Online Appendix for

Cooperation in the Finitely Repeated Prisoner's Dilemma

Matthew Embrey
U. of Sussex

Guillaume R. Fréchette
NYU

Sevgi Yuksel
UCSB

Contents

A.1 Literature Review 4

A.2 Further Analysis of the Meta Data 9

A.3 Further Analysis of the Experiment 16

A.4 Robustness Checks 19

A.4.1 Meta Data 21

A.4.2 Experiment Data 23

A.5 Further Details and Analysis of the Learning Model 26

A.5.1 Estimates 26

A.5.2 Figures 28

A.6 Sample Instructions: D8 Treatment 38

List of Figures

A1 Normalized Game Parameters . . . 10
A2 Comparison of the Size of the Basin of Attraction of AD and the Horizon . 11
A3 Evolution of Cooperation by Round and First Defection . . . 13
A4 (LHS) Mean Round to First Defection: All Pairs Versus Those That Cooperated in Round 1; (RHS) Probability of Breakdown in Cooperation . . 14
A5 Mean Cooperation Rate by Round. (See Figure 6.) . . . 17
A6 (LHS) Mean round to first defection: all pairs versus those that cooperated in round 1; (RHS) Probability of breakdown in cooperation . . . 18
A7 Average Cooperation: Simulation Versus Experimental data . . . 28
A8 Mean Round to First Defection by Supergame: Simulation versus Experimental Data . . . 29
A9 Long Term Evolution of Mean Round to First Defection by Supergame . . 29
A10 Average Cooperation Rate by Supergame: Simulation versus Experimental Data for Each Round in the Short Horizon Treatments . . . 30
A11 Long Term Evolution of Cooperation Rate for Each Round of the Short Horizon Treatments . . . 31
A12 Average Cooperation Rate by Supergame: Simulation versus Experimental Data for Each Round in D8 . . . 32
A13 Long Term Evolution of Aggregate cooperation For Each Round In E8 . . 32
A14 Long Term Evolution of Aggregate cooperation For Each Round In D8 . . 33
A15 Cumulative Distribution of Cooperation Against an AD type In E8 (Supergames 250-300) . . . 34
A16 Long term Evolution of Aggregate Cooperation For Each Round In E8 By Subset . . . 35
A17 Long term Evolution of Aggregate Cooperation . . . 35
A18 Frequency and Expected Payoff of Each Strategy . . . 36
A19 Effects of Constraining the Decline (with Experience) of Implementation Noise . . . 37

List of Tables

A1 Summary of Experiments and Sessions Included in the Meta-Study 9
A2 Marginal Effects of Correlated Random Effects Regressions for the *Standard Perspective*. (See last paragraph of page 13.) 10
A3 Marginal Effects of Correlated Random Effects Probit Regression of the Probability of Cooperating in Round One. (See Table 2.) 11
A4 Marginal Effects of Correlated Random Effects Probit Regression of the Probability of Cooperating in Round One. (Alternative Specification for Table A3.) 12
A5 Consistency of Play with Threshold Strategies. (See Table 5.) 15
A6 Session Characteristics 16
A7 Cooperation Rates and Mean Round to First Defection 16
A8 Pair-Wise Comparison of Measures of Cooperation Across Treatments. . . 16
A9 Alternative Specifications for Table A2: Marginal Effects of Correlated Random Effects Regressions for the *Standard Perspective* 21
A10 Alternative Specifications for Table A4: Marginal Effects of Correlated Random Effects Probit Regression of the Probability of Cooperating in Round One 22
A11 Alternative Specifications for Table A5: Consistency of Play with Threshold Strategies 22
A12 Alternative Specifications for Table 3: Cooperation Rates: Early Supergames (1–15) vs Late Supergames (16–30) 23
A13 Alternative Specifications for Table 4: Cooperation Rate for All Rounds in Supergames 1, 2, 8, 20 and 30 24
A14 Alternative Specifications for Table A8: Pair-Wise Comparison of Measures of Cooperation Across Treatments 25
A15 Summary statistics for long horizon treatments 26
A16 Summary statistics for short horizon treatments 27

A.1. Literature Review

Selten & Stoecker (1986) study behavior in a finitely repeated PD with a horizon of 10 rounds. Subjects play 25 supergames where they are rematched between every supergame. They observe that behavior converges to a specific pattern with experience: joint cooperation in early rounds followed by joint defection in subsequent rounds once defection is initiated by either player.[56] Importantly, they state that the point at which subjects intend to first deviate moves earlier with experience.[57] Roth (1988) summarizes these observations as follows: "in the initial [supergames] players learned to cooperate [...]. In the later [supergames], players learned about the dangers of not defecting first, and cooperation began to unravel."[58] The impression at the time is that unraveling comes about with experience. We should point out that Selten and Stoecker in their paper do not take a position on whether, with more experience, unraveling would lead to complete defection in this game. They acknowledge that unraveling might slow down such that cooperation could stabilize at some level. Furthermore, their analysis is based on results from a single set of parameters, a point noted by Selten and Stoecker, as well as Roth. Hence, to what extent these results would be robust to variations is not clear. In addition, the observation about the evolution of intended deviation round is not directly linked to the pattern of play observed in the game, as it is in part based on inferences about how players *expected* to play.[59]

Andreoni & Miller (1993) and Kahn & Murnighan (1993) directly investigate whether cooperation in the finitely repeated PD is consistent with the incomplete information model of Kreps et al. (1982). Both papers involve varying the probability that subjects interact with a pre-programmed opponent to affect the subjects' beliefs over the value of building a reputation. Because we use their data in our meta-study, we focus on Andreoni & Miller (1993).

Andreoni & Miller (1993) conducted four treatments all involving 200 choices in total. In the *Partners* treatment, these were 20 finitely repeated PDs with a horizon of 10 rounds. In the *Strangers* treatment, these were 200 one-shot PDs. The two additional treatments are variations on the *Partners* treatment, where subjects are probabilistically matched

[56]In the last five supergames, 95.6% of the supergames are consistent with that pattern, although only 17.8% of the data fits that requirement in the first five supergames.

[57]It is on average round 9.2 in supergame 13, and steadily moves down to 7.4 in supergame 25. The *intended deviation period* is computed for a subset of the data which changes by supergame, but includes almost all of the data by the end of the experiment. If there was no prior defection by either player, it is taken to be the first period at which a player defects; otherwise it is either obtained from the written comments of the subject, or inferred from reported expectations about the opponents combined in an unspecified way with past behavior and past written comments. If no defection happens, the deviation period is recorded as 11.

[58]p.1000. Roth, as Selten and Stoecker, uses the word round, to which we substituted supergame for clarity.

[59]Moreover, it is calculated on a subsample of the subject population that changes in every supergame. Consequently, it is not clear if the diminishing average is a result of subjects defecting in earlier rounds as required for unraveling, or a by-product of the changing subsample.

to play against a computer that follows the Tit-For-Tat (TFT) strategy.[60] Cooperation rates are highest in the treatment where subjects are most likely to be playing against the computer, and lowest in the *Strangers* treatment.[61] By the end of the session, in all treatments except *Strangers*, cooperation rates are above 60% in round one, stay above 50% for at least 6 rounds, then fall under 10% in the last round. In the *Strangers* treatment, cooperation rates fall below 30% in the last 10 rounds. These facts, which are consistent with the findings of Kahn & Murnighan (1993), imply that subjects' choices depend on their beliefs about the type of opponent they are faced with. This is interpreted by both papers as evidence consistent with the reputation building hypothesis of Kreps et al. (1982). Andreoni & Miller (1993) go on to note that in two of the treatments where subjects play the 10-round finitely repeated PD,[62] the mean round at which the first defection is observed in a pair increases over the course of the experiment, starting below two in the first supergame and ending above 5 in the last.[63] This observation is inconsistent with unraveling, and in contrast to the result of Selten & Stoecker (1986). Again, both papers considered a single set of payoffs and a single horizon. Hence, the contrasting results could be due to the payoffs or the different ways in which each research group constructed the relevant statistic.

Cooper et al. (1996) design an experiment to separate the reputation building and altruism hypotheses. They compare a treatment with 20 one-shot PDs to a treatment where subjects play two finitely repeated PDs with an horizon of 10 rounds.[64] They observe higher cooperation rates in the finitely repeated PD than in the one-shot treatment. Cooperation rates start above 50% in the finitely repeated game and end below, but are always lower for the one-shot game. However cooperation is significantly above zero in both treatments. Due to the limited number of repetitions, they cannot analyze the evolution of behavior. They conclude that there is evidence of both reputation building and altruism; and that neither model can explain all the features of the data on its own. As with previous studies, a single set of payoffs and horizon was considered.

Hauk & Nagel (2001) study the effect of entry-choice on cooperation levels in the finitely repeated PD with an horizon of 10 rounds.[65] A control *lock-in* treatment (with

[60] In the *Computer50* treatment this probability is 50%; in *Computer0* it is 0.1%. The TFT strategy starts by cooperating and from then on matches the opponent's previous choice.

[61] In almost all rounds cooperation rates averaged over all supergames are ordered as *Computer50* > *Partners* > *Computer0* > *Strangers*. The exception are the final two rounds where it is more or less equal in most treatments and round one where *Computer50* and *Partners* are inverted. Cooperation rates are not statistically different between the *Computer0* and *Partners* treatments, but in both cases they are significantly higher than for *Strangers* and significantly less than for *Computer50*.

[62] The *Partners* and *Computer50* treatments.

[63] The first defection round is set to 11 for a subject that never defects, otherwise it is simply the first round in which a subject defects.

[64] They use a turnpike protocol to avoid potential contagion effects (McKelvey & Palfrey 1992).

[65] Certain design choices for this paper differ significantly from the other papers discussed. Each session had seven subjects; and each subject played 10 supergames simultaneously against the remaining 6 players. ID numbers for partners were used to separate the different partners, and were randomly reassigned in the following supergame.

no ability to choose partners) is compared to two *choice* treatments where subjects are unilaterally and multilaterally given an exit opportunity with a sure payoff instead of playing the PD game. The exit option yields higher payoffs than mutual defection. Hence, an entry decision reveals intentions on how to play the game, and beliefs about how other subjects might play. Results show that entry-choice can have ambiguous effects on welfare: Conditional on entering, cooperation levels are much higher in the choice treatments. However, when the entry-choice is taken into account, overall cooperation levels are indistinguishable (unilateral choice), or significantly lower (mutual choice). The treatment differences in this paper suggest that a subject's decision to cooperate changes with beliefs about what type of opponent he is facing.

Bereby-Meyer & Roth (2006) compare play in the one-shot PD to play in the finitely repeated PD with either deterministic or stochastic payoffs.[66] The one-shot condition involves 200 rounds with random rematching whereas the finitely repeated PD has 20 supergames with an horizon of 10 rounds. They report more cooperation in round one of the repeated games than in the one-shot games. They also find that in the repeated games, with experience, subjects learn to cooperate more in the early rounds and less towards the end of the supergame. This effect is dampened with stochastic payoffs. They interpret these observations to be consistent with models of reinforcement learning: adding randomness to the link between an action and its consequences, while holding expected payoffs constant, slows learning.

Dal Bó (2005) and Friedman & Oprea (2012) both conduct finitely repeated PD experiments as controls for their respective studies, the first on infinitely repeated games and the second on continuous time games. Dal Bó (2005) looks at two stage-game payoffs with horizon of one, two or four rounds. The main focus of the paper is to compare behavior in finitely repeated games to behavior in randomly terminated repeated games of the same expected length. The results establish that cooperations rates in the first round are much higher when the game is indefinitely repeated. In the finitely repeated games, aggregate cooperation rates decline with experience. Within a supergame, there is a sharp decline in cooperation in the final rounds. However, consistent with previous findings, first round cooperation rates are higher in games with a longer horizon, and the cooperation rates in the four-horizon game is at 20% even after 10 supergames.

Friedman & Oprea (2012) study four stage-game payoffs with an horizon of 8 rounds. They find cooperation rates to increase with experience when payoffs of the stage-game are conducive to cooperation (low temptation to defect, and high efficiency gains from cooperation), but to decrease otherwise. They conclude that "even with ample opportunity to learn, the unraveling process seems at best incomplete in the laboratory data". When behavior in these treatments is compared to the continuous time version with flow payoffs, they find cooperation rates to dramatically increase. They conclude that the unraveling argument of backward induction loses its force when players can react quickly.

[66] In addition, Bereby-Meyer & Roth (2006) also vary the feedback in the stochastic condition.

