LUCIANA MOSCOSO BOEDO, LUCIA QUESADA AND MARCELA TARAZONA

Cooperation among Strangers in the Presence of Defectors: An Experimental Study

CIDE

www.cide.edu

Abstract

Does the rotten apple spoils his companions? This is the question we analyze in the context of a repeated population game with behavioral types. Our experimental results show that the inclusion of a non-cooperative player in an anonymous community makes cooperation much more difficult to sustain but that individuals still manage to trust some of the permanent players of society. The rotten apple lowers the quality of the companions, but is not able to completely spoil them.

Keywords: *Population games, anonymous random matching, social norms.*

Resumen

¿Arruina la manzana podrida a todo el cajón? Esta es la pregunta que analizamos en el contexto de un juego poblacional repetido con tipos conductuales. Nuestros resultados experimentales muestras cómo la inclusión de un jugador no cooperativo en una comunidad anónima complica el mantenimiento de la cooperación. También mostramos que los individuos de todas formas logran confiar en los miembros permanentes de la sociedad. La manzana podrida baja la calidad del cajón, pero no logra arruinarlo por completo.

Palabras Clave: *Juegos poblacionales, normas sociales y emparejamientos aleatorios anónimos.*

Cooperation among Strangers in the Presence of Defectors:

An Experimental Study*

Luciana Moscoso Boedo[†]   Lucía Quesada[‡]   Marcela Tarazona[§]

October 25th, 2013

Abstract

Does the rotten apple spoils his companions? This is the question we analyze in the context of a repeated population game with behavioral types. Our experimental results show that the inclusion of a non-cooperative player in an anonymous community makes cooperation much more difficult to sustain but that individuals still manage to trust some of the permanent players of society. The rotten apple lowers the quality of the companions, but is not able to completely spoil them.

1 Introduction

Cooperation among strangers in a community might be difficult to sustain. If we add to a community a rotten apple, i.e. a short-run player that always plays opportunistically, can we expect cooperation to arise among the permanent members of the society? This is the question we analyze in the context of a repeated population game with behavioral types. We are interested in knowing whether permanent members of a society are able to overcome the noise introduced by short-run players and sustain a cooperative outcome. We find that even though cooperation is much more complicated to sustain, long-run players are able to build a society in which they trust each-other. We analyze a community composed by heterogeneous players and look at the sustainability of cooperation using decentralized informal mechanisms. Internet transactions, tips and some leasing-borrowing activities are some of the situations in which this result can be relevant. We propose a robustness check to the well known result that anonymous communities are able to cooperate even if they only have private information.

In order to answer our question, we design an experiment to test cooperation in heterogeneous communities. In our experiment, each subject belongs to a community of four players and in each period they are matched anonymously to play a prisoner's dilemma game. Each community plays an infinitely repeated game. We are interested in analyzing the average behavior of the population and in figuring out whether the short-run players generate a disruption in cooperation that long-runs are not able to overcome.

We analyze player's individual behavior in a repeated game and check how this behavior changes in a setting where information is even more imperfect. Our results show that different strategies coexist in a given community and that the simple one-shot game Nash equilibrium strategy is more attractive in a noisier context. We also show that the contagious (grim) strategy is used less often when it is not an equilibrium strategy.

Our results add to an active literature that studies the sustainability of cooperation in populations both theoretically and experimentally. Our data suggest that people do worse in a nosier environment but that they are nevertheless able to reach an outcome, that even if it is not efficient, it is far away from the one shot game inefficient outcome.

The paper is organized as follows: Section 2 discusses the related literature; Section 3 presents the experiment, Section 4 presents the results and Section 5 concludes.

2   Related Literature

2.1   Theoretical Results

The Folk Theorem for repeated games states that when two individuals play an infinitely repeated version of the prisoner's dilemma, if they are sufficiently patient, there is an equilibrium in which the cooperative action is taken every period by both individuals.[1] In this literature an important assumption is that each player always interacts with the same opponent and hence is able to observe his opponent's complete history of moves.

