict assignment game is obtained from changing the value of θ —and only θ —from $7/3$ in the Low Conflict assignment game to 6.⁶

Table 2: A Different Specification of the CPR assignment game

1\2	Good Spot (Tough)	Bad Spot (Soft)
Good Spot (Tough)	$\left(0.5 \frac{\theta\lambda}{1+\theta}, 0.5 \frac{\theta\lambda}{1+\theta}\right)$	$\left(\frac{\theta\lambda}{1+\theta}, \frac{\lambda}{1+\theta}\right)$
Bad Spot (Soft)	$\left(\frac{\lambda}{1+\theta}, \frac{\theta\lambda}{1+\theta}\right)$	$\left(0.5 \frac{\lambda}{1+\theta}, 0.5 \frac{\lambda}{1+\theta}\right)$

Table 3: The Low Conflict Assignment Game and the High Conflict Assignment Game

1\2	Tough	Soft
Tough	(49, 49)	(98, 42)
Soft	(42, 98)	(21, 21)

The Low Conflict assignment game with
 $\lambda = 140$ and $\theta = 7/3$

1\2	Tough	Soft
Tough	(60, 60)	(120, 20)
Soft	(20, 120)	(10, 10)

The High Conflict assignment game
with $\lambda = 140$ and $\theta = 6$

In this study, we consider how learning and teaching promote the adoption of the following efficient turn-taking equilibrium: One player plays the sequence $(S, T, S \dots)$ (i.e., plays Soft in the odd periods) while the other players plays the sequence $(T, S, T \dots)$ (i.e., plays Soft in the even periods), with any defection leading to the play of the unique stage game Nash

⁶ Thus, a change from the Low Conflict game to the High Conflict game can be thought as representing a change in the physical environment, where the total amount of fish available in the community remains unchanged, but some fish had migrated to the good spot. The laboratory allows us to test comparative statics results in a controlled environment where clean *ceteris paribus* counter-factual changes in the environment faced by the players may be hard to observe in the field.

equilibrium (T, T) forever. Because Tough is the dominant strategy in the stage game, a player has the incentive to defect from turn taking behavior when she “takes the bad turn,” and one can show that supporting this equilibrium requires that the discount factor be larger than $\left(1 - 2\frac{l}{h}\right)$.

This implies that the required critical discount factors for the Low Conflict treatment and the High Conflict treatment are $1/7$ and $2/3$, respectively. In the experiment, each of the 7 groupings (“matches”) in a session is a repeated game with random termination, using a $9/10$ continuation probability. This continuation probability is chosen to ensure that it is larger than $1/7$ and $2/3$.

We chose the parameter values in Table 3 to give subjects an expected payoff equal to 70 in both treatments, if successful turn taking is established. On the other hand, the costs and benefits in deciding whether to cooperate differ in the two games. As illustrated in Table 3, if a subject in the Low Conflict treatment expects that her opponent will play the dominant strategy Tough in the current period, by playing Soft instead of Tough, she is incurring a payoff loss equals 7 in the current period. On the other hand, a subject in the High Conflict treatment who expects that her opponent will play Tough incurs a higher payoff loss of 40 by playing Soft, which suggests that teaching is more costly in the High Conflict treatment.⁷ Since the non-cooperative benchmark of repeating the stage game Nash equilibrium involves a higher payoff of 60 in the High Conflict treatment (which is higher than the Nash equilibrium payoff of 49 in the Low Conflict treatment), the relative gain from successful turn taking is also lower in the High Conflict treatment. These observations motivate our conjecture that teaching will be less likely in

⁷ More generally, suppose that a subject in the High Conflict treatment has a belief p_H that her opponent will play Tough, and a subject in the Low Conflict treatment has a belief p_L . Then the difference in the cost of teaching for such two subjects will be $(110 - 70p_H) - (77 - 70p_L)$. By design, this difference in teaching costs is constant at 33 if the two subjects hold the same belief $p_H = p_L$ in both treatments. While differences in the treatments’ degree of conflict might lead to differences in beliefs across treatments, so long as $p_H - p_L < 33/70$, the subject in the High Conflict treatment will face a higher cost of teaching.

the High Conflict treatment.⁸

A novel feature of our design is that the Low Conflict treatment and the High Conflict treatment were conducted simultaneously in a session. Because subjects play a repeated game with random termination, the realized length of the repeated game can vary significantly across matches. By conducting the two treatments in the same session, this *simultaneous treatments* design ensures that the realized lengths of the relevant supergames are identical across treatments. In the beginning of each session, the 16 participants were randomly assigned to one of two equal-sized “clusters”, with 8 participants in each cluster. The instructions (available as supplementary material from the journal website) explained that participants in both clusters make decisions using exactly the same rules, except that participants in each cluster use an earnings table that differs from the earnings table used by participants in the other cluster. Each session consists of 7 “groupings,” and a participant in a cluster is randomly matched with every other participant in the same cluster once and only once (that is, perfect strangers matching). All this information is common knowledge to the participants.

The experiment was conducted in English. The instructions employed neutral terminology; for example, the two available actions in each stage game were simply labeled as a choice between X or Y, and their playing partner was described as “the other person you are grouped with” rather than “opponent” or “partner.” All 16 participants were given the same set of instructions, and they learned the actual payoff table they would use throughout the experimental session when the instructions were completed. Subjects in one cluster did not know the payoff table used in the other cluster. At the conclusion of the instructions subjects completed a 5-question computerized quiz to ensure that they understood how to read their assigned payoff

⁸ This discussion assumes risk neutrality, but allowing for risk aversion will not change the implication that the differences in cost and benefit imply that teaching is more likely in the Low Conflict treatment.

table and other aspects of the instructions. They received HK\$3 for each correct answer on the quiz, and for any incorrect answer the subject's computer reviewed the correct answer by referencing the relevant part of the instructions.⁹ The average number of correct quiz answers was 4.7, and 76 percent of subjects answered all 5 questions correctly.

Each of the 7 groupings ("matches") in a session is a repeated game with random termination, using a 9/10 continuation probability. At the end of every period, subjects learned all actions and monetary payoffs for both persons in their grouping, and they recorded these choices and their own earnings on a hardcopy record sheet so they always had easy access to their complete history. The experimenter then rolled a ten-sided die in front of subjects to determine termination, and the match was terminated if and only if 0 was rolled. The instructions explained that re-grouping would stop after 7 matches, or if too little time remained in the session to initiate a new match. All matches in our experiment were terminated randomly according to the above procedure, and every session completed all 7 matches. (The final match of the first session had to be discarded, however, due to a software bug.) The match lengths varied from 1 period to 50 periods, with a mean of 10 periods and a median of 7 periods.¹⁰ A typical session lasted for 60 to 80 minutes. Earnings typically ranged between HK\$76 to HK\$156, with mean=HK\$104.

3. Results

We first investigate how the degree of conflict and experience affect turn taking behavior in subsection 3.1. We then consider teaching and learning in subsection 3.2.

⁹ The exchange rate was $7.8 \text{ HK\$} \approx 1 \text{ US\$}$ when the experiment was conducted.

¹⁰ The mean match length was 10.1 periods, with a median of 7 periods and an interquartile range of 4 to 13 periods. The maximum match length was 50 periods. The 12 sessions each had 7 matches, and the mean total periods per session was 69.9 with a median of 65.5 and an interquartile range of 51 to 82.5 periods.

3.1 Turn Taking: The Degree of Conflict and the Role of Experience

Result 1: Turn taking occurs more frequently in the Low Conflict treatment than in the High Conflict treatment.

Support: We define a pair of subjects as engaging in taking turn in a match if they take turns for at least two consecutive periods and continue to alternate between X and Y. When participants accomplish turn taking, they very rarely “fall off” the turn taking path: Only 9 out of the 664 matches (1%) reached and then fell off a turn taking path. We find that turn taking occurs more frequently in the Low Conflict treatment than in the High Conflict treatment, leading to greater efficiency in the Low Conflict treatment. It is difficult for subjects to reach turn taking in very short supergame matches, so we focus on turn taking rates for matches that are longer than four periods. For these matches, summarized in Figure 1, the turn taking rate is 40% (92 out of 232) for the Low Conflict treatment. This rate falls by more than one-half to 19% (43 out of 232) in the High Conflict treatment. Conservative nonparametric Wilcoxon tests, using the independent session cohorts as the unit of observation, indicate that these differences in turn taking rates are highly statistically significant (sample sizes $n=m=12$, $p\text{-value}=0.014$). We also present regression estimates that provide additional support of this result below.