They formalize this idea in terms of ϵ-equilibrium (Radner 1986). Agents determine their optimal first defection point in a supergame by balancing two opposing forces: incentives to become the first defector, and potential losses from preempting one's opponent to start defecting early. The capacity to respond rapidly weakens the first incentive and stabilizes cooperation. Both Friedman & Oprea (2012) and Dal Bó (2005) use a within subjects design, making it difficult to isolate the effect of experience.

In addition to repeated PD experiments, backward induction has been extensively studied in the centipede game.[67] In the many experimental studies on the game, subjects consistently behave in stark contrast to the predictions of backward induction.[68] The pattern of behavior observed in this game share many features to the experimental findings on the finitely repeated PD. First, round 1 behavior diverges from the predictions of subgame-perfection. In the seminal paper on the game, McKelvey & Palfrey (1992) find that even after 10 supergames, less than 10% of subjects choose to stop the game in the first round. Second, the horizon of the centipede game has a significant impact on initial behavior: the stopping rate in the first round is significantly lower in the longer horizon games (less than 2% after 10 supergames.) Third, there is heterogeneity in the subject pool with respect to how behavior changes in response to past experience. While most subjects learn to stop earlier with experience, at the individual level, some subjects never choose to stop despite many opportunities to do so. Motivated by this observation, McKelvey & Palfrey (1992) show that an incomplete information game that assumes the existence of a small proportion of altruists in the population can account for many of the salient features of their data.[69]

Several recent papers study heterogeneity in cooperative behavior and the role of reputation building in the finitely repeated PD. Schneider & Weber (2013) allow players to select the interaction length (horizon of each supergame). They find commitment to long-term relationships to work as a screening device. Conditionally cooperative types are more likely to commit to long term relationships relative to uncooperative types. While longer interactions facilitate more cooperation even when the interaction length is exogenously imposed, endogenously chosen long-term commitment yields even higher cooperation rates.

Kagel & McGee (2016) compare individual play and team play in the finitely-repeated

[67] The standard centipede game consists of two players moving sequentially for a finite number of rounds, deciding on whether to stop or continue the game. In every round, when it is one's turn to make a decision, the payoff from stopping the game is greater than the payoff associated with continuing and letting the opponent stop in the next round, but lower than the payoff associated with stopping the game in two rounds if the game continues that far. Applying backward induction gives the unique subgame perfect Nash equilibrium for the game which dictates the first player to stop in the first round.

[68] McKelvey & Palfrey (1992), Nagel & Tang (1998); Fey et al. (1996); Zauner (1999); Rapoport, Stein, Parco & Nicholas (2003); Bornstein et al. (2004).

[69] Subsequent experimental papers on the centipede game have focused on identifying how beliefs about one's opponent affects play to provide evidence for the reputation hypothesis (Palacios-Huerta & Volij (2009), Levitt, List & Sadoff (2011)).

PD.[70] Although under team play defection occurs earlier and unraveling is faster, cooperation persists in all treatments. Subjects attempt to anticipate when their opponents might defect and try to defect one period earlier, without accounting for the possibility of their opponents thinking similarly. This is interpreted to be consistent with a strong status quo bias in when to defect across super-games. The authors interpret these results as a failure of common knowledge of rationality. Analysis of team dialogues reveal beliefs regarding the strategies of the others to change significantly across supergames. This observation is in contrast to standard models of cooperation in the finitely repeated PD,

Finally, Cox et al. (2015) test the reputation building hypothesis in a sequential-move finitely-repeated PD. Cooperation can be sustained in this setting if the first-mover has uncertainty about the second mover's type. To eliminate this channel, they reveal second-mover histories from an earlier finitely repeated PD experiment to the first-mover. In contradiction to standard reputation-building explanations of cooperation in finitely repeated PDs, they find higher cooperation rates when histories are revealed. They provide a model of semi-rational behavior that is consistent with the pattern of behavior observed in the experiment. According to the model, players use strategies that follow TFT until a predetermined round and then switch to AD. Players decide how long to conditionally cooperate in each supergame based only on *naive* prior beliefs about what strategy their opponent is playing. Similar to the Kagel & McGee (2016) findings, the model does not assume any higher-level reflection about the rationality or best-response of the opponent.[71]

Mao et al. (2017) study long-term behavior in the finitely repeated prisoner's dilemma by running a virtual lab experiment using Amazon's Mechanical Turk in which 94 subjects play up to 400 supergames of a 10-round prisoner's dilemma (with random matching) over the course of twenty consecutive weekdays. While the first defection round moves earlier with experience, partial cooperation mostly stabilizes by the end of the first week. Cooperation is sustained by about 40% of the population who behave as conditional cooperators never preempting defection even when following this strategy comes with significant payoff costs.

[70]In the team play treatments each role is played by two subjects who choose their common action together after free form communication.

[71]Recently, Kamei & Putterman (2015) investigate reputation building in a finitely repeated PD where there is endogenous partner choice, and the parameters of the game allow for substantial gains from cooperation. While subjects repeatedly observe end-game effects, under the right information conditions (how much is revealed about subject's past history of play), learning to invest in building a cooperative reputation becomes the dominant force. This leads to higher cooperation rates with experience.

A.2. Further Analysis of the Meta Data

Henceforth, Andreoni & Miller (1993) will be identified as AM1993, Cooper, DeJong, Fosythe & Ross (1996) as CDFR1996, Dal Bó (2005) as DB2005, Bereby-Meyer & Roth (2006) as BMR2006, and Friedman & Oprea (2012) as FO2012.

Table A1: Summary of Experiments and Sessions Included in the Meta-Study

Experiment	Sessions	Subjects	Supergames	Horizon	g	ℓ	Within-subject variation
DB2005	4	192					horizon
	2	108	8-10	2	0.83	1.17	
	2	84	5-9	2	1.17	0.83	
	2	108	8-10	4	0.83	1.17	
	2	84	5-9	4	1.17	0.83	
FO2005	3	30					stage-game
	3	30	8	8	0.67	0.67	
	3	30	8	8	1.33	0.67	
	3	30	8	8	2.00	4.00	
	3	30	8	8	4.00	4.00	
BMR2006	4	74	20	10	2.33	2.33	
AM1993	1	14	20	10	1.67	1.33	
CDFR1996	3	30	2	10	0.44	0.78	
Total	15	340					

Just over a quarter of the sessions come from BMR2006, which implemented a stage-game with both larger gain and loss parameters. The sessions that implemented a shorter horizon–just over a quarter of the sessions–come from DB2005, which also varied horizon within subject. By varying the stage-game within-subjects, the study of FO2012 includes most of the extreme points of the set of normalized parameter combinations that have been studied.

Table A2: Marginal Effects of Correlated Random Effects Regressions for the *Standard Perspective*. (See last paragraph of page 13.)

	Round 1	Cooperation Rate Last Round	Average	Mean Round to First Defection
g	-0.04^{***} (0.008)	-0.03^{***} (0.003)	-0.01 (0.013)	-0.43^{***} (0.041)
l	-0.02^{***} (0.005)	-0.01^{***} (0.002)	-0.05^{***} (0.004)	-0.16^{***} (0.032)
Horizon	0.03^{***} (0.004)	0.01^{**} (0.002)	0.05^{***} (0.008)	0.37^{***} (0.058)
Supergame $\times \{H = 2\}$	-0.02^{***} (0.001)	-0.02^{***} (0.001)	-0.03^{***} (0.003)	-0.00 (0.009)
Supergame $\times \{H = 4\}$	-0.00^{***} (0.001)	-0.01^{***} (0.001)	-0.01^{***} (0.001)	-0.05^{***} (0.009)
Supergame $\times \{H = 8\}$	0.03^{***} (0.002)	-0.01^{***} (0.001)	0.00 (0.003)	0.25^{***} (0.015)
Supergame $\times \{H = 10\}$	0.02^{***} (0.001)	-0.01^{***} (0.001)	0.02^{***} (0.004)	0.21^{***} (0.011)
Initial Coop. in Supergame 1	0.23^{***} (0.042)	0.04^{***} (0.004)	0.16^{***} (0.025)	0.63^{***} (0.138)

Notes: For the cooperation rates, the regression model is a probit; for the mean round to first defection, it is linear. Standard errors clustered (at the study level) in parentheses. ***1%, **5%, *10% significance.
The Supergame $\times$ $\mathbf{1}\{\cdot\}$ variable takes the value of the supergame number only for those observations with the relevant horizon.
The total number of supergames varies between 5 to 10 for sessions with $H = 2$ and $H = 4$, is 8 for sessions with $H = 8$, and is either 2 or 20 when $H = 10$.

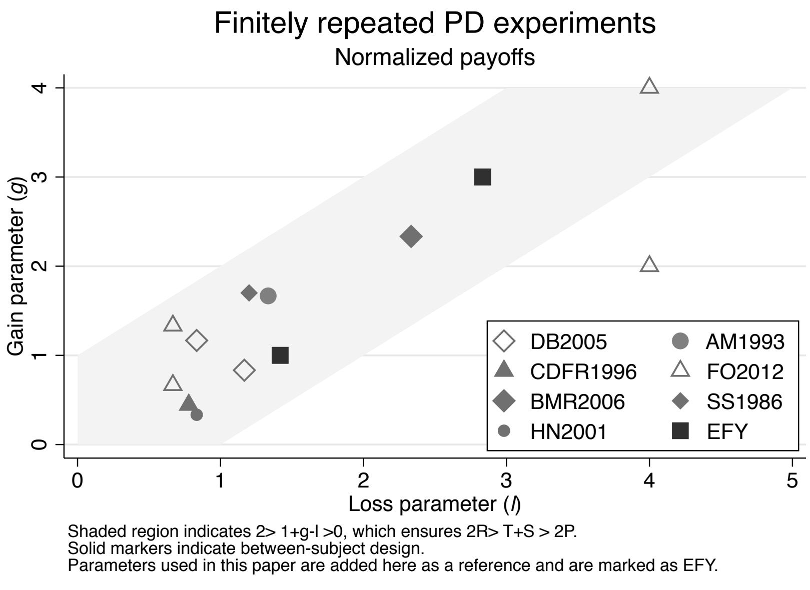

Figure A1: Normalized Game Parameters

The shaded region indicates the set of parameters for which (1) The mutual cooperation payoff is larger than the average of the sucker and temptation payoffs, thus ensuring cooperation is more efficient in the repeated game than any alternating behavior; (2) The mutual defection payoff is lower than the average of the sucker and temptation payoffs, thus ensuring that the average payoff always increases with cooperation. The sixth set of sessions included in the diagram are from our own experiment, labelled EFY.

Table A3: Marginal Effects of Correlated Random Effects Probit Regression of the Probability of Cooperating in Round One. (See Table 2.)

	(1)	(2)
g	-0.04^{***} (0.009)	-0.03^{***} (0.006)
ℓ	-0.02^{***} (0.005)	0.00 (0.005)
Horizon	0.03^{***} (0.004)	0.01 (0.005)
$sizeBAD$		-0.24^{***} (0.025)
Supergame $\times \mathbf{1}\{H=2\}$	-0.02^{***} (0.001)	-0.01^{***} (0.001)
Supergame $\times \mathbf{1}\{H=4\}$	-0.00^{***} (0.001)	-0.01^{***} (0.000)
Supergame $\times \mathbf{1}\{H=8\}$	0.03^{***} (0.002)	0.03^{***} (0.002)
Supergame $\times \mathbf{1}\{H=10\}$	0.02^{***} (0.001)	0.02^{***} (0.001)
Other Initial Coop. in Supergame - 1	0.04^{***} (0.007)	0.04^{***} (0.007)
Initial Coop. in Supergame 1	0.16^{**} (0.049)	0.15^{**} (0.049)
Observations	5398	5398

Notes: Standard errors clustered (at the study level) in parentheses. ***1%, **5%, *10% significance.
The Supergame $\times$ $\mathbf{1}\{\cdot\}$ variable takes the value of the supergame number only for those observations with the relevant horizon.
The total number of supergames varies between 5 to 10 for sessions with $H = 2$ and $H = 4$, is 8 for sessions with $H = 8$, and is either 2 or 20 when $H = 10$.