Kandori (1992) extends this result to a community of M individuals in which individuals are randomly matched in pairs at the beginning of each period to play the prisoner's dilemma. The big difference is that individuals do not know the identity of the opponent in each round of the game, nor they know his history of moves. Individuals only observe the outcomes of the games in which they were involved. However, cooperation can be sustained through a simple symmetric strategy called the contagion strategy. According to this strategy, players start cooperating and they continue to cooperate until they observe a defection. An observed defection triggers defection forever. The seminal paper by Ellison (1994) generalizes this result to populations of any size by allowing punishment length to depend on a public randomization device.

Moscoso Boedo (2010) changes Kandori's set up by introducing short-run players in the population. Each period, this short-runs are replaced by others with the same discount factor. She shows that in this context the contagious strategy proposed by Kandori is not an equilibrium strategy and that, furthermore, in this context there is no symmetric pure strategy with infinite punishments that can sustain cooperation.

There are many recent contributions to the literature on repeated population games with anonymous random matching. Among those, Deb (2012) proves a Folk Theorem by allowing for pre-play communication and Dal Bó (2007) provides a Folk Theorem when allowing for some information transmission among players. Other have extended the results beyond the prisoners' dilemma such as Takahashi (2010), who proposed belief free equilibria in a context when first order information is available and Deb and González-Díaz (2011), who proposed a strategy with a stage of trust building. All of these proposals are not robust to the inclusion

---

[1]See, Fudenberg and Maskin (1986) or Fudenberg, Kreps, and Maskin (1990) for reference.

of a short-run player. Dilmé (2012) proposes affine strategies that are robust to the addition of behavioral types. He finds belief free equilibria supported by memory one strategies.

2.2   Experiments on Repeated Games

Our work adds to the experimental literature on infinitely repeated games and specifically to that of repeated randomly matched population games. Dal Bó (2005) stablished that the incentives in an infinitely repeated games are different that those of a finitely repeated game, reporting higher cooperation rates in the indefinite duration game. In that paper there is a repeated interaction with a fixed partner. Duffy and Ochs (2009) study the effect of random matching in a population game. They found that cooperation is higher in fixed matching economies. They found that a cooperative norm does not emerge and that only with experience are players able to increase cooperation in the population (of 14 and 6 players respectively).

The closest work to ours is Camera and Casari (2009). Our baseline treatment is a replication of the private monitoring treatment of their study. They show that in an economy of four players that meet anonymously and interact privately, players are able to cooperate and that, cooperation increases for later cycles (that is, in each session of the experiment they form five supergames where a different set of four players are randomly matched with each other, each cycle is a supergame). In their paper they are mostly interested in analyzing how monitoring and punishment institutions are able to promote cooperation. For the treatment that we replicate they found that the behavior of the representative agent corresponds to a reactive strategy. In a later paper, Camera, Casari and Bigoni (2012) study the individual strategies that are followed in the private monitoring treatment of the previous paper, and show that there is a wide heterogeneity of strategies that describe individual play. They also find that individuals do not play the grim-trigger strategy even when the representative player does. We follow their estimation procedure for individual strategy to figure out how individual behavior change in the presence of short-run players.

The importance of experience in repeated game was stablished by Dal Bó and Fréchette (2011a). In their paper they analyze how cooperation evolves with experience and show that, cooperation may not prevail even when it is an equilibrium outcome. They study a repeated prisoners' dilemma game with fixed pairings and provide a novel method to estimate

4

strategies that allows player to sometimes make mistakes when choosing the action. We used their estimation procedure in this paper. In a different paper, Dal Bó and Fréchette (2011b) elicit strategies in a infinitely repeated prisoners' dilemma experiment. They find that subjects use some common strategies such as Grim trigger and Tit-for-Tat but that other with desirable properties such as Win-Stay-Lose-Shift are not prevalent. Fudenberg, Rand and Dreber (2012) add noise to the environment of Dal Bó and Fréchette (2011a) and extend the set of strategies to more than memory one. They found that when adding noise to a two-player infinitely repeated prisoners' dilemma makes players more lenient and more forgiving, thus allowing them to support cooperative outcomes.