Pairs who did not adopt a turn taking strategy chose the stage-game dominant strategy of Tough in 93 percent of the periods. They chose Tough in 88 percent of periods for the Low Conflict treatment and in 95 percent of the periods in the High Conflict treatment, and this difference is statistically significant using a non-parametric Wilcoxon tests that employ the independent session cohorts as the unit of observation ($n=m=12$, $p\text{-value}<0.01$).

Result 2: Prior experiences in turn taking significantly increase the incidence of turn taking behavior, for both the High and the Low Conflict treatments.

Support: Figure 1 suggests that turn taking increases in late matches in both the High Conflict treatment and the Low Conflict treatment, providing preliminary evidence that learning is important in affecting behavior and increasing efficiency. Figure 2 illustrates how experience affects the incidence of turn taking for matches that are longer than four periods. For a given match, we define a participant as *experienced in turn taking* if the participant has ever engaged in successful turn taking in *any* previous match. In the Low Conflict treatment, turn taking occurs 21 out of 103 times (20.4%) when no member of the pair has experience in turn taking. This rate increases to 28.4% (23 out of 81 times) when one member has experience in turn taking, and is 100% (48 out of 48 times) when both members have experience in turn taking. A similar pattern holds for the High Conflict treatment, although the rate of turn taking is lower in this treatment compared to the Low Conflict treatment when either no member or only one member is experienced.

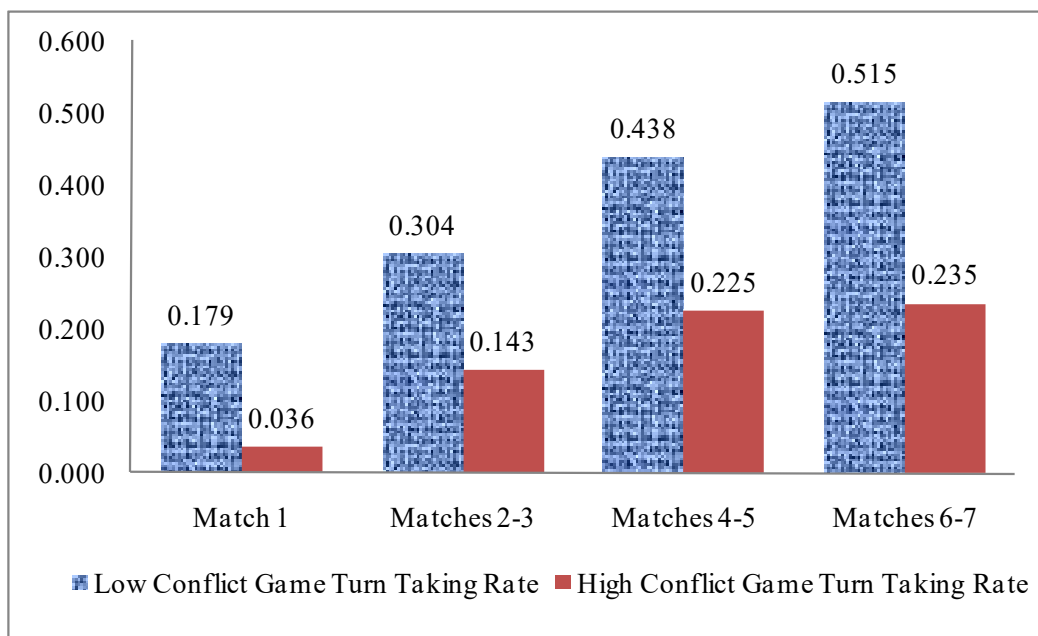


Figure 1: Rates that Pairs Adopted Turn Taking, by Treatment and Match Order, for Matches that Continued for More than 4 Periods

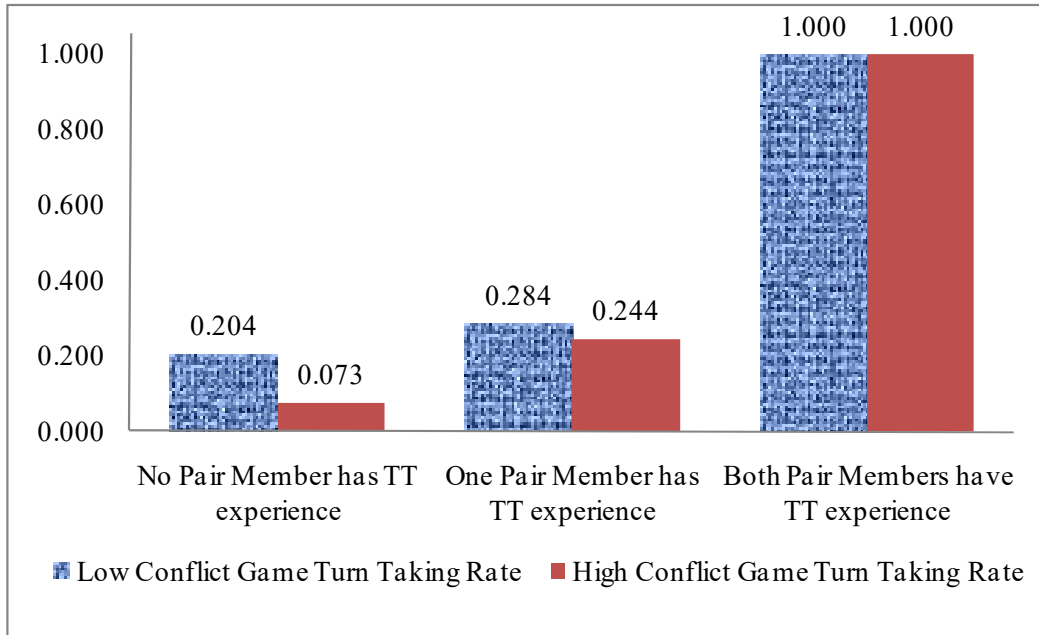


Figure 2: Rates that Pairs Adopted Turn Taking, by Treatment and Turn Taking Experience, for Matches that Continued for More than 4 Periods

Clearly experience significantly increases turn taking, and Figure 2 suggests that experience could have a more dominant effect than the differences in the degree of conflict. While a lower degree of conflict promotes turn taking when both members are inexperienced in turn taking, this difference disappears when members are experienced in turn taking.

In the logit regression shown in column (1) of Table 4, the coefficients on the turn taking experience variables are both highly significant. When both members have turn taking experience the impact on the probability that the pair will engage in turn taking is greater than when only one member has turn taking experience, providing evidence that turn taking is especially likely if both have turn taking experience (likelihood ratio test p -value <0.01). The High Conflict dummy variable is negatively significant while the “Period length of match”

variable is positive and significant, showing that turn taking is more likely for the Low Conflict treatment (Result 1) and for longer matches. This regression also includes 1/Match to allow for a nonlinear time trend across the session in the rate the pairs adopt turn taking, but this trend is never significant after accounting for experience and the degree of conflict.

Table 4: Random-Effects Logit Models of Turn Taking
Dependent Variable = 1 if the pair engaged in successful Turn Taking

	All Data (1)	No member has TT experience (2)	One member has TT experience (3)
One pair member has TT experience	1.03** (0.32)		
Both pair members have TT experience	3.77** (0.40)		
High Conflict Game (dummy)	-0.49* (0.25)	-1.05* (0.07)	-0.08 (0.44)
Period Length of Match	0.08** (0.01)	0.07** (0.01)	0.06** (0.02)
1/Match	0.06 (0.60)	0.07 (0.67)	1.07 (2.26)
Constant	-3.08** (0.44)	-2.74** (0.51)	-2.38** (0.71)
ρ (random effects)	0.004	0.006	0.082
Likelihood ratio test of $\rho=0$, p -value	0.452	0.467	0.125
Log-likelihood	-240.8	-99.7	-101.5
Observations	664	366	212

Notes: Models are estimated with random session effects. Standard errors are shown in parentheses. ** denotes significance at the one-percent level; * denotes significance at the five-percent level (all two-tailed tests).

Columns (2) and (3) report estimates separately for the cases in which no or one member of the pair has turn taking experience. (Insufficient variation exists in turn taking realizations for the case when both members have turn taking experience for reliable estimates.) The incidence of turn taking is significantly lower for the High Conflict game only when neither member has

turn taking experience in their previous supergames. In summary, these results suggest that once subjects learn about the efficiency-enhancing benefits of turn taking, experience is more important for explaining turn taking adoption than the degree of conflict. They further imply that the persistent differences between the incidence of turn taking across the Low and High Conflict treatments in the late matches illustrated in Figure 1 occur because fewer subjects experienced successful turn taking in the early matches in the High Conflict treatment.