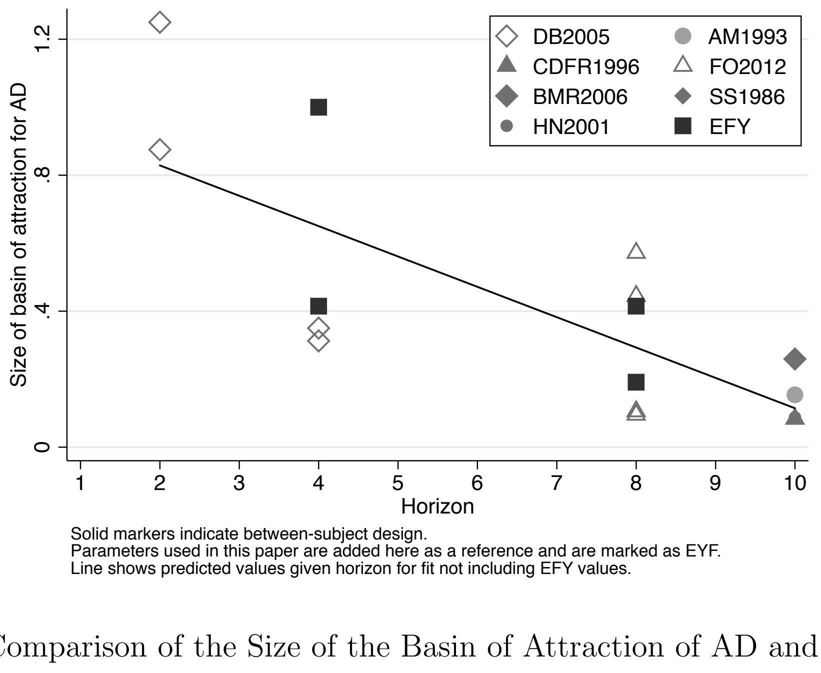

Figure A2: Comparison of the Size of the Basin of Attraction of AD and the Horizon

Table A4: Marginal Effects of Correlated Random Effects Probit Regression of the Probability of Cooperating in Round One. (Alternative Specification for Table A3.)

	(1)	(2)
g	-0.10^{***} (0.026)	-0.05^{***} (0.012)
l	0.02^{*} (0.010)	0.04^{***} (0.009)
Horizon	0.03^{***} (0.008)	-0.00 (0.009)
sizebad		-0.35^{***} (0.039)
Other Initial Coop. in Supergame - 1	0.04^{***} (0.008)	0.04^{***} (0.008)
Initial Coop. in Supergame 1	0.16^{***} (0.049)	0.16^{***} (0.051)
Supergame $\times \mathbf{1}\{g = 0.83,\ \ell = 1.17,\ H = 2\}$	-0.02^{***} (0.001)	-0.01^{***} (0.001)
Supergame $\times \mathbf{1}\{g = 1.17,\ \ell = 0.83,\ H = 2\}$	-0.02^{***} (0.003)	-0.00^{***} (0.001)
Supergame $\times \mathbf{1}\{g = 0.83,\ \ell = 1.17,\ H = 4\}$	-0.00^{*} (0.001)	-0.01^{***} (0.001)
Supergame $\times \mathbf{1}\{g = 1.17,\ \ell = 0.83,\ H = 4\}$	-0.01^{***} (0.001)	-0.02^{***} (0.001)
Supergame $\times \mathbf{1}\{g = 0.67,\ \ell = 0.67,\ H = 8\}$	0.03^{***} (0.008)	0.04^{***} (0.005)
Supergame $\times \mathbf{1}\{g = 1.33,\ \ell = 0.67,\ H = 8\}$	0.03^{***} (0.006)	0.03^{***} (0.004)
Supergame $\times \mathbf{1}\{g = 2,\ \ell = 4,\ H = 8\}$	0.01^{***} (0.002)	0.01^{***} (0.002)
Supergame $\times \mathbf{1}\{g = 4,\ \ell = 4,\ H = 8\}$	0.04^{***} (0.008)	0.03^{***} (0.005)
Supergame $\times \mathbf{1}\{g = 0.44,\ \ell = 0.78,\ H = 10\}$	-0.03 (0.041)	0.02 (0.022)
Supergame $\times \mathbf{1}\{g = 1.67,\ \ell = 1.33,\ H = 10\}$	0.02^{***} (0.002)	0.02^{***} (0.001)
Supergame $\times \mathbf{1}\{g = 2.33,\ \ell = 2.33,\ H = 10\}$	0.02^{***} (0.002)	0.03^{***} (0.002)
Observations	5398	5398

Notes: Standard errors clustered (at the study level) in parentheses. ***1%, **5%, *10% significance.
The Supergame $\times \mathbf{1}\{\cdot\}$ variable takes the value of the supergame number only for those observations with the relevant with the relevant parameters.
The total number of supergames varies between 5 to 10 for sessions with $H = 2$ and $H = 4$, is 8 for sessions with $H = 8$, and is either 2 or 20 when $H = 10$.

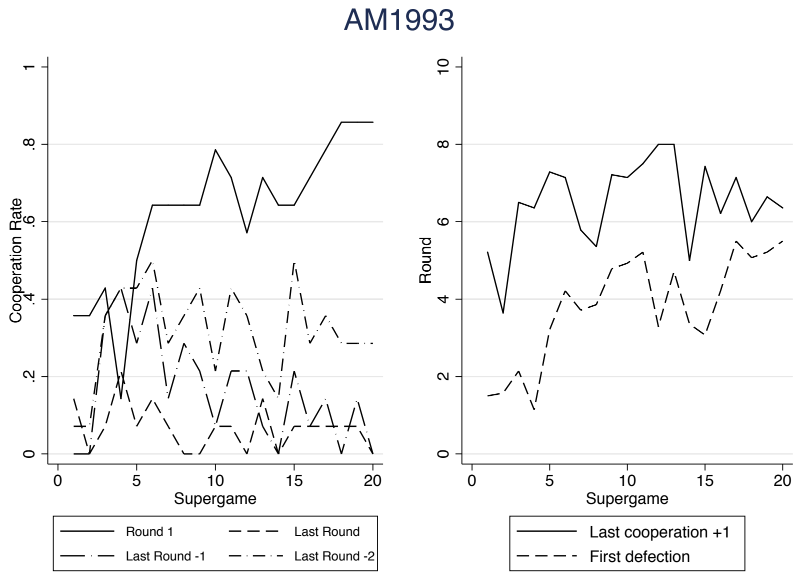

(a)

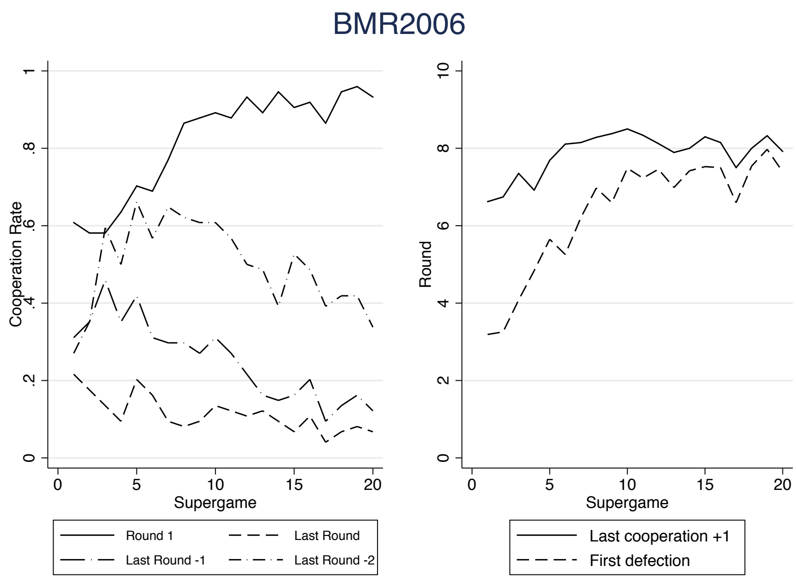

(b)

Figure A3: Evolution of Cooperation by Round and First Defection

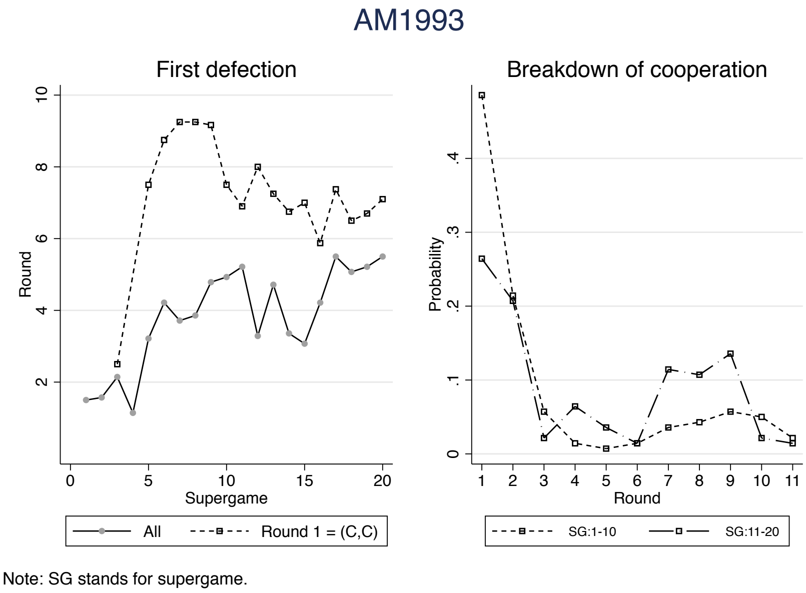

(a)

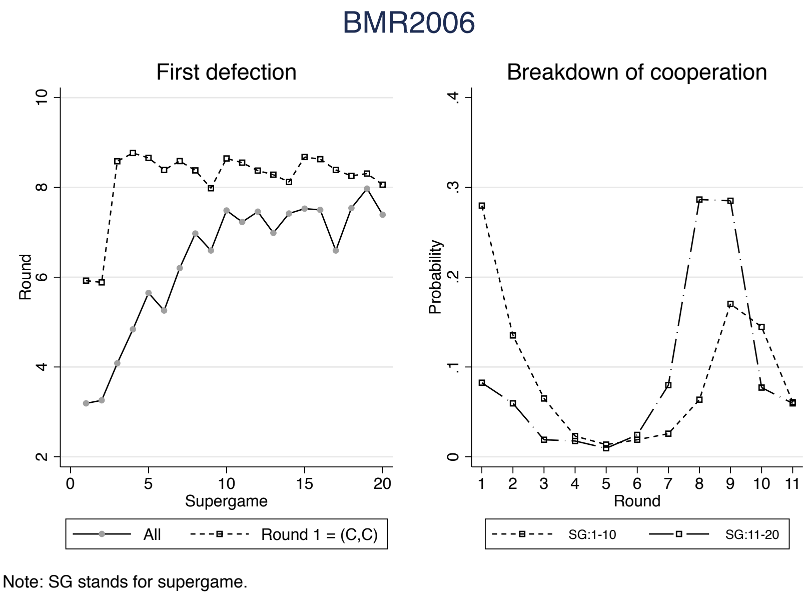

(b)

Figure A4: (LHS) Mean Round to First Defection: All Pairs Versus Those That Cooperated in Round 1; (RHS) Probability of Breakdown in Cooperation

Table A5: Consistency of Play with Threshold Strategies. (See Table 5.)

Experiment	Horizon	g	ℓ	Play Consistent With Threshold Strategy: First Supergame		Play Consistent With Threshold Strategy: Last Supergame
DB2005	2	1.17	0.83	–		–
	2	0.83	1.17	–		–
	4	1.17	0.83	0.68	$<^{***}$	0.80
	4	0.83	1.17	0.68	$<^{***}$	0.78
FO2012	8	4.00	4.00	0.43	$<^{***}$	0.90
	8	2.00	4.00	0.43	$<^{***}$	0.90
	8	1.33	0.67	0.37	$<^{***}$	0.77
	8	0.67	0.67	0.47	$<^{***}$	0.87
BMR2006	10	2.33	2.33	0.42	$<^{***}$	0.81
AM1993	10	1.67	1.33	0.29	$<^{***}$	0.79
CDFR1996	10	0.44	0.78	0.30	$<^{***}$	0.50
Meta All	.	.	.	0.52	$<^{***}$	0.79
EFY (D4)	4	3.00	2.83	0.66	$<^{***}$	0.94
EFY (E4)	4	1.00	1.42	0.66	$<^{***}$	0.94
EFY (D8)	8	3.00	2.83	0.50	$<^{***}$	0.65
EFY (E8)	8	1.00	1.42	0.57	$<^{***}$	0.89
EFY All	.	.	.	0.60	$<^{***}$	0.85

Notes: Supergame refers to supergame within a set of payoff and horizon parameters. Significance reported using subject random effects with standard errors clustered at the study level. In the meta study, the total number of supergames varies between 5 to 10 for sessions with $H = 2$ and $H = 4$, is 8 for sessions with $H = 8$, and is either 2 or 20 when $H = 10$. For the EFY experiments, the total number of supergames is either 20 or 30 for all parameter combinations.