3   The Experiment

3.1   Experimental Design

The experiment was designed to analyze cooperation in a group where players play an indefinitely repeated prisoner's dilemma under anonymous random matching. We are interested in the effect of introducing a short-run player in each group. Our baseline is based on the design for the private monitoring treatment in Camera and Casari (2009) . Subjects interacted anonymously through computer terminals. Each subject was assigned to a group of four players. Subjects never knew the composition of their group. In each session, we ran 5 cycles (supergames). Each cycle consisted of a random number of periods. In each period, one subject was randomly matched with another player of the same group to play the prisoner's dilemma in Table 1. The identity of the partner was never revealed. We set the probability of continuation at $\delta = 0.95$ and this was announced to the participants at the beginning of the session. Once a cycle is over, participants are reassigned to new groups of four players. The composition of groups was made in order to minimize the possibility of contagion from one cycle to another. We had 16 players in each session. Hence, in each cycle we had four groups of 4 players.

Table 1

Even Player

		Y	Z
Odd Player	Y	25, 25	5, 30
	Z	30, 5	10, 10

We ran two different treatments. In the first treatment (baseline) all 16 players were identical in the sense that they faced the same payoff matrix and the same probability of continuation. In the second treatment, players were heterogeneous in the following sense. In each group, one player was a short-run player, who plays then non-cooperative strategy Z every period. In each session with treatment, there were four players of this type. The behavior of these players was simulated by the program. The 12 remaining players were informed about these short-run players. They also knew there was one in each group. However, in a given period they were never informed about the identity of their opponent or if they were facing a short-run or a long-run player.

The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). Our participant were undergraduate students from various disciplines at Universidad Torcuato Di Tella (UTDT) in Buenos Aires. The 96 subjects were recruited through e-mail and in-class-announcements.

All payoffs during the game were expressed in points. At the end of the session, a cycle was randomly chosen and the total number of points of that cycle was converted to Argentinean pesos (AR$) at a rate of 10 points = AR$ 1. Subjects were also paid an AR$10 show-up fee.[2]

The summary of the sessions is in Table 1 Appendix B. There were 16 players in each session. In the baseline those players were 16 recruited students, for the treatment we recruited 12 students and simulated 4 short-run players with the computer. Each session lasted on average 2 hours including instructions reading, a quiz and a survey. We planned to have 48 subjects in each treatment, and for that run 3 sessions of the baseline and 4 sessions of the treatment. We run three mock sessions to test the program. The subjects that participated in those sessions did not participate in any of the studied sessions. In session 9 of the baseline, we were able to run only three cycles of the game within the two hours (cycles were unusually long: cycle 1 was 70 periods long, cycle 2 was 105 periods long and cycle 3 was 10 periods long). We replaced session 9 with session 11 of the baseline, in order to have a session with

---

[2]The exchange rate at the moment of the sessions was about AR$4 for each US dollar.

all 5 cycles. Average earnings were AR\$ 43.4 (41.92 in baseline, 44.88 in treatment) with a maximum of AR\$ 84 and a minimum of AR\$ 20. The average session lasted 25.3 periods (19.1 periods in baseline and 29.9 periods in treatment).