3.2 Teaching and Learning

We now investigate the conjecture that teaching is more likely in the Low Conflict treatment discussed above. Since Tough is the stage-game dominant strategy in the assignment game, a participant's choice to play Soft alternating with Tough provides a relatively clear indication that she is trying to teach the other pair member to adopt the efficiency-enhancing turn taking strategy. A player who intends to teach others to play the efficient turn-taking equilibrium can do so by alternating between Soft and Tough. For the following analysis, we define the teacher as the pair member who first begins an alternating cycle of Soft-Tough-Soft or Tough-Soft-Tough in the periods preceding the (successful) turn-taking phase. This simple rule identifies a teacher in 84 of the 145 turn taking matches. In the remaining 61 matches the two pair members began alternating simultaneously, so this simple rule is unable to determine who is the teacher. For 57 of these 61 cases, however, one member began a Soft-Tough-Soft alternation beginning in period 1 while the other began the opposite Tough-Soft-Tough alternation. In these cases we classified the member who began with Soft-Tough-Soft as the teacher, since the other member who chose Tough in the period 1 is playing her stage-game dominant strategy.¹¹

¹¹ If the pair members who start with Tough-Soft-Tough pattern were actually teachers, they might just as well start with Soft-Tough-Soft as Tough-Soft-Tough, which would result in many "ties" where both pair members play Soft-Tough-Soft or both pair members play Tough-Soft-Tough simultaneously. But this is *not* commonly observed in the data; the vast majority of "ties" are cases in which one pair member plays Soft-Tough-Soft and the other plays

Alternative classifications of the teacher, such as the individual who first plays soft or the pair member who has lower profit in the periods preceding the turn taking phase, usually result in the same set of subjects identified as teachers and also provide identical conclusions for the key results presented below.

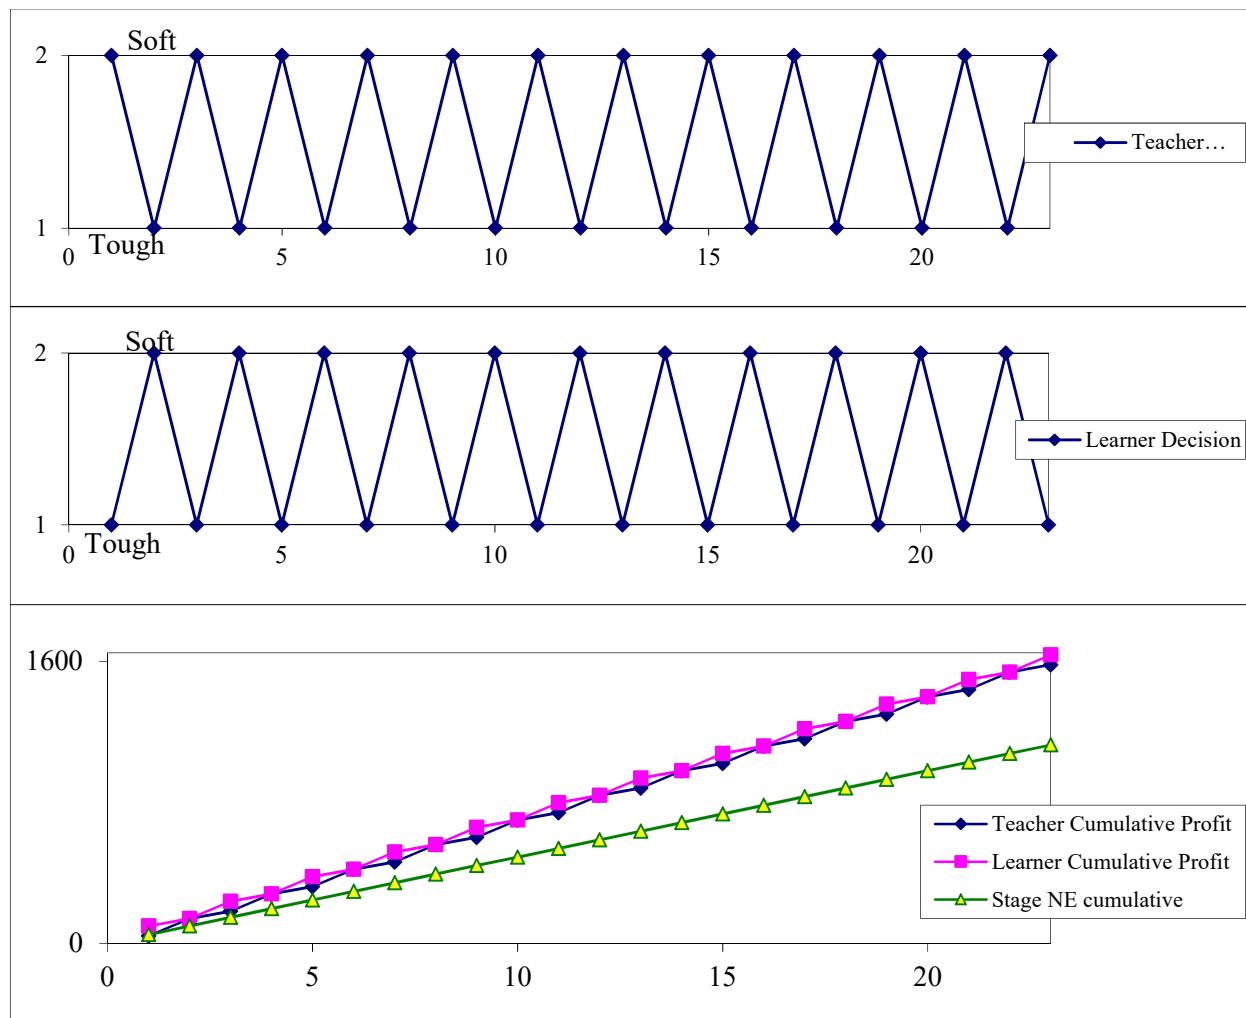


Figure 3: Fast Learning Example (Low Conflict, Match 5, Subjects 904 and 908)

Tough-Soft-Tough simultaneously. This suggests that those who play Tough-Soft-Tough are much more typically “fast learners” rather than teachers.

While there are several “patterns” in reaching the turn-taking path, two stylized patterns—“fast learning” and “slow learning”—appear in the data, with “fast learning” being much more common. As just noted, 57 turn taking matches began with one pair member choosing Soft-Tough-Soft and the other choosing Tough-Soft-Tough during the first 3 periods, immediately initiating the turn taking pattern. Figure 3 provides an example of such “fast learning” from the Low Conflict treatment. In this pair, the teacher (shown on the top panel) is a participant who has turn taking experience, and the learner (shown on the middle panel) is inexperienced. The bottom panel illustrates the gain from turn taking for both the teacher and the learner relative to the case when they play the non-cooperative strategy of Tough each period. Note that because learning was fast, the difference in cumulative payoff between the teacher and the learner is negligible.

Figure 4 shows an example of a teacher who is quite persistent, and faces a “slow” learner. The teacher (top panel) has turn taking experience, while the learner (middle panel) is inexperienced. This figure shows that successful turn taking only occurred after more than 15 periods, which included 11 periods of alternation by the teacher while the learner continued to play her stage-game dominant strategy. This teaching required modest investment costs in this Low Conflict treatment, as shown by the teacher’s cumulative profit lagging behind the stage game equilibrium profit. Only after the teacher apparently gave up this teaching did the learner begin alternating with the Soft action. The right side of the bottom panel shows that this teaching investment paid off in the long run, with more rapid growth of the teacher’s cumulative profit. It also shows that the difference in cumulative payoff between the teacher and the learner is much larger for this slow learner example than the difference shown in Figure 3 for the “fast” learning episode.