A.3. Further Analysis of the Experiment

Table A6: Session Characteristics

	Number of		Earnings		
Treatment	Sessions	Subjects	Avg ($)	Min ($)	Max ($)
D4	3	50	14.67	12.29	17.04
D8	3	54	31.10	27.41	34.46
E4	3	62	14.92	13.34	16.28
E8	3	46	32.83	30.40	34.70

Table A7: Cooperation Rates and Mean Round to First Defection

					Cooperation Rate (%)						Mean Round to First Defection	
					Average		Round 1		Last Round			
Treatment	Supergames	H	g	ℓ	1	L	1	L	1	L	1	L
D4	30	4	3	2.83	0.32	0.07	0.48	0.15	0.14	0.00	2.0	1.1
D8	30	8	3	2.83	0.36	0.33	0.44	0.58	0.22	0.03	2.7	3.0
E4	30	4	3	1.42	0.28	0.20	0.45	0.45	0.18	0.00	1.8	1.7
E8	30	8	1	1.42	0.48	0.52	0.57	0.88	0.26	0.09	3.7	5.1

Notes: First defection is set to Horizon + 1 if there is no defection. 1: First Supergame; L: Supergame 30.

Table A8: Pair-Wise Comparison of Measures of Cooperation Across Treatments.

	All rounds				Round 1				First defect			
	D4	D8	E4	E8	D4	D8	E4	E8	D4	D8	E4	E8
Supergames 1–15												
D4	15.4	$<^{**}$	$<^{*}$	$<^{***}$	29.1	$<^{**}$	$<^{*}$	$<^{***}$	1.5	$<^{***}$	$<^{**}$	$<^{***}$
D8		34.6	$>$	$<^{***}$		49.3	$>$	$<^{***}$		2.8	$>^{***}$	$<^{***}$
E4			28.0	$<^{***}$			49.0	$<^{***}$			1.9	$<^{***}$
E8				60.1				79.7				5.3
Supergames 16–30												
D4	9.0	$<^{***}$	$<^{**}$	$<^{***}$	19.5	$<^{***}$	$<$	$<^{***}$	1.3	$<^{***}$	$<^{**}$	$<^{***}$
D8		33.2	$>^{***}$	$<^{***}$		57.1	$>$	$<^{***}$		3.1	$>^{***}$	$<^{***}$
E4			21.2	$<^{***}$			45.2	$<^{***}$			1.7	$<^{***}$
E8				55.2				88.2				5.3

Notes: The symbol indicates how the cooperation rate of the row treatment compares (statistically) to the column treatment. Significance reported using subject random effects and clustered (session level) standard errors.
***1%, **5%, *10%.

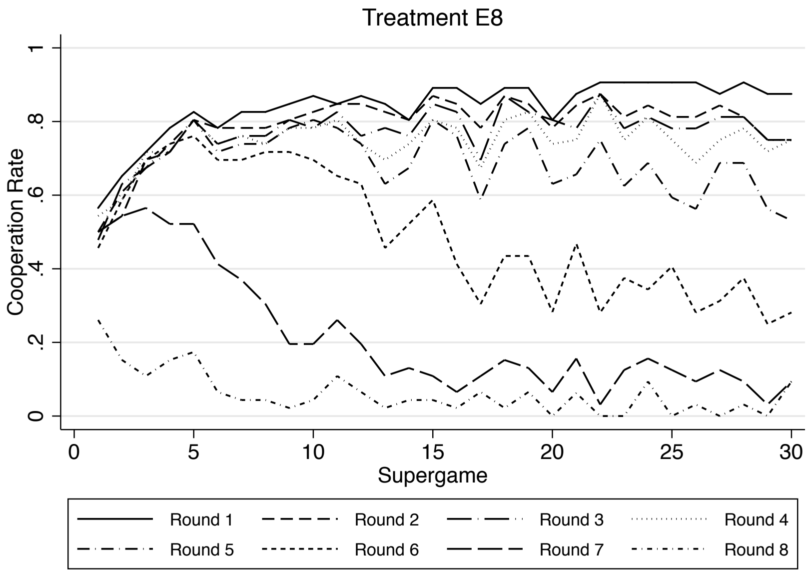

Figure A5: Mean Cooperation Rate by Round. (See Figure 6.)

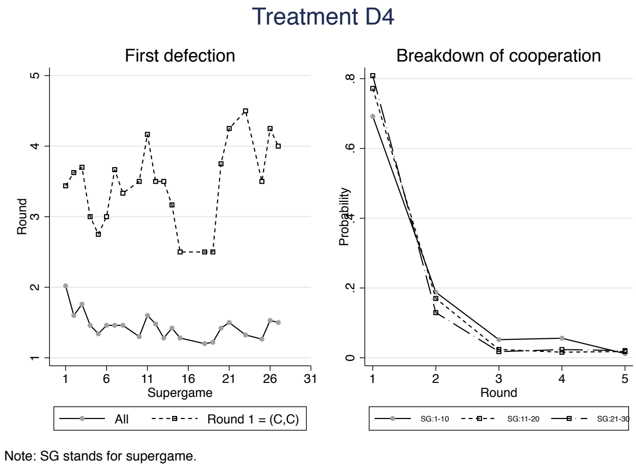

(a)

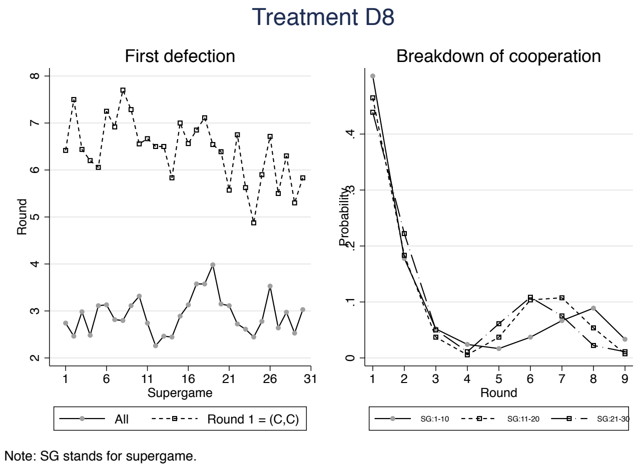

(b)

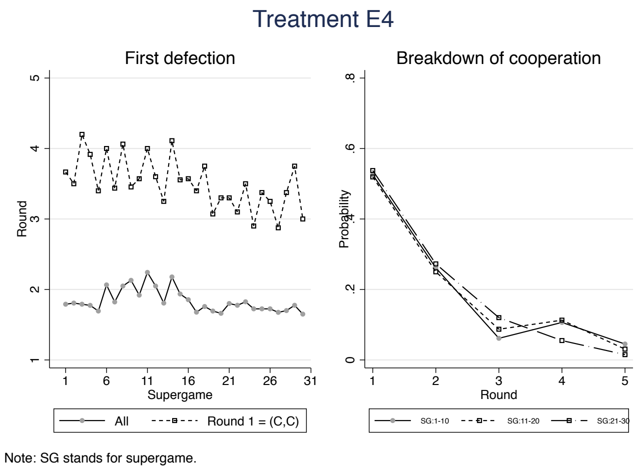

(c)

Figure A6: (LHS) Mean round to first defection: all pairs versus those that cooperated in round 1; (RHS) Probability of breakdown in cooperation

A.4. Robustness Checks: Alternative Specifications to Evaluate Statistical Significance

The data analysis reported in the main body of the text uses two main specifications: probit with subject-level random effects and variance-covariance clustered at the level of the paper for the meta data, and probit with subject-level random effects and variance-covariance clustered at the level of the session for analysis of the data from our own experiment (or for paper specific tests from the meta).[72] These specifications are meant to account for heterogeneity across subjects as well as potential, unmodeled correlations that emerge due to the interactions of subjects within a session, or to study-specific idiosyncrasies (see Fréchette 2012, for a discussion of session-effects). One potential concern with this approach is that having a low number of clusters can lead to corrected standard-errors that do not have the correct coverage probability in finite samples–see, for example, Cameron & Miller (2015) for a recent survey.

Papers that establish the extent of the problem and the effectiveness of various alternatives mostly rely on simulation studies (see, for example, Bertrand et al. 2004, Cameron et al. 2008). These simulations, however, are not geared towards data typically arising from laboratory experiments. For example, the extent of heterogeneity across clusters in the number of observations, the realisation of covariates and the error variance-covariance matrix are all important factors for understanding the potential for over-rejection when using cluster robust standard errors (see, amongst others, Imbens & Kolesar 2016, MacKinnon & Webb 2017, Carter et al. 2017). These are all dimensions on which data from laboratory studies can be expected to differ substantially from the data for which these simulation studies were designed for–indeed, these dimensions are likely to vary between laboratory studies given that details such as matching group size and number, re-matching protocol, and feedback are all experimental design choices.

Nonetheless, this a potential concern, and this appendix explores alternative specifications for the results reported in the paper. One approach is to model within cluster dependency more explicitly. We do this by estimating specifications with paper, session, and subject random effects, or session and subject random effects, as the case may be. Another approach is to remain agnostic about the form of the dependency between observations at the highest level (paper or session), while using bootstrap methods that are designed to provide proper coverage in cases with a small number of clusters. For this, we use a score-based wild bootstrap procedure (Kline et al. 2012) with a six point random weight distribution (Webb 2014).[73] To our knowledge, this is the only bootstrap-based

[72] A few specifications involve the equivalent linear version of these two when the dependent variable is not binary, such as the first round of defection.

[73] For the specifications that use a linear model, a wild bootstrap t-testing procedure (Cameron et al. 2008) is used, again with a six point random weight distribution (Webb 2014). For both the score and wild bootstrap-t procedures, the null hypothesis is imposed before applying random weights to residuals or scores.

procedure developed so far to deal with a small number of clusters that can be used when estimating a probit (see also Cameron & Miller 2015).[74] However, for these specifications we have to drop the subject-level random effect, thus ignoring a main feature of the panel structure of the data. We note that we find a great deal of evidence for the importance of subject-level random effects in our data, which are typically more important than session or paper level effects in the models that we estimate with multiple levels. By not explicitly taking into account an important source of within cluster error correlation, this potentially magnifies the small cluster problem. Nonetheless, this agnostic specification provides a useful benchmark as the non-panel estimator of the coefficients is necessarily less efficient than the panel estimator under the usual exogeneity assumption of the random effects model.

The tables in this appendix reproduce all of the main statistical tests. All tables report the p-value for the t-test of the approach in the main text (labeled CR-t for cluster robust), the p-value for the t-test of the multi-level random effects model (labeled RE-t for random effect model for the highest cluster level), and the p-value for the score-based wild bootstrap t-test of the probit specification (labeled Bt-t for bootstrap). In the few cases where the dependent variable is not dichotomous, the linear version of these is reported. In cases where estimates of the regression are of interest, we also report the respective marginal effects, to show how the magnitude of the estimated effects vary with the specification. Note that the p-values are not of the marginal effects, but of the actual coefficients from the underlying model estimated.

Although the p-values vary with the estimation method, the main results of the paper remain. For instance, here are some of the important results: The fact that most of the impact of the horizon on round one cooperation rates in the meta is absorbed by $sizeBAD$ remains true in all estimations (see Table A10). The finding that round one cooperation rates are not statistically different when comparing treatments D8 and E4 from our experiment is true in all specifications (see the D8 vs E4 rows for the Round 1 block in Table A14, which shows this separately for early and late supergames; the result also holds combining all supergames, with p-values for the CR-t, Bt-t and RE-t tests of 0.45, 0.54 and 0.61, respectively). The observation that the play of threshold strategies increases between the first and last supergame of a session is true in all specifications for the data of our experiment, as well as the meta data.

[74] An alternative bootstrap method that is generally applicable for a wide variety of estimators is the pairs cluster bootstrap, which resamples with replacement from the sample of clusters. However, with very few clusters this method can run into a number of implementation problems. See, for example, Cameron & Miller (2015) for details. Another alternative is to use the linear probability model instead of the probit, and then use the more commonly applied wild bootstrap t-testing procedure (Cameron et al. 2008), again with a six point random weight distribution (Webb 2014). Given this approach did not produce any notable differences in robustness of the main results–as well as the potential problems for the linear probability model when the regressors are no longer just a complete set of treatment indicator variables, as is the case with regressions including g, ℓ and $Horizon$–we only report the results of a bootstrap method that keeps the functional form of the limited dependent variable model fixed throughout.