3.2 Theoretical Predictions

The theoretical predictions are based on the works in Kandori (1992), Ellison (1994) and Moscoso Boedo (2010). Our subjects play the Prisoner's dilemma presented in Table 1 above. For this game, (D, D) is the unique one-shot dominant strategy equilibrium and (C, C) Pareto dominates (D, D). The work by Fudenberg and Maskin (1986) demonstrates a folk theorem for two player games in which cooperation can be sustained if the game is repeated an indefinite number of times and the players are sufficiently patient and there are multiple other equilibria. For the case of anonymous random matching where agents interact with different partners each time and do not observe the opponent's identity, Kandori (1992) and Ellison (1994) have shown that the cooperative outcome can be sustained by a contagious strategy (with or without public randomization device respectively). In our baseline treatment, previous experimental studies have shown Camera and Casari (2009) that cooperation increases through cycles, that the average behavior satisfies a contagion process and that, there is a coexistence of heterogeneous individual strategies (this is studied more specifically by Camera, Casari and Bigoni (2012)). Our treatment involves an heterogeneous population, which adds some behavioral players to the community, short-run players. With this kind of heterogeneity, Moscoso Boedo (2010) shows that the contagious strategy is not an equilibrium strategy and that there is no pure strategy equilibrium with infinite punishment that can support cooperation along the equilibrium path (See Appendix A for the proof for the game of the experiment). Given these cited theoretical results, we expect to have the following results in our experiment:

- More cooperation in the baseline than in the treatment. This follows from the fact that the treatment case corresponds to a sharper imperfect monitoring environment (a defection corresponds not only to a personal deviation and social punishments, but also it can be due to the fact that you met a short-run player). We know that there is no simple pure strategy that supports cooperation along the equilibrium path in the treatment, and thus we expect players to perform worse in the noisier environment.

- If the grimm strategy is used, cooperation should decrease during a cycle. We expect to see the strategy "always defect" used more often in the treatment than in the baseline sessions, and the opposite with "always cooperate" (notice that as the contagious strategy is an equilibrium strategy in our baseline, when there are no deviations, this is observationally equivalent to the strategy always cooperate). When we try to classify observed strategy by memory-one strategies, we expect to see classify strategies in the baseline performing better than in the treatment, given that, except for the strategy "always defect" the simple memory-one strategies are not equilibrium strategies in the treatment, so the payoffs for following a non-equilibrium plan should be low. We expect the opposite for the baseline.

3.3   Estimation Procedure for Individual Strategies

Maximum Likelihood Estimation

In our first estimation, to assess the prevalence of each strategy in our data we follow Dal Bó and Fréchette (2011b) and assume that each subject chooses a fixed strategy at the beginning of each cycle. This is to make our results comparable to them and to Camera, Casari and Bigoni (2012) and because after each cycle players are assigned to a new group with different partners. We allow players to make errors in actions, that is with some probability in every period they may choose an action that is not described by the strategy. Specifically, if subject i chooses strategy s, her chosen action in round r of group k is C if $s_{ikr}(s) + \gamma \varepsilon_{ikr} \geq 0$ where $s_{ikr}(s) = 1$ if strategy s ask players to cooperate in round r of group k given the history to that point and $s_{ikr}(s) = -1$ if strategy s asks the player to defect. The error ε_{ikr} is assumed to be independent across subjects, rounds, groups and histories. γ parameterizes the probability of errors and the density of the error is such that the overall likelihood that subject i uses strategy s is

$$p_i(s) = \prod_k \prod_r \left(\frac{1}{1 + exp(-s_{ikr}(s)/\gamma)} \right)^{y_{ikr}} \left(\frac{1}{1 + exp(s_{ikr}(s)/\gamma)} \right)^{1 - y_{ikr}}$$

where y_{ikr} is 1 is the player chooses C and 0 it it chooses D. We use maximum likelihood estimation to estimate the prevalence of the memory 1 strategies (strategies that only depend on previous period own and opponents actions) and bootstrapping for the associated standard errors.