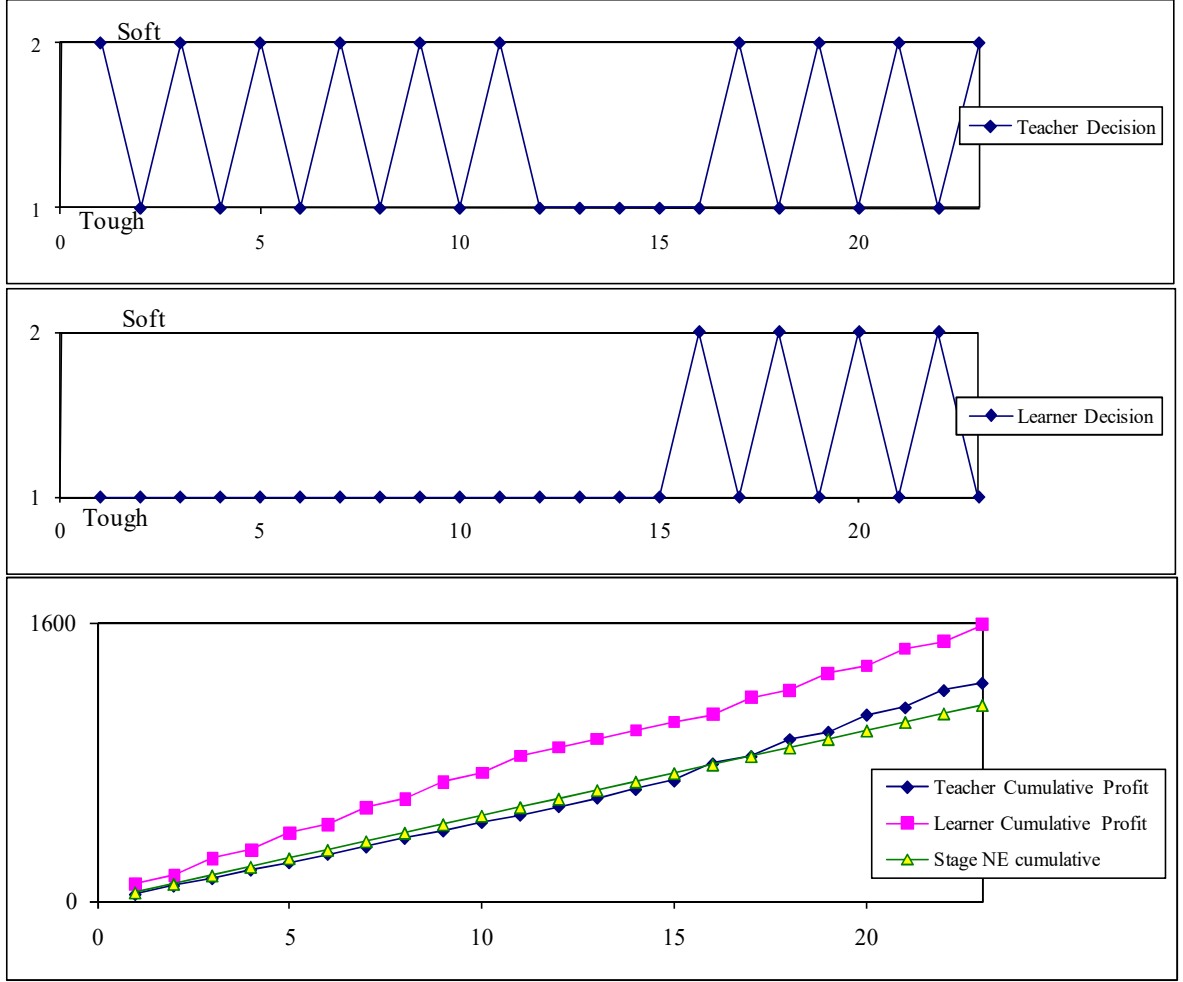


Figure 4: Slow Learning Example (Low Conflict, Match 5, Subjects 901 and 906)

Result 3: Successful turn taking often involves fast learning.

Support: The slow learning in Figure 4 represents the exception rather than the rule. The modal teaching episode is short: The teacher chooses Soft in the first period, and the learner gets the hint and chooses Soft in the second period. Overall the median number of periods required to reach the turn taking path was 3, which corresponds to the sequence (S, T, S...) for one player and the sequence (T, S, T...) for the other beginning in period 1, and the average number of periods equal to 4.2. Figure 5 shows that the average number of periods required to reach the turn taking path declines with player turn taking experience. Even when no member has turn

taking experience, nearly half of the pairs who reach the turn taking path do so within three periods. Over 80 percent of pairs who adopt turn taking when only one pair member has experience do so within five periods.¹²

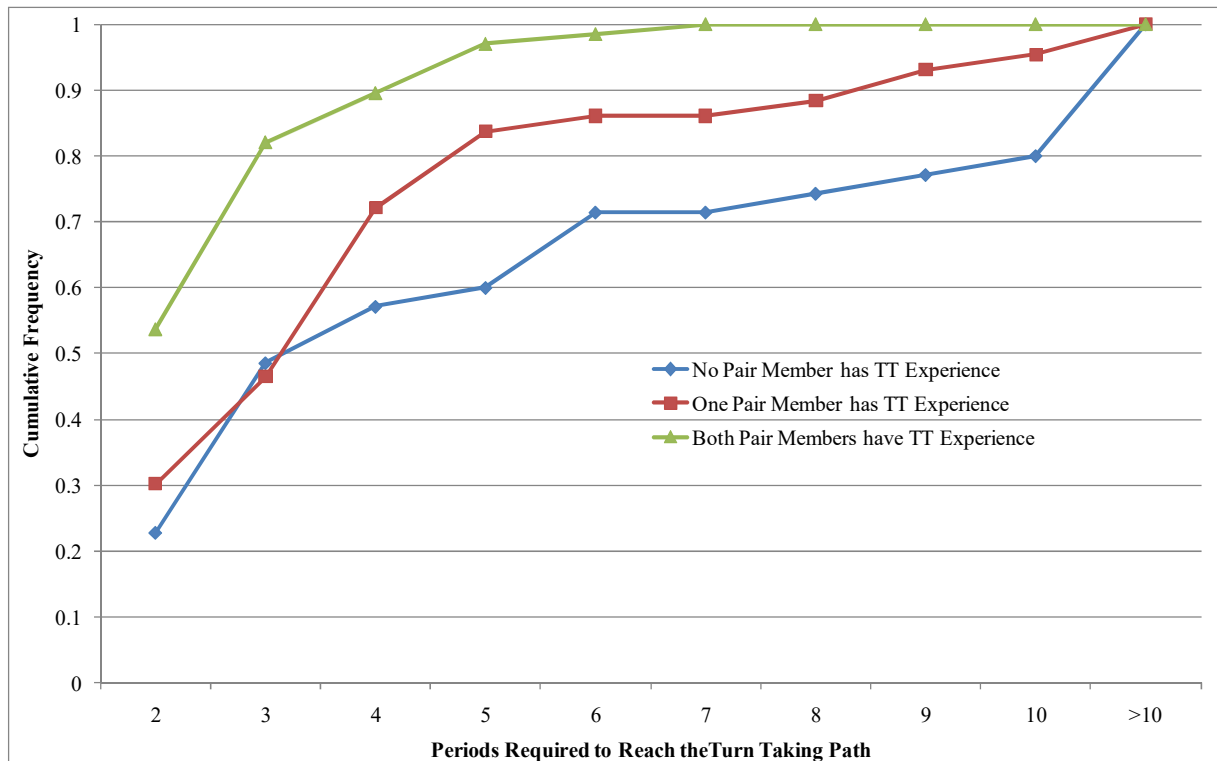


Figure 5: Cumulative Distribution of Periods Required to Reach the Turn Taking Path, by Experience

Result 4: Individuals with turn taking experience are more likely to be teachers than inexperienced individuals.

¹² Pair members who may attempt to teach others to take turns by alternating Soft-Tough-Soft sometimes encounter subjects who simultaneously play Soft in early periods. These periods of miscoordination in which both pair members choose Soft occur less than one percent of the time (65 out of 6712 period-pair observations), however. When miscoordination occurs pairs nevertheless usually reach the turn taking path, and 52 of the 65 periods of miscoordination occur during the first three periods of a match—nearly always in the first or second period. When both pair members have turn taking experience, the miscoordination rate increases to 2.6 percent in all periods of matches, and it is over 16 percent in the first period of matches.

Support: Using the earlier definition of the teacher discussed above, as the pair member who first chooses an alternating pattern, and the pair member who first chose Soft-Tough-Soft in the cases where both began alternating simultaneously, identifies a specific teacher in 141 of the 145 turn taking matches. Exactly one of the two pair members has previous turn taking experience in 42 of these 141 turn taking episodes in which the teacher is identifiable. The pair member who is experienced in turn taking is the teacher in 29 of these cases (69%).