A.4.1 Meta Data

In addition to what is described above, Tables A9 and A10 also report the p-value for the t-test on the estimated coefficients of the non-panel probit model using the standard cluster robust variance-covariance estimator (in addition to the bootstrapped version). This is to give a sense of what drives the changes between the random effects probit CR-t in the text and the probit Bt-t: part of it is from the bootstrapping, but part of it is simply the result of dropping the subject random effects. Table A9 shows variations across specifications. In particular, none of g, ℓ, or $Horizon$ are statistically significant when using bootstrapped standard errors for any of round one, the last round, all rounds, or the round of first defection. On the other hand, the multi-level random effects almost exclusively finds statistically significant effects. Importantly, note that the lack of significance when bootstrapping does not mean that g, ℓ, and $Horizon$ do not matter as the next table makes it clear.

Indeed, Table A10 revisits the estimation of the determinants of round one cooperation controlling for experience. The main result is the significance of $sizeBAD$ in all specifications. This confirms that the effect of g, ℓ, and $Horizon$ can be summarized by how it affects the value of cooperation. Clearly there could be additional effects of these parameters that $sizeBAD$ does not fully capture, but it accounts for an important part of the variation.

Table A9: Alternative Specifications for Table A2: Marginal Effects of Correlated Random Effects Regressions for the *Standard Perspective*

	RE Probit		Multiple REs		Probit		
	ME	CR-t	ME	RE-t	ME	CR-t	Bt-t
Round 1							
g	−0.04	0.00	−0.05	0.00	−0.04	0.00	0.31
ℓ	−0.02	0.00	−0.02	0.07	−0.01	0.77	0.92
Horizon	0.03	0.00	0.03	0.00	0.02	0.00	0.12
Last Round							
g	−0.03	0.00	−0.02	0.04	−0.02	0.00	0.17
ℓ	−0.01	0.00	−0.01	0.36	−0.01	0.30	0.66
Horizon	0.01	0.01	0.00	0.10	0.00	0.13	0.23
All Rounds							
g	−0.04	0.00	−0.04	0.00	−0.03	0.00	0.39
ℓ	−0.03	0.00	−0.03	0.00	−0.02	0.47	0.83
Horizon	0.04	0.00	−0.02	0.00	0.03	0.00	0.17
First Defect							
g	−0.43	0.00	−0.46	0.00	−0.32	0.06	0.37
ℓ	−0.16	0.00	−0.17	0.01	−0.06	0.80	0.78
Horizon	0.37	0.00	0.29	0.00	0.33	0.02	0.12

Notes: Additional controls include experience variables (supergame interacted with horizon) and an indicator variable for whether the player cooperated initially in the first supergame. The ME columns give the average marginal effect of each explanatory variable.

Table A10: Alternative Specifications for Table A4: Marginal Effects of Correlated Random Effects Probit Regression of the Probability of Cooperating in Round One

	RE Probit		Multiple REs		Probit		
	ME	CR-t	ME	RE-t	ME	CR-t	Bt-t
Independent Variable Specification (1)							
g	−0.10	0.00	−0.11	0.00	−0.09	0.00	0.32
ℓ	0.02	0.09	0.02	0.48	0.03	0.17	0.35
Horizon	0.03	0.00	0.04	0.00	0.03	0.00	0.14
Independent Variable Specification (2)							
g	−0.05	0.00	−0.05	0.16	−0.04	0.13	0.80
ℓ	0.04	0.00	0.05	0.10	0.04	0.01	0.16
Horizon	−0.00	0.60	−0.01	0.65	−0.01	0.48	0.48
$sizeBAD$	−0.35	0.00	−0.39	0.00	−0.36	0.00	0.10

Notes: Additional controls include experience variables (supergame interacted with each combination of stage-game and horizon parameters) and choice history variables (whether the player cooperated in the first supergame and whether the player they were matched with cooperated in the round one of the last supergame). The ME columns give the average marginal effect of each explanatory variable.

Table A11: Alternative Specifications for Table A5: Consistency of Play with Threshold Strategies

				Play Consistent With Threshold Strategy			
Experiment	Horizon	g	ℓ	1 v L Difference	CR-t	Bt-t	RE-t
DB2005	4	1.17	0.83	0.12	0.00	0.86	0.05
DB2005	4	0.83	1.17	0.10	0.00	0.34	0.06
FO2012	8	4.00	4.00	0.47	0.00	0.16	0.00
FO2012	8	2.00	4.00	0.47	0.00	0.33	0.00
FO2012	8	1.33	0.67	0.40	0.00	0.84	0.00
FO2012	8	0.67	0.67	0.40	0.00	0.33	0.00
BMR2006	10	2.33	2.33	0.39	0.00	0.50	0.00
AM1993	10	1.67	1.33	0.50	0.00	0.66	0.00
CDFR1996	10	0.44	0.78	0.20	0.00	0.34	0.08
Meta All	.	.	.	0.27	0.00	0.05	0.00
EFY	4	3.00	2.83	0.28	0.00	0.09	0.00
EFY	4	1.00	1.42	0.27	0.00	0.12	0.00
EFY	8	3.00	2.83	0.15	0.00	0.12	0.11
EFY	8	1.00	1.42	0.33	0.01	0.12	0.00
EFY All	.	.	.	0.25	0.00	0.00	0.00

Notes: The 1 v L Difference column gives the difference between the first and last supergames. Supergame refers to supergame within a set of payoff and horizon parameters. Where possible, the CR-t and Bt-t columns use standard errors clustered at the session level (for AM1993 row, standard errors are clustered at subject level since there is only one session; for the Meta All row, standard errors are clustered at the study level). In the meta study, the total number of supergames varies between 5 to 10 for sessions with $H = 2$ and $H = 4$, is 8 for sessions with $H = 8$, and is either 2 or 20 when $H = 10$. For the EFY experiments, the total number of supergames is either 20 or 30 for all parameter combinations.

A.4.2 Experiment Data

Table A12: Alternative Specifications for Table 3: Cooperation Rates: Early Supergames (1–15) vs Late Supergames (16–30)

	Round 1				Last Round			
Treatment	Diff	CR-t	Bt-t	RE-t	Diff	CR-t	Bt-t	RE-t
D4	−9.6	0.20	0.33	0.00	−0.9	0.05	0.13	0.25
D8	7.9	0.01	0.09	0.00	−3.9	0.00	0.34	0.00
E4	−3.8	0.26	0.43	0.03	−6.6	0.00	0.13	0.00
E8	8.5	0.00	0.08	0.00	−5.7	0.00	0.09	0.00
4	−6.6	0.09	0.18	0.00	−4.1	0.00	0.02	0.00
8	8.3	0.00	0.10	0.00	−4.8	0.00	0.03	0.00
All	0.6	0.67	0.86	0.23	−4.4	0.00	0.00	0.00
	All Rounds				First defect			
	Diff	CR-t	Bt-t	RE-t	Diff	CR-t	Bt-t	RE-t
D4	−6.3	0.03	0.13	0.00	−0.2	0.24	0.36	0.16
D8	−1.4	0.26	0.85	0.00	0.3	0.21	0.36	0.00
E4	−6.8	0.00	0.09	0.00	−0.2	0.02	0.12	0.01
E8	−4.9	0.01	0.31	0.00	−0.0	0.74	0.84	0.56
4	−6.7	0.00	0.02	0.00	−0.2	0.02	0.05	0.01
8	−2.9	0.02	0.45	0.00	0.2	0.40	0.62	0.05
All	−4.1	0.00	0.12	0.00	−0.0	0.78	0.97	0.55

Notes: For the cooperation measures, the regression model is a random effects probit on an indicator variable for late supergames, with standard errors clustered at the session level; for first defect, the regression model is a linear equivalent. The Diff column gives the difference in the measure between early and late supergames.

Table A13: Alternative Specifications for Table 4: Cooperation Rate for All Rounds in Supergames 1, 2, 8, 20 and 30

	SG 1 vs. SG 2				SG 1 vs. SG 8			
Treatment	Diff	CR-t	Bt-t	RE-t	Diff	CR-t	Bt-t	RE-t
D4	−10.5	0.01	0.12	0.01	−19.0	0.00	0.11	0.00
D8	0.5	0.92	0.89	0.91	−0.7	0.93	0.96	0.77
E4	1.6	0.72	0.64	0.75	2.0	0.79	0.87	0.74
E8	6.2	0.04	0.13	0.06	13.9	0.00	0.08	0.00
4	−3.8	0.25	0.35	0.14	−7.4	0.15	0.21	0.00
8	3.1	0.23	0.28	0.16	6.0	0.35	0.38	0.01
	SG 1 vs. SG 20				SG 1 vs. SG 30			
Treatment	Diff	CR-t	Bt-t	RE-t	Diff	CR-t	Bt-t	RE-t
D4	−20.0	0.01	0.07	0.00	−24.9	0.00	0.12	0.00
D8	−1.2	0.86	0.88	0.71	−3.7	0.01	0.62	0.00
E4	−8.9	0.20	0.38	0.01	−8.2	0.07	0.36	0.03
E8	3.8	0.64	0.73	0.26	4.0	0.18	0.64	0.16
4	−13.8	0.01	0.07	0.00	−15.8	0.00	0.04	0.00
8	1.1	0.84	0.85	0.62	0.0	0.59	1.00	0.26

Notes: The regression model is a random effects probit on an indicator variable for the later supergame, with standard errors clustered at the session level. The Diff column gives the difference in the all-rounds cooperation rate in supergame 1 versus the later supergame.

Table A14: Alternative Specifications for Table A8: Pair-Wise Comparison of Measures of Cooperation Across Treatments

	Round 1				Last Round			
	Diff	CR-t	Bt-t	RE-t	Diff	CR-t	Bt-t	RE-t
Supergames 1–15								
D4 vs D8	−20.2	0.05	0.10	0.01	−3.8	0.08	0.19	0.05
D4 vs E4	−20.0	0.07	0.13	0.01	−6.3	0.00	0.01	0.00
D4 vs E8	−50.6	0.00	0.01	0.00	−4.9	0.01	0.05	0.02
D8 vs E4	0.2	0.91	0.97	0.94	−2.5	0.13	0.36	0.18
D8 vs E8	−30.5	0.00	0.01	0.00	−1.1	0.62	0.68	0.69
E4 vs E8	−30.7	0.00	0.01	0.00	1.4	0.07	0.35	0.38
Supergames 16–30								
D4 vs D8	−37.7	0.01	0.02	0.00	−0.7	0.65	0.72	0.60
D4 vs E4	−25.7	0.11	0.02	0.01	−0.6	0.58	0.78	0.56
D4 vs E8	−68.7	0.00	0.01	0.00	−0.1	0.68	0.97	0.73
D8 vs E4	11.9	0.17	0.26	0.12	0.2	0.97	0.94	0.97
D8 vs E8	−31.0	0.00	0.01	0.00	0.7	0.85	0.75	0.87
E4 vs E8	−43.0	0.00	0.01	0.00	0.5	0.76	0.75	0.84
	All Rounds				First Defect			
	Diff	CR-t	Bt-t	RE-t	Diff	CR-t	Bt-t	RE-t
Supergames 1–15								
D4 vs D8	−19.2	0.01	0.03	0.00	−1.3	0.00	0.01	0.00
D4 vs E4	−12.6	0.07	0.09	0.02	−0.5	0.02	0.11	0.06
D4 vs E8	−44.7	0.00	0.01	0.00	−3.9	0.00	0.01	0.00
D8 vs E4	6.6	0.16	0.26	0.26	0.8	0.00	0.01	0.00
D8 vs E8	−25.5	0.00	0.01	0.00	−2.5	0.00	0.01	0.00
E4 vs E8	−32.1	0.00	0.01	0.00	−3.4	0.00	0.01	0.00
Supergames 16–30								
D4 vs D8	−24.1	0.00	0.02	0.00	−1.8	0.00	0.01	0.00
D4 vs E4	−12.2	0.02	0.09	0.00	−0.4	0.03	0.16	0.16
D4 vs E8	−46.2	0.00	0.02	0.00	−4.0	0.00	0.01	0.00
D8 vs E4	12.0	0.00	0.06	0.01	1.4	0.00	0.01	0.00
D8 vs E8	−22.0	0.00	0.02	0.00	−2.2	0.00	0.02	0.00
E4 vs E8	−34.0	0.00	0.01	0.00	−3.6	0.00	0.01	0.00

Notes: For all cooperation measures, the regression model is a random effects probit on a complete set of treatment dummies, with standard errors clustered at the session level; for first defect, the model is the linear equivalent. The Diff column gives the difference in the measure between the measures for the comparison treatments.