8

Two-State Automata with transitions errors

Following the methodology of Camera, Casari and Bigoni (2012) and of Dal Bó and Fréchette (2011b) we coded each of the memory one strategy by a binary number of 5 digits, where a 1 reflects cooperation, C, and a 0 reflects a defection, D. The first digit reflects the action to be taken in the first period, then the action that needs to be taken in period t if the outcome of $t-1$ was: (C,C) for the second digit, (C,D) for the third digit, (D,C) for the fourth digit and (D,D) for the fifth digit. Thus, for any previous period play, the strategy determined next period play. For instance the strategy tit for tat corresponds to 11010. Given this classification of strategies we take each sequences of actions by a subject in a cycle as an observation. We have 240 observation for each treatment. For each observation we have the realized string of actions by the opponent and are able to simulate the behavior prescribed by each of the 32 memory one strategies. After that we assign a score to each strategy that depends on how it compares to actual behavior (we assign the number 1 to each action that matches the prescribed action by a strategy and then calculate the average over those numbers) and select the strategies with the highest score. We consider that a strategy classifies an observed string of actions if the score is one. We also take into account the possibility of mistakes in transitions: players do not transition correctly from period t to period $t+1$ by a strategy. If this probability of mistakes is sufficiently high (when it reaches 35% we are able to classify all observations). We will be interested in low probabilities of mistakes (5%).

4   Results

Average Cooperation on the Population

We first present results regarding the average cooperation on the population. Our baseline replicates Camera and Casari (2009) but we do not get increasing cooperation in later cycles as they do, instead we replicate Duffy and Ochs (2009) that get declining trend across supergames (see Figure 1 and 2 in Appendix B). We do get this increasing cooperation across cycles in our treatment with short-run players. In our data the coordination in cooperation was remarkably high. When we calculate average cooperation considering each cycle in each session a unit of observation, i.e. we take into account that each cycle in each session has a different duration, and we give each of them equal weight, we get that average cooperation in

the baseline was 75.8% and in the treatment was 61.6% which is significantly lower at the 5% level. If we just calculate the average of the variable choice among all observations (without considering that some cycles lasted longer than the others) we get that average cooperation was 40.31 in the baseline and 23.86 in the treatment, which is line with the result that Duffy and Ochs (2009) for the random matching game (they get an average cooperation of 42.9%). (See Tables 1 and 2 in Appendix B). Both methods of calculating average cooperation support our prediction regarding the difficulty of sustaining cooperation in our treatment. The fact that the average of cooperation in the treatment is significantly different from zero in all cycles and all periods, implies that behavioral types are not completely eroding trust among permanent members of the society. Our results show that when we evaluate average cooperation as Camera and Casari (2009) we get a significant higher cooperation rate in the first cycle of the baseline compared to the treatment but those differences are not significant in later cycles. We get a differential response of average cooperation to the first defection, as shown in the table below:

Table 2

	Baseline	*Treatment*
Before 1st D	27.77%	48.39%
After 1st D	3.17%	1.79%

This differential response to a defection is consistent with a contagious strategy, which is not an equilibrium strategy in the treatment. Notice that by the nature of our treatment, the proportion of players who have observed a defection is above 95% in period 4 while in the baseline this happens in period 8. It is worth noting that the proportion reaches 1 in period 8 for treatment but in period 28 for the baseline. See Figure 3 in Appendix B. Analyzing the probability of cooperation conditional on own previous actions and opponent previous actions, we find that both in the baseline and in the treatment, the best predictor of a player period t action is the player's period $(t-1)$ action, and that it is more probable for players to cooperate after they have defected in the baseline than in the treatment.

Table 3

	Baseline	*Treatment*
Prob. C after own C	71%	72%
Prob. C after own D	19%	12%

The difference in the probability of cooperating after defecting is significant and we interpret this as players understanding that in the more noisy environment, trust is more likely to be broken once they defect(given that, 50% of the population is defecting at a given period). See Figure 4 in Appendix B. Experience does not seem to give rise to cooperation even when it is an equilibrium result, as Dal Bó and Fréchette (2011a) present it. Our results show that cooperation is decreasing over time for both baseline and treatment (see figure 5 in Appendix B).