Figure 6 illustrates how teaching and turn taking “spreads” through the population of subjects, displaying the frequencies that subjects who have different experience chose the Soft action across periods of a match for the High Conflict treatment. The low line marked with triangles indicates the low rate of Soft when neither pair member has turn taking experience, consistent with the low rate of turn taking for these cases (cf. Figure 2). For the cases when only one member of a pair has turn taking experience, the experienced subject (indicated by the line marked with diamonds) chooses Soft at a much higher rate in the early periods. In period 1, for example, this member chooses Soft nearly 40 percent of the time. The pair member without experience, marked with squares, does not choose Soft more frequently than the low rate chosen in completely inexperienced pairs during the initial periods—about 10 percent. By period 4, however, even in these aggregate data a turn taking pattern begins to emerge. The experienced and inexperienced pair members’ average Soft rates begin to alternate in a turn taking pattern for the remaining periods, with the experienced member choosing Soft more often in odd periods and the inexperienced member choosing Soft more often in even periods. Subjects in pairs where both members have turn taking experience choose Soft at rates that converge to one-half, reflecting those pairs’ uniform adoption of turn taking. Similar patterns emerge for the aggregate data in the Low Conflict treatment, although the alternating pattern in the matches with one

experienced member is less pronounced.

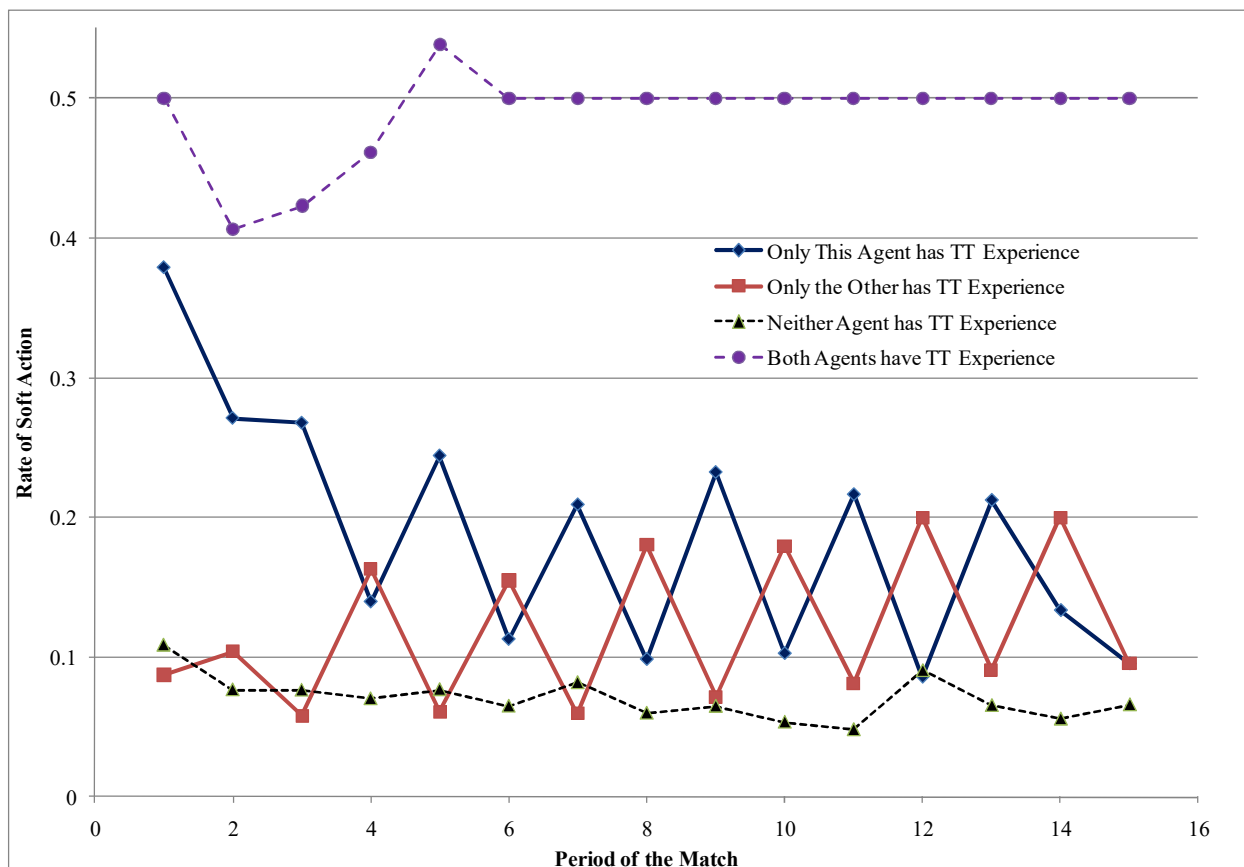


Figure 6: Rates that Subjects Chose the Soft Action by Experience, Across Different Periods of the Match, for the High Conflict Treatment

Table 5 presents a random effects logit model of subjects' choice of the Soft action to document how the propensity to play Soft depends on previous turn taking experience and on whether the period is even or odd in the case of a single pair member with experience. The dependent variable is equal to 1 when Soft is chosen, and the case of no turn taking experience for either pair member is the omitted case. The estimates show that any form of turn taking experience raises the likelihood of playing Soft, particularly when both members have experience, and Soft is more common in the low conflict game. Soft is also more likely when

only one pair member has experience, but the interaction term with the even numbered period shows that the experienced member is more likely to choose Soft only in odd numbered periods and the inexperienced member is particularly likely to choose Soft in even numbered periods. This is consistent with the visual impression provided by the alternation shown in Figure 6. A drawback of this pooled model, however, is that it does not identify how long it takes for the inexperienced members to learn turn taking and begin alternating.

Table 6 quantifies the diffusion of this turn taking adoption by re-estimating the pooled model specification of Table 5 for each of the first 15 periods of a match. The leftmost column shows that subjects who had turn taking experience are more likely to choose Soft in the first period, but of course their propensity to play Soft in the first period is unaffected if only the other subject has turn taking experience. Subjects in pairs where both members have turn taking experience are more likely to choose Soft than the omitted case of no turn taking experience throughout all periods. When a subject has experience but the other player does not, she chooses Soft at a consistently higher rate throughout all the early periods of the match. By the fourth period, this has its first statistically significant influence on the inexperienced player, increasing her likelihood of choosing Soft. This evidence is consistent with the fast learning described above. After the seventh period of the match, the estimates indicate a pattern where the experienced pair member is more likely to choose Soft in the odd periods, and the (previously) inexperienced pair member chooses Soft in the even periods. This reflects the alternating pattern often established by the experienced and teaching player, through choosing Soft in odd periods. This often leads the inexperienced, learning player to choose Soft in subsequent even periods.¹³

¹³ Lau and Mui (2012) studies how players may use the Turn Taking with Independent Randomization (TTIR) strategy to support turn taking as equilibrium in infinitely repeated 2x2 games such as the assignment game. The TTIR strategy specifies that players randomize independently between Tough and Soft in the initial phase of the repeated assignment game, and then engage in turn taking once the asymmetric outcome of either (T, S) or (S, T) is

Table 5: Random-Effects Logit Models of Subjects' Choice of the Soft Action Choice**Dependent Variable = 1 if Individual Chooses Action Soft**

Only this Individual has TT experience	0.24* (0.10)
Only Other Pair Member has TT experience	0.56** (0.14)
Both Members have TT experience	1.09** (0.10)
Only this Individual has TT experience*Even period dummy	-0.54** (0.10)
Only Other Pair Member has TT experience* Even period dummy	0.78** (0.16)
High Conflict Game (dummy)	-1.29** (0.31)
1/Match	-0.04 (0.15)
1/Period	0.10 (0.09)
Constant	-2.16** (0.23)
ρ (random effects)	0.55
Likelihood ratio test of $\rho = 0$, p -value	< 0.01
Observations	13424

Notes: Models are estimated with random subject effects. Standard errors are shown in parentheses. ** denotes significance at the one-percent level; * denotes significance at the five-percent level (all two-tailed tests).

The kind of slow learning depicted in the example shown in Figure 4 raises the possibility that “playing dumb,” that is, pretending to be a slow learner, is actually a smart strategy. It could be considered as a best response to an opponent who engages in persistent

reached, with any defection from turn taking punished by the play of the Nash equilibrium (T, T) forever. The TTIR equilibrium is designed to generate predictions about how changes in payoff parameters affect the incidence of successful turn taking for players who have not played the repeated assignment game before, and by design does not take into account how experience may affect behavior. The TTIR equilibrium correctly predicts that the incidence of turn taking is higher in the Low Conflict treatment than in the High Conflict treatment. The data, however, show that for both treatments, the TTIR equilibrium over-predicts the incidence of turn taking for the early matches in each session, but under-predicts the actual incidence of turn taking when both subjects have turn taking experience in later matches.