A.5. Further Details and Analysis of the Learning Model

A.5.1 Estimates

Tables A15 and A16 report summary statistics for the estimates of the learning model for each treatment. To facilitate comparison, the parameters representing initial beliefs in supergame 1 are normalized. $\bar{\beta}$ denotes $\sum_k \beta_{k0}$, as defined in the learning model. Using this, $\tilde{\beta}_k = \frac{\beta_{k0}}{\bar{\beta}}$ so that $\sum_k \tilde{\beta}_k = 1$.

E8

Variable	Mean	Std. Dev.
λ	0.83	0.86
θ	0.83	0.22
σ	0.16	0.17
κ	4.22	2.76
$\bar{\beta}$	4.45	2.92
$\tilde{\beta}_1$	0.14	0.27
$\tilde{\beta}_2$	0.03	0.07
$\tilde{\beta}_3$	0.03	0.06
$\tilde{\beta}_4$	0.06	0.11
$\tilde{\beta}_5$	0.08	0.19
$\tilde{\beta}_6$	0.04	0.06
$\tilde{\beta}_7$	0.06	0.1
$\tilde{\beta}_8$	0.1	0.12
$\tilde{\beta}_9$	0.07	0.11
$\tilde{\beta}_{10}$	0.16	0.24
$\tilde{\beta}_{11}$	0.21	0.17
ll	54.77	28.77

D8

Variable	Mean	Std. Dev.
λ	2.68	4.9
θ	0.62	0.34
σ	0.22	0.18
κ	$33.\times 10^{12}$	$233.\times 10^{12}$
$\bar{\beta}$	9.04	5.36
$\tilde{\beta}_1$	0.31	0.39
$\tilde{\beta}_2$	0.07	0.19
$\tilde{\beta}_3$	0.02	0.09
$\tilde{\beta}_4$	0.02	0.05
$\tilde{\beta}_5$	0	0.01
$\tilde{\beta}_6$	0.02	0.04
$\tilde{\beta}_7$	0.03	0.06
$\tilde{\beta}_8$	0.05	0.14
$\tilde{\beta}_9$	0.09	0.18
$\tilde{\beta}_{10}$	0.12	0.14
$\tilde{\beta}_{11}$	0.26	0.3
ll	91.12	43.92

Table A15: Summary statistics for long horizon treatments

D4

Variable	Mean	Std. Dev.
λ	2.2	4.83
θ	0.71	0.32
σ	0.16	0.18
κ	1.12	7.88
$\bar{\beta}$	$4.\times 10^{14}$	3.24
$\tilde{\beta}_1$	0.32	0.35
$\tilde{\beta}_2$	0.07	0.19
$\tilde{\beta}_3$	0.05	0.16
$\tilde{\beta}_4$	0.05	0.09
$\tilde{\beta}_5$	0.05	0.08
$\tilde{\beta}_6$	0.12	0.19
$\tilde{\beta}_7$	0.34	0.32
ll	29.39	19.02

E4

Variable	Mean	Std. Dev.
λ	8.17	11.9
θ	0.46	0.31
σ	0.22	0.19
κ	$31.\times 10^{12}$	$17.\times 10^{13}$
$\bar{\beta}$	1.98	1.01
$\tilde{\beta}_1$	0.15	0.32
$\tilde{\beta}_2$	0.18	0.33
$\tilde{\beta}_3$	0.02	0.04
$\tilde{\beta}_4$	0.09	0.19
$\tilde{\beta}_5$	0.11	0.25
$\tilde{\beta}_6$	0.16	0.31
$\tilde{\beta}_7$	0.29	0.37
ll	27.64	18.76

Table A16: Summary statistics for short horizon treatments

A.5.2 Figures

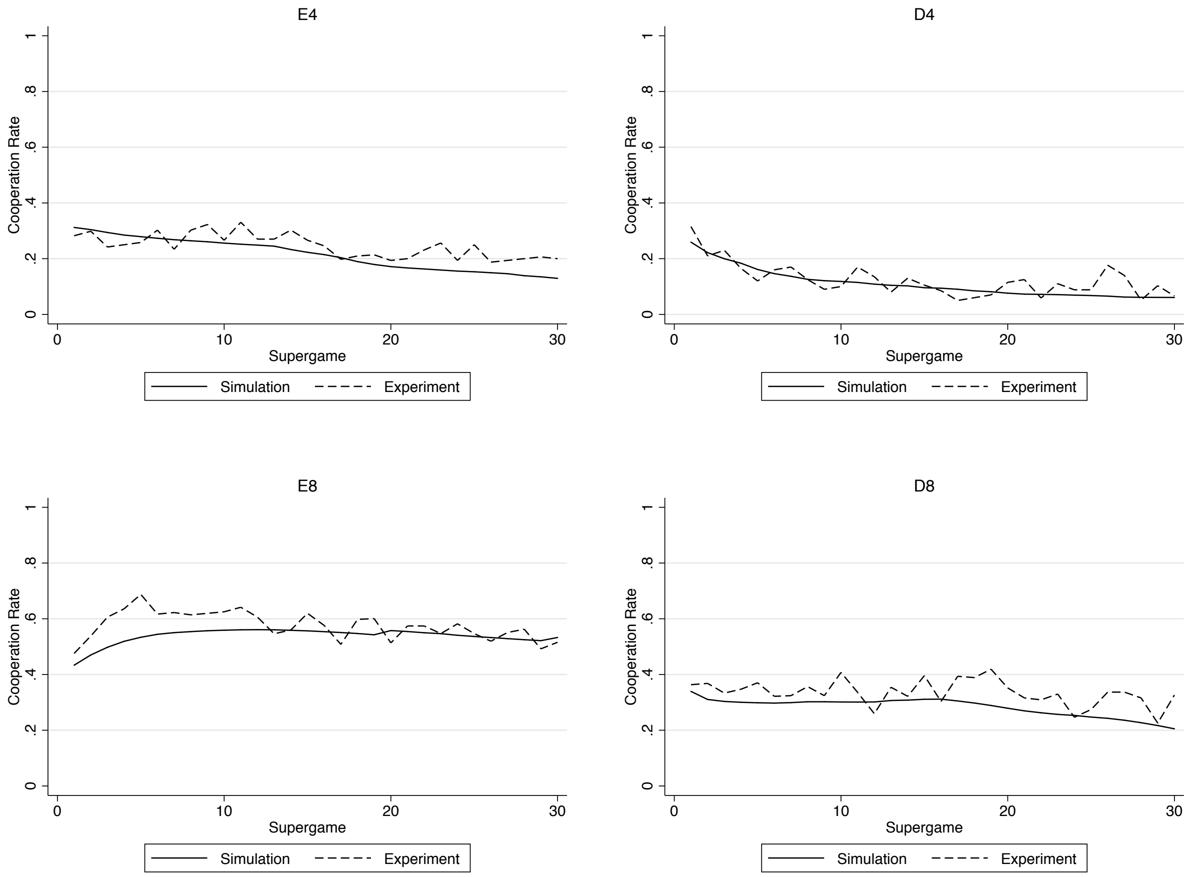

Figure A7: Average Cooperation: Simulation Versus Experimental data

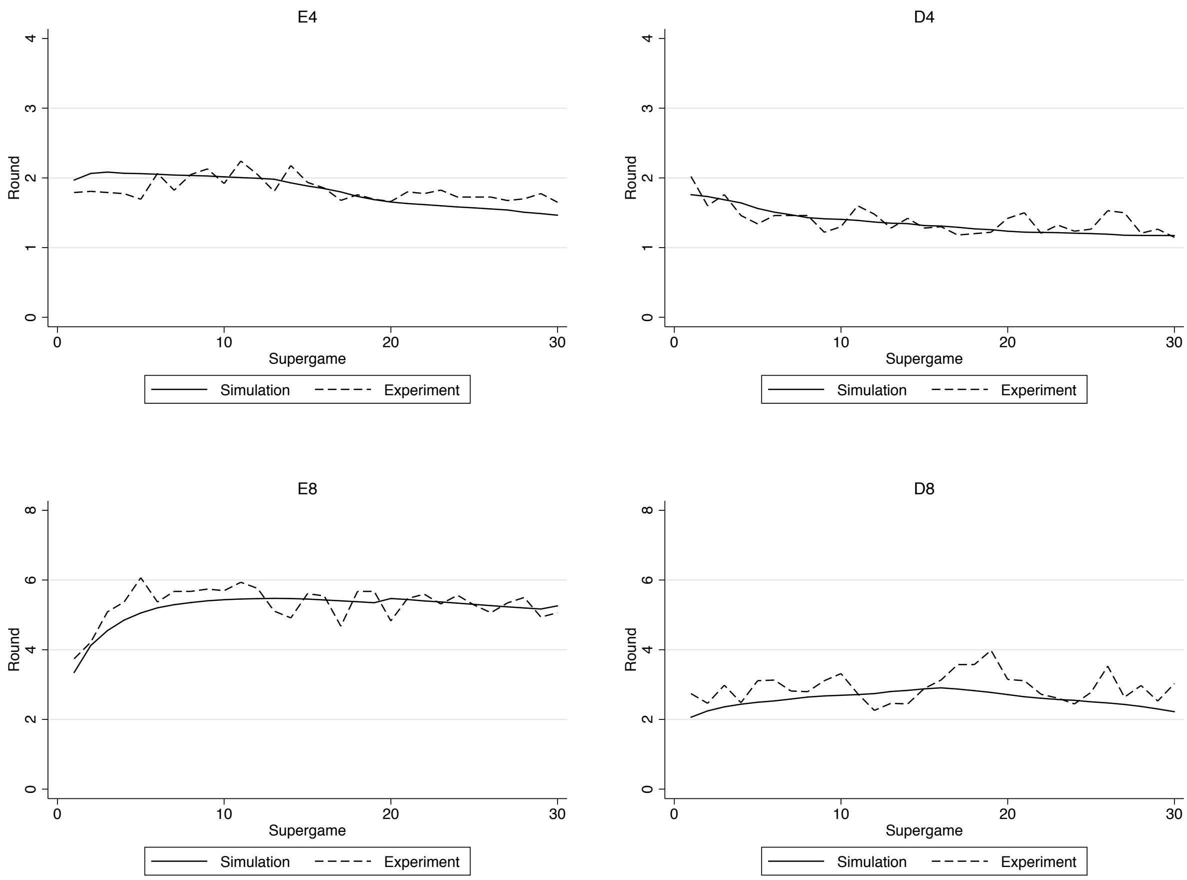

Figure A8: Mean Round to First Defection by Supergame: Simulation versus Experimental Data