Eliciting Individual Strategies

We have also analyzed the individual strategies players use. For this we estimated the importance of each candidate strategy with a maximum likelihood approach as in Dal Bó and Fréchette (2011b) and the section 7 of Camera, Casari and Bigoni (2012). In our setting we estimate the importance of each of a set of 26 candidate strategies. In this approach every individual is classified, and the prevalence of errors in actions is estimated; i.e. players might with some probability choose the wrong action relative to the strategy. We allow subjects to change strategies from cycle to cycle, to make our results comparable to those of Camera, Casari and Bigoni (2012). Table 8 in Appendix B shows the results. In the baseline only two strategies are significant and those are the unconditional strategies "Always Defect" and "Always Cooperate" with a predominance of the former. This is in line with Camera, Casari and Bigoni (2012) but in our settings the importance of those two strategies is much higher (more than 70% of individuals are represented by these. It is interesting to notice that in the treatment, the strategy "Always Defect" is the most likely strategy but a forgiving strategy such as Tit For Tat is also significant, these being consistent with after experiments replays by subjects saying that they had to forgive the computer for misbehaving. We expanded the set of strategies to include some longer memory cooperative strategies that are more lenient or forgiving (as presented by Fudenberg, Rand and Dreber (2012)). None of those strategies resulted significant. We also classified strategies for individuals per cycle following the strategy of Camera, Casari and Bigoni (2012). These results are shown in table 4 of Appendix B. In this approach individuals are allowed to make mistakes in transitions given a strategy, and if we increase the probability of mistakes we increase the percentage of the population that we can classify using this strategy (for our experiment, if the probability of mistakes in transition is of the .35% then we classify all the observations, see figure 6 in Appendix B). Our results show that with a 5% probability of mistakes we

can classify 70% of the observed strategies both in the baseline and the treatment. The equilibrium strategy "Always Defect" classifies more individuals in the treatment than in the baseline (37.9% vs 28.3%). It is interested to look at payoffs associated with individual strategies. Of course, given the presence of the behavioral type, average payoffs are lower in the treatment, but, interestingly average payoffs are higher for classified individuals in the baseline (where many of the presented simple strategies are equilibrium strategies) but for unclassified individuals in the treatment (where the only simple strategy that is an equilibrium strategy is "Always Defect"). This is in line with the findings of Camera, Casari and Bigoni (2012) and with our prediction that in the noisy environment, equilibrium strategies are much more complicated. The Grim Strategy is much more frequent in the baseline (17.9%) than in the treatment (3.3%), showing that players do not play non- equilibrium strategies as often, as proven in a two-player context by Dal Bó and Fréchette (2011b). In both the baseline and the treatment, there is no single strategy that can classify the majority of individuals, showing that heterogeneity of behavior is a characteristic of this games. There is exogenous heterogeneous behavior in the treatment, but heterogeneity arrises endogenously in the baseline also. The strategy "Always Defect" is used more often in the treatment, showing the spoiling power of behavioral types.

5   Conclusion

We have studied the sustainability of cooperation in a community with anonymous random matching and anonymity. We have furthermore shown that the inclusion of behavioral types makes cooperation more difficult to sustain, but does not destroy all cooperative results. Our baseline treatment replicates Camera and Casari (2009) results for the private monitoring case. We obtain higher levels of cooperation but those are not increasing over time, making our results consistent with those obtained by Duffy and Ochs (2009). We have shown that the inclusion of a short-run player in a small population lowers the level of cooperation, there is a first period effect, but cooperation increases over time. Our explanation is that players learn to deal with a noisier world as time goes by, as found by Fudenberg, Rand and Dreber (2012). We have also elicited individual strategies in our treatments. For that, we have considered memory-1 strategies and analyzed the possibility that players make mistakes in actions (for which we have identified the most frequent strategies using a maximum likelihood estimator)