Table 6: Random-Effects Logit Models of Subjects' Choice of the Soft Action Choice in the First 15 Periods of a Match

Dependent Variable = 1 if Individual Chooses Action Soft

	Period of the Match														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Only this Individual	1.64**	1.13**	1.66**	1.02**	1.32**	0.83*	1.33**	0.30	1.04**	0.78	1.20**	0.26	1.21**	0.28	1.18*
has TT experience	(0.34)	(0.31)	(0.29)	(0.33)	(0.30)	(0.35)	(0.34)	(0.44)	(0.37)	(0.50)	(0.44)	(0.50)	(0.44)	(0.62)	(0.60)
Only Other Pair Mem-	0.07	0.65	-0.44	1.01**	-0.16	0.56	0.06	1.12**	0.16	1.37**	0.47	1.11*	0.45	1.55**	0.98
ber has TT experience	(0.39)	(0.33)	(0.42)	(0.34)	(0.37)	(0.37)	(0.42)	(0.41)	(0.42)	(0.50)	(0.49)	(0.48)	(0.50)	(0.58)	(0.67)
Both Members have	2.08**	2.07**	2.21**	2.31**	2.14**	2.13**	2.35**	2.35**	2.16**	2.58**	2.31**	2.03**	2.30**	2.66**	2.84**
TT experience	(0.37)	(0.32)	(0.32)	(0.33)	(0.33)	(0.35)	(0.36)	(0.42)	(0.40)	(0.53)	(0.51)	(0.54)	(0.52)	(0.63)	(0.71)
High Conflict Game	-0.27	-0.62*	-0.61*	-0.66*	-0.70**	-0.60*	-0.58*	-0.58	-0.76*	-1.14**	-1.12**	-0.65	-0.60	-0.64	-0.72
(dummy)	(0.43)	(0.29)	(0.24)	(0.26)	(0.25)	(0.28)	(0.27)	(0.34)	(0.30)	(0.42)	(0.38)	(0.39)	(0.34)	(0.46)	(0.43)
1/Match	-0.08	0.41	0.86*	0.93*	0.09	-0.90	0.95	-0.17	0.35	1.17	-0.02	0.28	1.21	1.16	2.43*
	(0.48)	(0.42)	(0.41)	(0.45)	(0.50)	(0.72)	(0.57)	(0.81)	(0.77)	(0.87)	(0.92)	(0.88)	(0.94)	(1.14)	(1.05)
Constant	-3.04**	-2.57**	-2.55**	-2.64**	-2.06**	-2.01**	-2.46**	-2.46**	-2.10**	-2.76**	-2.09**	-2.33**	-2.38**	-2.86**	-3.08**
	(0.44)	(0.35)	(0.33)	(0.35)	(0.34)	(0.40)	(0.38)	(0.46)	(0.42)	(0.55)	(0.51)	(0.53)	(0.53)	(0.67)	(0.73)
ρ (random effects)	0.63	0.37	0.20	0.24	0.20	0.21	0.15	0.30	0.18	0.39	0.29	0.37	0.06	0.35	0.14
Likelihood ratio test of $\rho=0$, p -value	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	0.02	<0.01	<0.01	<0.01	<0.01	<0.01	0.31	<0.01	0.23
Observations	1328	1232	1104	1024	928	768	704	640	592	480	464	448	336	320	256

Notes: Models are estimated with random subject effects. Standard errors are shown in parentheses. ** denotes significance at the one-percent level; * denotes significance at the five-percent level (all two-tailed tests).

teaching. So it is natural to ask whether such “sophisticated dumb” play is observed more frequently in late periods among subjects who may understand the benefits of turn taking and think others may try to teach them to take turns. The evidence, however, suggests that such sophisticated exploitation of the teachers is not widely used. First, as already documented the delay before turn taking begins is very short when subjects have turn taking experience. Second, among pairings that lasted longer than four periods but failed to result in turn taking, there were 120 cases where only one member has turn taking experience. The experienced member of the pair does not apparently try to play dumb systematically, because this member chose Soft at least once in 110 out of these 120 cases (92%). Third, among the 151 matches lasting more than 4 periods involving subjects who have turn taking experience, those with such experience play Soft at least once in 148 matches (98%). The data thus indicate that the experienced subjects try to teach, rather than play dumb.¹⁴

To provide further statistical support for Result 4, Table 7 presents evidence regarding the type of agent who tends to play the teaching role for successful turn taking matches, or who attempts to initiate turn taking for unsuccessful matches. Column (a) shows that subjects who have successful turn taking experience are more likely to be classified as the teacher in the successful matches. Column (b) provides an indication of the types of subjects who attempt to initiate turn taking by examining the factors influencing the propensity to play Soft in pairings that do not adopt turn taking. (The sample size is much larger because the unit of observation is the choice in an individual period, rather than the outcome of a multi-period match.) Subjects are more likely to play Soft, a necessary step to initiate turn taking, when they have previous turn

¹⁴ We also do not find evidence that “playing dumb” is empirically profitable. The average earnings per period for subjects who never played Soft, compared to those earned by subjects who play Soft in matches after they experience turn taking, are significantly less in the Low Conflict game (56.39 versus 62.59), and an equivalent amount in the High Conflict game (63.05 versus 63.03).

taking experience and in the Low Conflict treatment. The interaction term added in column (c) indicates that this experience effect is particularly pronounced in the High Conflict treatment.¹⁵

Result 5: Teaching attempts occur more frequently in the Low Conflict treatment than the High Conflict treatment among pairs who fail to adopt turn taking. For teaching to be profitable relative to the uncooperative equilibrium that repeats the Nash equilibrium (Tough, Tough) every period, based on realized payoffs the likelihood of successful teaching must be over ten times greater in the High Conflict compared to the Low Conflict treatment.

Support: Columns (b) and (c) of Table 7 already document the greater likelihood of choosing Soft in the Low Conflict game for pairs who do not adopt turn taking. While choosing Soft in any period is a simple indication of teaching, in order to identify which subject should be considered the teacher of the pair in matches that fail to adopt turn taking, we need to account for which subject first chooses Soft. We also need to define what matches constitute unsuccessful teaching. We consider that unsuccessful teaching occurs in a match when (i) we do not observe turn taking, (ii) one subject of the pair chooses Soft earlier than the other subject, and (iii) the match lasts for more than four periods. Criterion (iii) is included because satisfying the definition of turn taking is difficult in matches that terminated quickly. A total of 201 matches were classified as unsuccessful teaching by this definition.¹⁶ The teacher is the subject who first chooses Soft, and the non-teacher is the other subject.

¹⁵ The same interaction term is not significant in a similar specification for column (a) based on successful turn taking matches (p -value=0.20). Note that male subjects and those with a high grade point average (GPA) are more likely to choose Soft in match pairings that do not result in turn taking, and that students majoring in Economics and Finance are less likely to be identified as teachers in turn taking matches. These gender and major results are consistent with other studies concluding that women are less willing to incur risks than men (Croson and Gneezy 2009) and that economics majors tend to be less cooperative than non-economics majors (e.g., Faravelli 2007).

¹⁶ Alternatively, we could use a more stringent definition to classify a teacher as a subject who alternates between Soft and Tough. For example, we considered the definition of teaching for the unsuccessful turn taking matches to be at least one pattern of Soft-Tough-Soft by the teacher. We only observed 74 matches that could be classified as unsuccessful teaching by this definition, however. Conclusions regarding the relative costs of teaching in the two games are qualitatively similar, so to conserve space we only report the version based on the first choice of Soft.

Table 7: Random-Effects Logit Models of Teacher Identity and Subjects' Choices of Soft Action

Column (a): Dependent Variable = 1 if subject is identified as the teacher in a match pairing that reaches the turn taking path, as defined by the subject who first completed an alternating Soft-Tough-Soft or Tough-Soft-Tough cycle, or who first completed Soft-Tough-Soft cycle when the other pair member simultaneously completed a Tough-Soft-Tough cycle.

Columns (b) and (c): Dependent Variable = 1 if the subject chooses Soft in a period, for match pairings that did not reach the turn taking path.