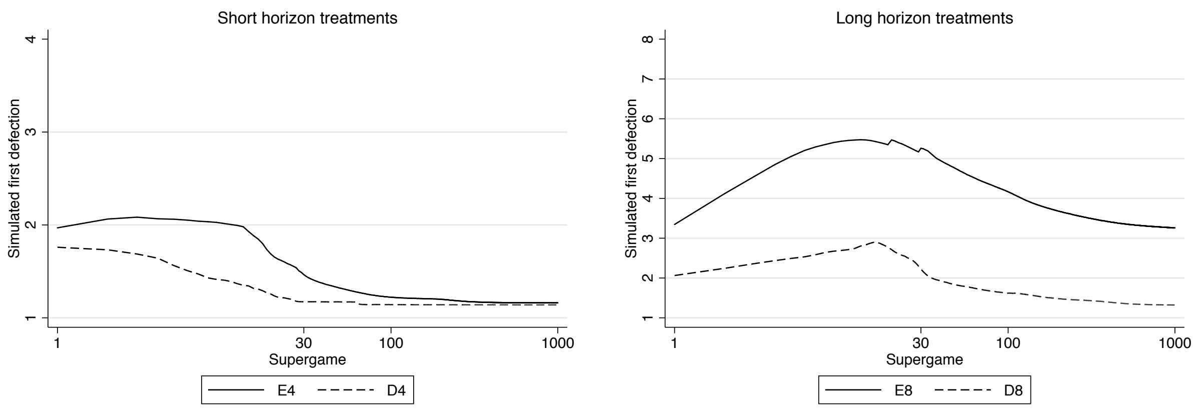

Figure A9: Long Term Evolution of Mean Round to First Defection by Supergame

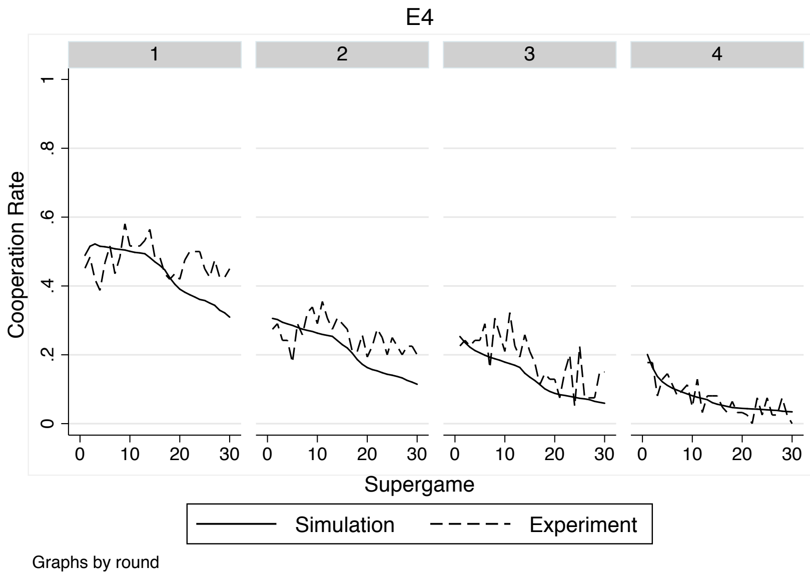

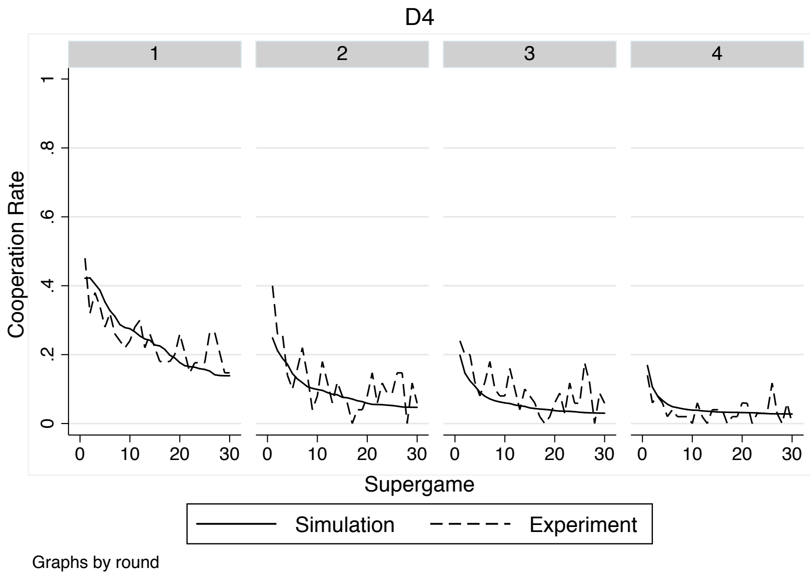

Figure A10: Average Cooperation Rate by Supergame: Simulation versus Experimental Data for Each Round in the Short Horizon Treatments

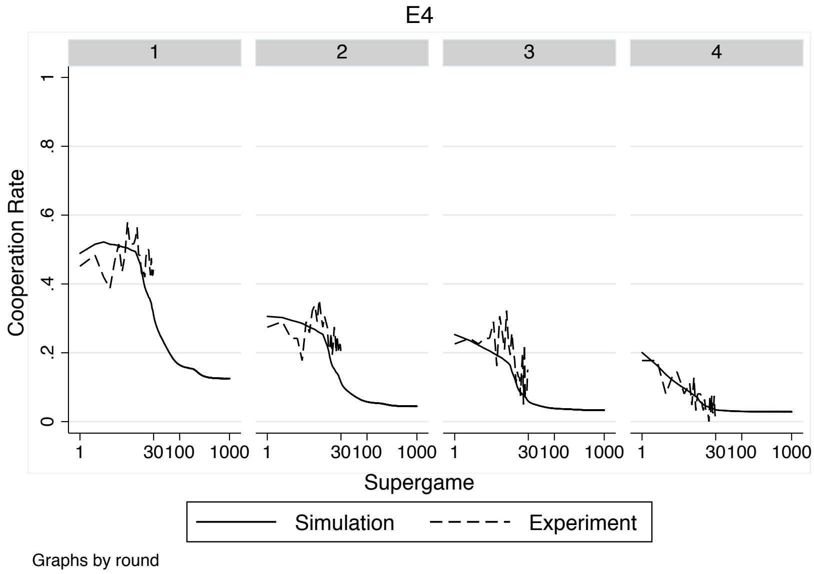

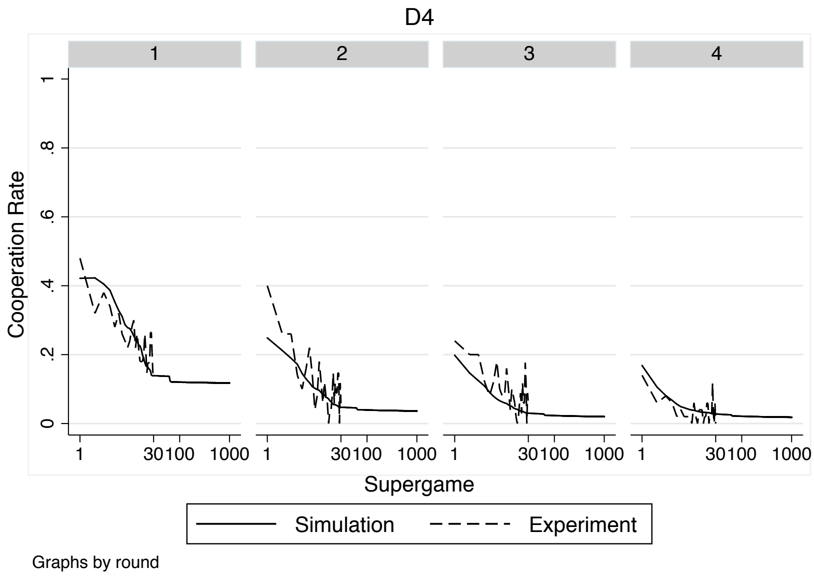

Figure A11: Long Term Evolution of Cooperation Rate for Each Round of the Short Horizon Treatments

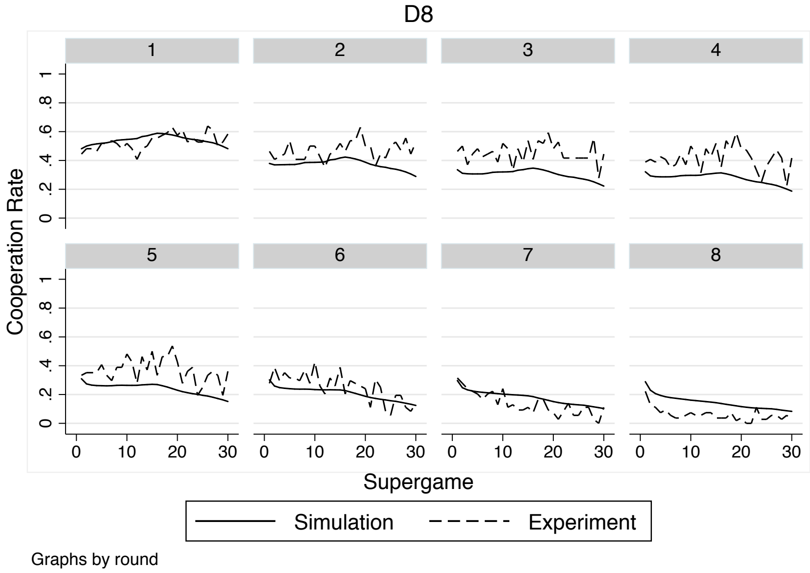

Figure A12: Average Cooperation Rate by Supergame: Simulation versus Experimental Data for Each Round in D8

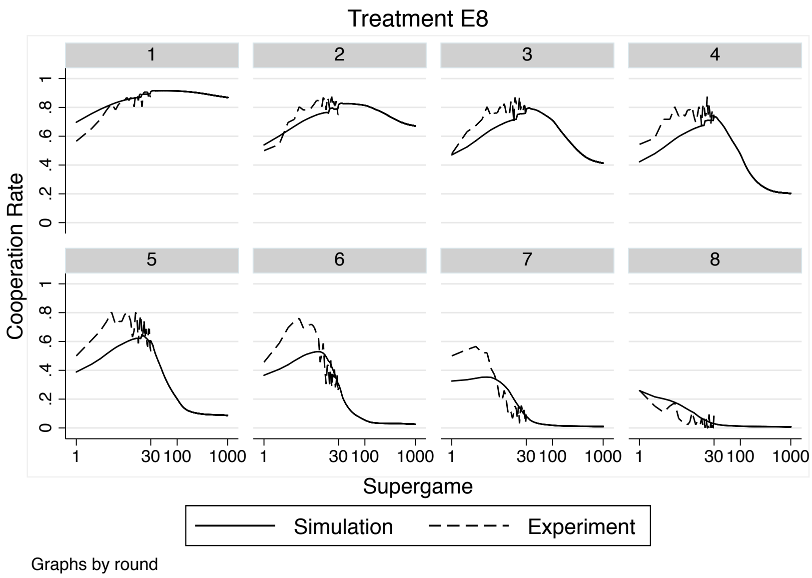

Figure A13: Long Term Evolution of Aggregate cooperation For Each Round In E8

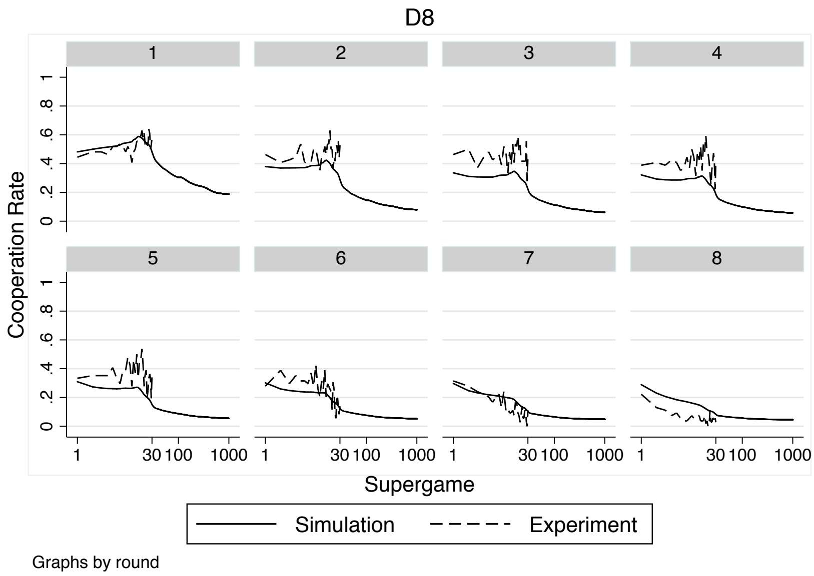

Figure A14: Long Term Evolution of Aggregate cooperation For Each Round In D8

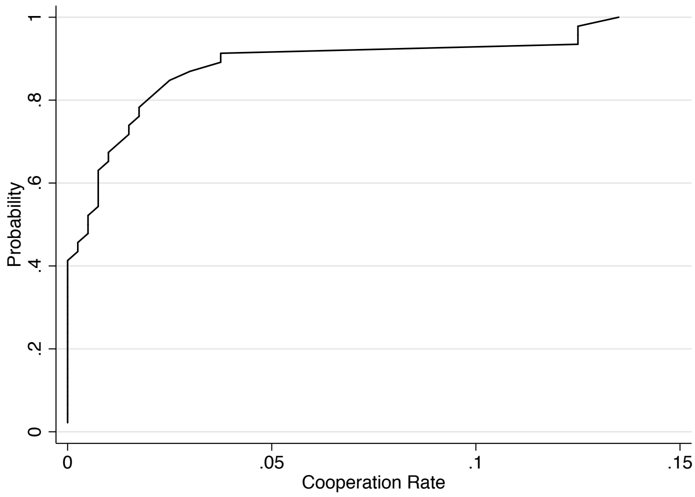

Figure A15: Cumulative Distribution of Cooperation Against an AD type In E8 (Supergames 250-300)

Each subject is simulated to play against an AD type–someone who defects in all rounds of a supergame regardless of past experience–for 300 supergames. The average cooperation rate for the subject from supergames 250-300 is taken as a measure of that subject's cooperativeness. Such a measure of cooperativeness combines the effects of the parameters estimated in the model in an intuitive way. It effectively captures how well a subject is able to learn to defect against a defector.[75,76]

Figure A15 plots the cumulative distribution of simulated cooperation rates after 250 supergames against a player who is following the AD strategy. The distribution has a mass point around 0 implying that about 40% of the subjects learn to defect perfectly with sufficient experience in this environment. There is limited but positive levels of cooperation for the remaining subjects. Note that this corresponds to subjects making cooperative choices after observing their partners defecting in every single round of 250 supergames; hence, this suggests the existence of cooperative types. The model allows for multiple kinds of cooperative types: some forces that can drive cooperative actions in such an extreme environment are strong priors, limited learning from past experiences, and noise in strategy choice and implementation.