and in transitions (for which we have used an automaton representation of strategies and calculated the number of classified observations by a given strategy). With both methods we have been able to determine that the one-shot game Nash-equilibrium strategy is used more often in the noisier environment, reducing expected payoffs. We have also shown that the grim trigger strategy is used less often in our treatment where there is no equilibrium strategy. The analysis presented in this paper improves our understanding of the behavior of community that are affected by noise (or frequently affected by behavioral types). We found that the rotten apple damage its companions, but is not able to completely spoil them. Given that in our treatment, none of the analyzed pure strategies, except for the strategy "Always Defect" is an equilibrium, we would like to match our experimental results with the use of belief-free equilibrium strategies. Our presumption is that players might learn to play a mixed equilibrium that allow them to coordinate on a higher payoffs outcome in the noisier environment.

References

CAMERA, G., AND M. CASARI (2009): "Cooperation among Strangers under the Shadow of the Future," *American Economic Review*, 99(3), 979–1005.

CAMERA, G., M. CASARI AND M. BIGONI (2012): "Cooperative Strategies in Groups of Strangers: an Experiment," *Games and Economic Behavior*, Forthcoming July 2012.

DAL BÓ, P. (2005): "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games," *American Economic Review*, 95(5), 1591–1604.

DAL BÓ, P. (2007): "Social Norms, Cooperation and Inequality,"*Economic Theory*, January.

DAL BÓ, P. AND G. FRÉCHETTE (2011a): "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence," *American Economic Review*, 101(1), 411–429.

DAL BÓ, P. AND G. FRÉCHETTE (2011b): "Strategy Choice In The Infinitely Repeated Prisoners Dilemma," Working Paper.

DEB, J. (2012): "Cooperation and Community Responsibility: A Folk Theorem for Repeated Random Matching Games," Working Paper.

DEB, J. AND J. GONZÁLEZ-DÍAZ (2011): "Community Enforcement beyond the Prisoner?s Dilemma," Working Paper.

DILMÉ, F. (2012): "Cooperation in Large Societies," PIER Working Paper.

DUFFY, J. AND J. OCHS (2009): "Cooperative Behavior and the frequency of social interaction," *Games and Economic Behavior*, 66, 785-812.

ELLISON, G. (1994): "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61(3), 567–588.

FISCHBACHER, U. (2007): "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10, 171–178.

FUDENBERG, D., D. KREPS, AND E. MASKIN (1990): "Repeated Games with Long-run and Short-run players," *Review of Economic Studies*, 57, 555–573.

FUDENBERG, D., AND E. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533–556.

FUDENBERG, D., RAND, D.AND A. DREBER (2012): "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World," *American Economic Review*, 102, 720–749.

KANDORI, M. (1992): "Social Norms and Community Enforcement," *Review of Economic Studies*, 59(1), 63–80.

MOSCOSO BOEDO, L. (2010): "Cooperation in the Prisoner's Dilemma with Short Run Players," CIDE- mimeo.

TAKAHASHI, S. (2010): "Community Enforcement When Players Observe Partners? Past Play," *Journal of Economic Theory*, 145 (1): 42-62.

Appendix A

The contagious strategy is an equilibrium strategy in the baseline

Here we prove that with a population of four players, the contagious strategy proposed by Kandori (1992) is an equilibrium strategy for sufficiently patient players. On the equilibrium path,the restriction that requests players to cooperate when they have seen no defection becomes:

$$\frac{(3+\delta)}{(1-\delta)(3-\delta)} \geq \frac{4}{3}.$$

Notice that the number $\frac{(3+\delta)}{(1-\delta)(3-\delta)} > 1$ $\forall \delta \in (0,1)$.Thus, this equation holds for any $\delta \geq .164$. It is sufficient to show that the off-equilibrium path condition hold for any number of deviators ($S \in \{1,2,3\}$).

For $S = 3$, $0 \geq -\frac{1}{3}$. For $S = 2$, the restriction is $\delta \leq \frac{9}{4}$, and for $S = 1$, it is $\delta \leq \frac{7}{4}$.