	(a)	(b)	(c)
The subject has any successful TT experience	0.84** (0.33)	1.61** (0.16)	2.08** (0.20)
High Conflict game (dummy)	0.17 (0.32)	-0.94** (0.21)	-0.58* (0.23)
High Conflict game * any TT Experience (interaction)			-1.04** (0.26)
1/Match	1.38 (0.84)	1.34** (0.20)	1.34** (0.20)
Male (dummy)	0.48 (0.32)	0.59** (0.22)	0.60** (0.22)
Economics & Finance major (dummy)	-0.96** (0.34)	-0.45 (0.23)	-0.47* (0.23)
Native of Hong Kong (dummy)	0.21 (0.31)	0.09 (0.23)	0.10 (0.23)
High GPA (3 to 4) (dummy)	-0.03 (0.31)	0.47* (0.22)	0.45* (0.22)
First Year Student (dummy)	0.88* (0.34)	0.21 (0.23)	0.21 (0.24)
Perfect score on quiz (dummy)	0.36 (0.39)	0.58* (0.26)	0.58* (0.26)
Constant	-1.59* (0.70)	-4.38** (0.37)	-4.58** (0.38)
ρ (random effects)	0.01	0.29	0.30
Likelihood ratio test of $\rho=0$, p -value	0.115	<0.001	<0.001
Log likelihood	-1872.83	-1910.8	-1902.5
Observations	290	9226	9226

Notes: Models estimated with subject random effects. Standard errors are shown in parentheses. ** denotes significance at the one-percent level; * denotes significance at the five-percent level (all two-tailed tests).

Table 8 reports the realized per-period payoffs for the entire match, separately for the successful and unsuccessful turn taking matches and for teachers, learners and non-teachers. First, note that on average, teachers earn 3 to 4 experimental francs less per period than learners in the successful turn taking matches, but they earn 12 to 15 experimental francs less per period than non-teachers in the unsuccessful turn taking matches. Useful benchmarks for evaluating the expected profitability of teaching are shown in the far left column, based on repeated play of the one-shot Nash equilibrium. Clearly teaching pays off when it is successful, but not when it is unsuccessful. In the Low Conflict treatment, as indicated on the far right column teaching must only be successful 4.2 percent of the time for it to generate expected profits that exceed the (Tough, Tough) equilibrium benchmark of 49.¹⁷ Since teaching was successful in 96 matchings and unsuccessful in 104 matches, this 48 percent (i.e., $96/(96+104)$) success rate in the Low Conflict treatment indicates that teachers were correct to be persistent in teaching. In contrast, in the High Conflict treatment teaching must be successful more than half of the time for it to generate higher expected payoffs than 60, the payoff from repeated play of the one-shot Nash equilibrium of (Tough, Tough). Since the success rate in the High Conflict treatment was only 32 percent (i.e., $45/(45+97)$), teachers should have been less persistent in this treatment.

Table 8: Per-Period Payoffs for the Entire Match, for Teachers, Learners and Non-Teachers

		Successful Turn-taking		Unsuccessful Teaching		Minimum Success Rate for Profitable Teaching
		Teachers	Learners	Teachers	Non-Teachers	
Low Conflict	Mean Payoff	65.3	68	48.3	60.5	0.042
(49 for Tough- Tough Equil.)	(Std. Dev.)	(4.6)	(5.5)	(3.6)	(7.9)	
	Observations	96	102	104	104	
High Conflict	Mean Payoff	65.5	69.1	54.2	69.1	0.514
(60 for Tough- Tough Equil.)	(Std. Dev.)	(4.6)	(6.1)	(4.9)	(6.4)	
	Observations	45	47	97	97	

¹⁷ Since the mean payoff of successful teachers is 65.3 and that of unsuccessful teachers is 48.3, the minimum success rate of 4.2 percent is obtained by solving $q(65.3) + (1-q)(48.3) = 49$.

4. Concluding Remarks

History-dependent strategies are often used to support cooperation in repeated game models. An emerging literature has suggested that teaching can be important in affecting the adoption of equilibrium behavior, but to our knowledge, there is no study that empirically assesses whether teaching is important in the adoption of efficiency-enhancing history-dependent strategies in repeated games. Using the repeated assignment game and a perfect stranger design, this paper reports novel evidence that teaching is important in affecting the adoption of efficiency-enhancing history-dependent strategies in supergames. The comparative statics results show that teaching in repeated games also responds to incentives, since teaching is more frequent in the Low Conflict treatment with higher benefits and lower costs. We also find that successful turn taking often involves fast learning, and individuals with turn taking experience are more likely to be teachers than inexperienced individuals. Furthermore, teaching attempts occur more frequently in the Low Conflict treatment among pairs who fail to adopt turn taking.

This paper is the first to explore the importance of teaching on the adoption of efficiency-enhancing history-dependent strategies in indefinitely repeated games. Our findings suggest that the more experienced players are in playing the repeated game, and the stronger the incentives for teaching, the more likely that the efficient cooperative equilibrium will emerge as the outcome in the repeated game under consideration. A natural direction for future research is to investigate whether this result generalize to other games, such as widely studied games like the Battle of the Sexes and the Game of Chicken in which cooperation in repeated games also requires that players choose different actions in each period.

As suggested by an anonymous referee, future research can also explore whether successful teaching in a less challenging environment may actually promote teaching and

cooperation in a more challenging environment. Cooper and Kagel (2008) study how learning to play strategically in one signaling game promotes strategic play in a related signaling game, and our study can be extended to examine teaching and transfer across repeated games. For example, one can consider a treatment in which players first play a sequence of x Low Conflict indefinitely repeated assignment supergames followed by a sequence of y High Conflict indefinitely repeated assignment supergames, and a baseline treatment in which subjects play a sequence of $x + y$ High Conflict indefinitely repeated assignment supergames. A comparison of behavior from match $x + 1$ to match $x + y$ across the two treatments would indicate whether successful teaching in a less challenging environment promotes teaching and cooperation in a more challenging environment. We leave the experimental evaluation of this conjecture for future research.

References

- Bednar, J., Chen, Y., Liu, T., Page, S.: Behavioral spillovers and cognitive load in multiple games: An experimental study. *Game. Econ. Behav.* 74, 12-31 (2012).
- Berkes, F.: Success and failure in marine coastal fisheries of turkey. In Bromley, D. (ed.), *Making the Commons Work: Theory, Practice, and Policy*, pp. 161-182. Institute for Contemporary Studies, San Francisco (1992).
- Blonski, M., Ockenfels, P., Spagnolo, G.: Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence. *AEJ: Micro.* 3, 164-192 (2011).
- Bornstein, G., Budescu, D., Zamir, S.: Cooperation in intergroup, n-person, and two-person games of chicken. *J. Conflict Resolut.* 41, 384-406 (1997).
- Camerer, C.: *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ (2003).
- Camerer, C., Ho, T., Chong, J.: Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *J. Econ. Theory* 104, 137-88 (2002).
- Cooper, D., Kagel, J.: Learning and transfer in signaling games. *Econ. Theor.* 34, 415-439 (2008).
- Croson, R., Gneezy, U.: Gender differences in preferences. *J. Econ. Lit.* 47, 448-474 (2009).
- Dal Bó, P., Fréchet, G.: The evolution of cooperation in infinitely repeated games: Experimental evidence. *Am. Econ. Rev.* 101, 411-429 (2011).
- Duersch, P., Kolb, A., Oechssler, J., Schipper, B.: Rage against the machines: How subjects play against learning algorithms. *Econ. Theor.* 43, 407-430 (2010).
- Duffy, J., Ochs, J.: Cooperative behavior and the frequency of social interaction. *Game Econ. Behav.* 66, 785-812 (2009).
- Duffy, J., Xie, H., Lee Y-J.: Social norms, information, and trust among strangers: Theory and evidence. *Econ. Theor.* (forthcoming). doi: 10.1007/s00199-011-0659-x.
- Engle-Warnick, J., Slonim, R.: Learning to trust in indefinitely repeated games. *Game. Econ. Behav.* 54, 95-114 (2006a).
- Engle-Warnick, J., Slonim, R.: Inferring repeated-game strategies from actions: Evidence from trust game experiments. *Econ. Theor.* 28, 603-632 (2006b).