[75] An horizon of 250-300 is chosen to correspond to the time frame we are analyzing in what follows, but the exercise can easily be repeated for a different range of supergames. Looking at cooperation rates in supergames 900-1000 gives very similar results.

[76] Focusing on cooperation in later supergames also dampens the effect of a strong prior and execution noise in early supergames. This exercise can be repeated by constructing a measure of cooperativeness by focusing on behavior in early supergames. As expected, removing subjects based on such a measure has a bigger impact on cooperation in earlier supergames, but the effect quickly disappears with experience.

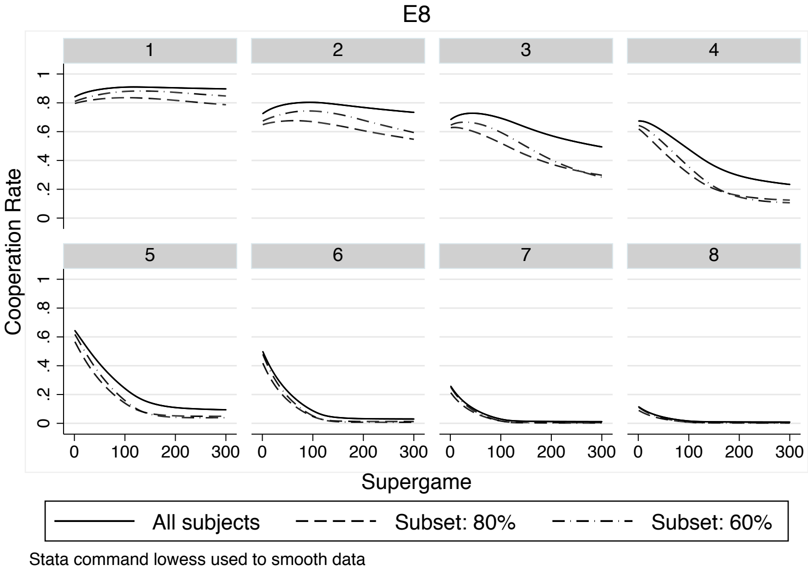

Figure A16: Long term Evolution of Aggregate Cooperation For Each Round In E8 By Subset

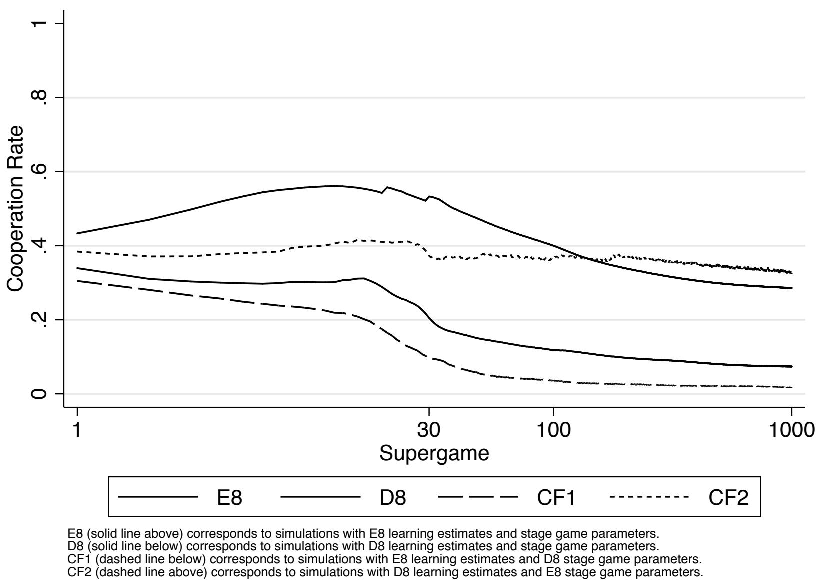

Figure A17: Long term Evolution of Aggregate Cooperation

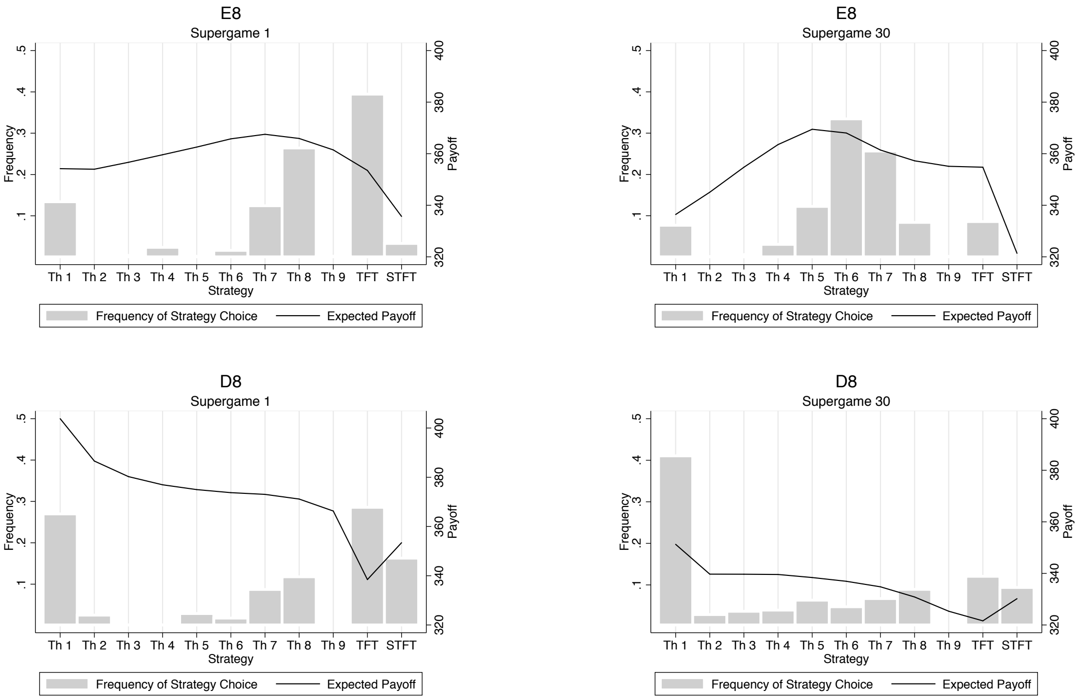

Figure A18: Frequency and Expected Payoff of Each Strategy

These values are estimated by simulating behavior in 1000 sessions composed of 14 randomly drawn subjects. The frequency of choice for each strategy is recorded, along with how well each strategy performs when played against each subject of the session.

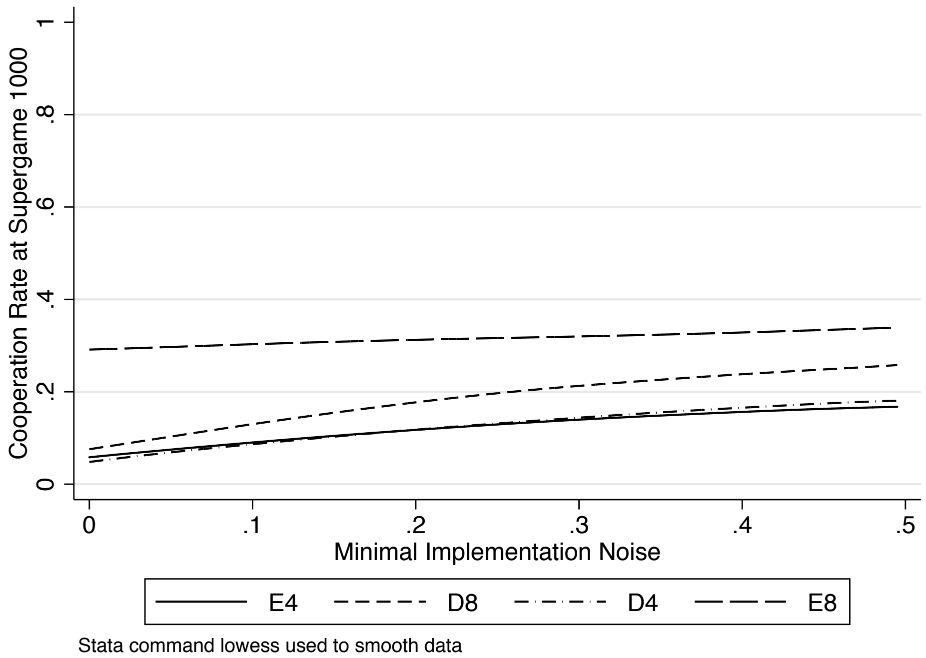

Figure A19: Effects of Constraining the Decline (with Experience) of Implementation Noise

One additional concern may be the robustness of the results to specific parameters. In particular, one may wonder what happens in the long run, if implementation error is not allowed to completely disappear with experience. To explore this possibility, we conduct additional simulations constraining how much the implementation error can decline as as a result of learning (through the κ parameter). Formally, if the constraint is set to σ_{min}, implementation noise in supergame t for subject i is calculated to be $\max\{\min(\sigma_{min}, \sigma_i), \sigma_i^{t^{\kappa_i}}\}$. According to this specification, σ_{min} does not constrain initial implementation noise σ_i, but limits how much it can decline over time with experience through the κ parameter. We recover our original simulation results when the constraint is never binding (set to 0), and we see what long term cooperation results would be like if the implementation noise never changed (corresponding to the case where the constraint is set to 0.5, which is equivalent to setting $\kappa_i = 0$ for all subjects).

The results show this constraint to have little effect on long term cooperation rates in the E8 treatment. In the case of treatments E4 and D4, looking directly at the experimental data reveals over 95% of play to be consistent with threshold strategies by supergame 30. Thus, persistent implementation error seems less of a concern in these treatments. D8 is the treatment where persistent implementation noise has the most effect. In this treatment, cooperation rates in the 1000th supergame are significantly affected by the constraint (although still remain below 30%) and our experimental data cannot inform us of the extent to which implementation errors may persist.

A.6. Sample Instructions: D8 Treatment

Welcome

You are about to participate in an experiment on decision-making. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. Please turn off cell phones and similar devices now. Please do not talk or in any way try to communicate with other participants.

We will start with a brief instruction period in which you will be given a description of the main features of the experiment. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear.

General Instructions

1. You will be asked to make decisions in several rounds. You will be randomly paired with another person in the room for a sequence of rounds. Each sequence of rounds is referred to as a *match*.

2. Each match will last for **8** rounds.

3. Once a match ends, you will be randomly paired with someone for a new match. You will not be able to identify who you've interacted with in previous or future matches.

Description of a Match

4. The choices and the payoffs in each round of a match are as follows:

	1	2
1	$51, 51$	$5, 87$
2	$87, 5$	$39, 39$

The first entry in each cell represents your payoff for that round, while the second entry represents the payoff of the person you are matched with.

(a) The table shows the payoffs associated with each combination of your choice and choice of the person you are paired with.

(b) That is, in each round of a match, if:

- $(1, 1)$: You select 1 and the other selects 1, you each make 51.
- $(1, 2)$: You select 1 and the other selects 2, you make 5 while the other makes 87.
- $(2, 1)$: You select 2 and the other selects 1, you make 87 while the other makes 5.
- $(2, 2)$: You select 2 and the other selects 2, you each make 39.

To make a choice, click on one of the rows on the table. Once a row is selected, it will change color and a red submit button will appear. Your choice will be finalized once you click on the submit button.

Once you and the person you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear.

End of the Session

5. The experiment will end after 30 matches have been played.

6. Total payoffs for each match will be the sum of payoffs obtained from each round of that match. Total payoffs for the experiment will be the sum of payoffs for all matches played. Your total payoffs will be converted to dollars at the rate of 0.003$ for every point earned.

Are there any questions?

Before we start, let me remind you that:

- Each match will last for 8 rounds. Payoffs in each round of a match, as given in the table above, depend on your choice and the choice of the person youre paired with.
- After a match is finished, you will be randomly paired with someone for a new match.