- Faravelli, M.: How context matters: A survey based experiment on distributive justice. *J. Public Econ.* 91, 1399-1422 (2007).
- Fischbacher, U.: z-Tree: Zurich toolbox for readymade economic experiments. *Exp. Econ.* 10, 171-178 (2007).
- Fudenberg, D., Levine, D.: *The Theory of Learning in Games*. MIT Press, Cambridge, MA (1998).
- Fudenberg, D., Rand, D., Dreber, A.: Slow to anger and fast to forgive: Cooperation in an uncertain world. *Am. Econ. Rev.* 102, 720-749 (2012).
- Harcourt, J., Sweetman, G, Manica, A., Johnstone, R.: Pairs of fish resolve conflicts over coordinated movement by taking turns. *Curr. Biol.* 20, 156–160 (2010).
- Helbing D., Schönhof, M., Stark, H-U., Holyst, J.: How individuals learn to take turns: Emergence of alternating cooperation in a congestion game and the prisoner's dilemma. *Adv. Complex Syst.* 8, 87-116 (2005).
- Hyndman, K., Ozbay, E., Schotter, A., Ehrblatt, W.: Convergence: An experimental study of teaching and learning in repeated games. *J. Eur. Econ. Assoc.* 10, 573-604 (2012).
- Hyndman, K., Terracol, A., Vaksmann, J.: Learning and sophistication in coordination games. *Exp. Econ.* 12, 450–472 (2009).
- Janssen, M., Ostrom, E.: Governing social-ecological systems. In Tesfatsion, L., Judd, K. (eds.) *Handbook of Computational Economics: Agent-Based Computational Economics*, Volume 2, pp.1465-1509. Elsevier, New York (2006).
- Kaplan, T., Ruffle, B.: Which way to cooperate? *Econ. J.* (2012). doi:10.1111/j.1468-0297.2011.02485.x.
- Lau, S.-H. Mui, V.-L.: Using turn taking to achieve intertemporal cooperation and symmetry in infinitely repeated 2x2 games. *Theor. Decis.* 72, 167-188 (2012).
- Leibbrandt, A., Sääksvuori, L.: Communication in intergroup conflicts. *Eur. Econ. Rev.* 56, 1136-1147 (2012).
- Mailath, G. J., Samuelson, L: *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, Oxford (2006).
- Ostrom, E., Gardner, R., Walker, J. (with Agrawal, A., Blomquist, W., Schlager, E., Tang, S. Y.): *Rules, Games, and Common-Pool Resources*. University of Michigan Press, Ann Arbor (1994).

- Prisbrey, J.: Cooperation in Reciprocity Games and in the Voluntary Contributions Mechanism. Ph.D. Dissertation, California Institute of Technology (1992).
- Roth, A., Murnighan, J.: Equilibrium behavior and repeated play of the prisoner's dilemma. J. Math. Psychol. 17, 189–98 (1978).
- Terracol, A., Vaksman, J.: Dumbing down rational players: Learning and teaching in an experimental game. J. Econ. Behav. Organ. 70, 54-71 (2009).

Experiment Instructions Appendix to
Learning, Teaching, and Turn Taking in the Repeated Assignment Game

by Timothy N. Cason, Sau-Him Paul Lau, and Vai-Lam Mui

Instructions

This is an experiment in the economics of strategic decision making. The Research Grant Council of Hong Kong has provided funds for this research. If you follow the instructions and make appropriate decisions, you can earn an appreciable amount of money. The currency used in the experiment is called “francs.” Your francs will be converted to Hong Kong Dollars at a rate of 60 francs to one dollar. At the end of today’s session, you will be paid in private and in cash.

It is important that you remain silent and do not look at other people’s work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

Please pay careful attention during these instructions. When the instructions are completed, you will take a short quiz on your computer to verify your understanding. You will be able to refer back to the instructions as you answer the quiz questions. The computer will record how many quiz questions you answer correctly, and you will be paid \$3 for correct answer(s) to each question.

The experiment consists of many separate decision making periods. At the beginning of the experiment you will be randomly grouped with another participant to form a two-person group. You will be grouped with this same participant for a random number of periods, as explained later. Although you will be grouped with someone in this room, you will never learn the identity of the person in your group.

Your Choice

During each period, you and the other person you are grouped with will make one choice, **X** or **Y**. You and the other person make this choice simultaneously; that is, you do not learn the choice of the other person until after you make your choice, and vice versa. Both you and the other person may choose either **X** or **Y**.

The computer program will display on the decision screen the earning table, which shows how the choice of you and the other person determine the earnings of each person, as illustrated in the following Figure.

Period

1 out of 1

Time Remaining [sec]: 27

Each period, your earning and the earning of the other person now depend on the choice made by you and the other person as shown below.

		The Other Person	
		X	Y
You	X	You receive: 1 The Other Person receives : 1	You receive: 2 The Other Person receives : 2
	Y	You receive: 3 The Other Person receives : 3	You receive: 4 The Other Person receives : 4

Please indicate your choice for this decision period. ☐ X ☐ Y

OK

Decision Screen for Person 1(Person 2's is very similar)

Your earnings from the choices each period are found in the box determined by you and the other person. If both you and the other person choose **X**, then earnings are paid as shown in the box in the upper left on the screen. If both you and the other person choose **Y**, then earnings are paid as shown in the box in the lower right on the screen. The other two boxes indicate earnings when one chooses **X** and the other chooses **Y**. To illustrate with a random example: given that earnings are determined as in the above figure, if you choose **X** and the other person, chooses **X**, then you receive 1 and the other person receives 1. You can find these amounts by looking at the appropriate box in the Figure.

The End of the Period

After everyone has made choices for the current period you will be automatically switched to the outcome screen, as shown below. This screen displays your choice as well as the choice of the other person in your group. It also shows your earnings for this period and your cumulative earnings for this grouping so far.

Period	
1 out of 1	Time Remaining [sec]: 26
<div><div>You chose</div><div>The Other Person chose</div><div>Your earnings this period</div><div>The Other Person's earnings this period</div><div>Your cumulative earnings for this grouping so far</div><div>Y</div><div>Y</div><div>21</div><div>21</div><div>21</div></div>	
<div>OK</div>	

Once the outcome screen is displayed you should record your choice and the choice of the other person in your group on your Personal Record Sheet. Also record your current and cumulative earnings for this grouping. Click on the *OK* button on the lower right of your screen when the experimenter instructs you.

The Random Ending to Each Grouping

At the beginning of the experiment, the computer will randomly match you with another participant to form a two person group. You will remain grouped with the same person in your two-person group for some random number of periods. At the end of each decision period, we will throw a ten-sided die on the floor in front of some of the participants. The outcome of the roll will be announced verbally to everyone. If the die comes up 1, 2, 3, 4, 5, 6, 7, 8 or 9, then you will remain grouped with the same participant for another period; at the end of the next period, the die will be thrown again, and again the grouping will continue for at least another round if a 1, 2, 3, 4, 5, 6, 7, 8 or 9, is thrown.

If the die comes up a 0 on any throw, then the current grouping is immediately terminated. The experiment will also be terminated at that time if one of the following conditions hold: (1) the total number of periods conducted in the experiment at that point exceeds 120 or (2) if you have already been grouped with seven different persons to form seven different two-person groups at that point, or (3) if less than 30 minutes remain in the two-hour block of time reserved for this lab session. Otherwise, you will be randomly re-grouped with another different participant to form a new two-person group. You will remain grouped with this same person for some random number of periods, with the same die-throwing rule to determine the termination of each random re-grouping of participants. Furthermore, the random grouping performed by the computer will ensure that *if you have been grouped with another participant to form a two-person group before, then you will never be grouped with this same participant in this experiment again*. Remember that you will never learn the actual identity of the individuals you are grouped with.

That is, if John is grouped with Rachel in the first grouping, John will remain grouped with Rachel for some random number of periods. When the grouping is terminated because the die comes up a 0 on a particular throw, the grouping is terminated. John will be re-grouped with another person other than Rachel, and will remain grouped with this new person for some random number of periods, and John will never be grouped with Rachel again for the rest of the experiment. Note that rule (2) above implies that in this experiment, you will at most be grouped with seven different participants to form different two-person groups.

Earning Tables and Exchange Rate

At the beginning of the experiment and before any two-person groups are formed, the 16 participants of today's experiment will be randomly divided into two equal-sized clusters, with 8 participants in each cluster. Participants in both clusters will be making decisions using exactly the same rules as explained above, except that participants in each cluster will be using an earning table that differs from the earning table used by participants in the other cluster. If a participant is randomly assigned to one of these two clusters in the beginning of the experiment, he/she remains in the same cluster for the whole experiment and he/she will only be matched with participants in the *same* cluster to form a two-person group. This means that you will be using the same earning table throughout the whole experiment, and that *whenever you are randomly matched with another participant to form a two person group, you will be matched with a participant who is also using the same earning table that you use throughout the whole experiment.*

Before we begin the experiment you will take a short quiz on your computer to verify your understanding of these instructions. Please feel free to refer back to the instructions as you answer the quiz questions. The computer will record how many quiz questions you answer correctly, and you will be paid \$3 for correct answer(s) to each question. Also feel free to raise your hand to summon an experimenter if you do not understand the explanation for a wrong answer. Please do not say anything before the experimenter comes to you, as the experimenter will answer your question in private.

The earning table in Figure 1 above provides an example regarding how the choice of you and the other person determine the earnings of each person. When we start the quiz, the *actual* earning table that indicates how your choice and the choice of that person you are grouped with determine the earnings of each person will be displayed. This is the earning table that you will be using throughout the whole experiment. Please record the information on your **Record for the Earning Table**.