It is worth noticing that whenever the restriction for $S = 1$ is satisfied, then the other ones are also satisfied and that all these restrictions are satisfied for $\delta = .95$.

The contagious strategy is not an equilibrium strategy in the treatment

The case analyzed in the treatment sessions when there is a population of 4 players and one short-run in each population is a special case because after a history of cooperative encounters a player knows for sure the number of non-cooperative players in the population. In this setup, a contagious strategy on the equilibrium path requires players to cooperate after a cooperative history instead of defecting forever. With only four players, the strategy asks long-run players to cooperate at the beginning of the game (when $S = 1$) and after any cooperative history (when $S = 2$). We call S the number of non-cooperative players.

When $S = 2$ this restriction implies that $\delta \geq \frac{9}{4}$ which can not be satisfied. Thus, given our payoffs, if a player beliefs there are two non-cooperative players in the population, and the others are playing a contagious strategy he wants to deviate and defect.

Appendix B: Tables and Figures

TABLE 1: Summary of Sessions

Characteristics of experimental sessions

Treatment	Session No.	No. Of subjects	Cycle	Length
Baseline	4	16	1	5
			2	12
			3	14
			4	23
			5	53
Treatment	5	12	1	8
			2	12
			3	50
			4	38
			5	45
Treatment	6	12	1	10
			2	32
			3	46
			4	27
			5	29
Treatment	7	12	1	41
			2	9
			3	16
			4	68
			5	29
Treatment	8	12	1	35
			2	16
			3	15
			4	20
			5	52
Baseline	10	16	1	48
			2	14
			3	7
			4	27
			5	62
Baseline	11	16	1	1
			2	2
			3	2
			4	3
			5	18
	7 sessions	96 subjects		Avg. = 25.4
	(3 BL, 4 Treatment)	48 per treatment		

TABLE 2: Average Cooperation

Average of Cooperation			Mean Difference Test			
	Baseline	Treatment	t-stat	p-value	N	Sgn. Level[†]
Cycle1, Period 1	0.563	0.396	1.640	0.104	96	
	(0.072)	(0.071)				
All Cycles, All Periods	0.758	0.616	2.452	0.015	480	**
	(0.040)	(0.042)				
All Cycles, Period 1	0.554	0.483	1.554	0.121	480	
	(0.032)	(0.032)				
All Cycles, No Period 1	0.802	0.623	2.794	0.005	464	***
	(0.048)	(0.043)				
Cycles 1 and 2	0.777	0.486	3.395	0.001	192	***
	(0.063)	(0.058)				
Cycles 4 and 5	0.703	0.746	-0.463	0.644	192	
	(0.061)	(0.071)				
Cycle 1, All Periods	0.726	0.355	3.849	0.000	96	***
	(0.078)	(0.056)				
Cycle 2, All Periods	0.829	0.618	1.508	0.135	96	
	(0.100)	(0.098)				
Cycle 3, All Periods	0.827	0.613	1.580	0.118	96	
	(0.098)	(0.094)				
Cycle 4, All Periods	0.712	0.770	-0.431	0.667	96	
	(0.092)	(0.098)				
Cycle 5, All Periods	0.694	0.723	-0.218	0.828	96	
	(0.082)	(0.103)				
Cycle 1, Period 1	0.563	0.396	1.640	0.104	96	
	(0.072)	(0.071)				
Cycle 2, Period 1	0.583	0.542	0.408	0.685	96	
	(0.072)	(0.073)				
Cycle 3, Period 1	0.604	0.438	1.640	0.104	96	
	(0.071)	(0.072)				
Cycle 4, Period 1	0.521	0.521	0.000	1.000	96	
	(0.073)	(0.073)				
Cycle 5, Period 1	0.500	0.521	-0.202	0.840	96	
	(0.073)	(0.073)				

[†] Mean difference test's significance level (Sgn. Level) . *,**,*** refer to 1%, 5% and 10% significance, respectively.
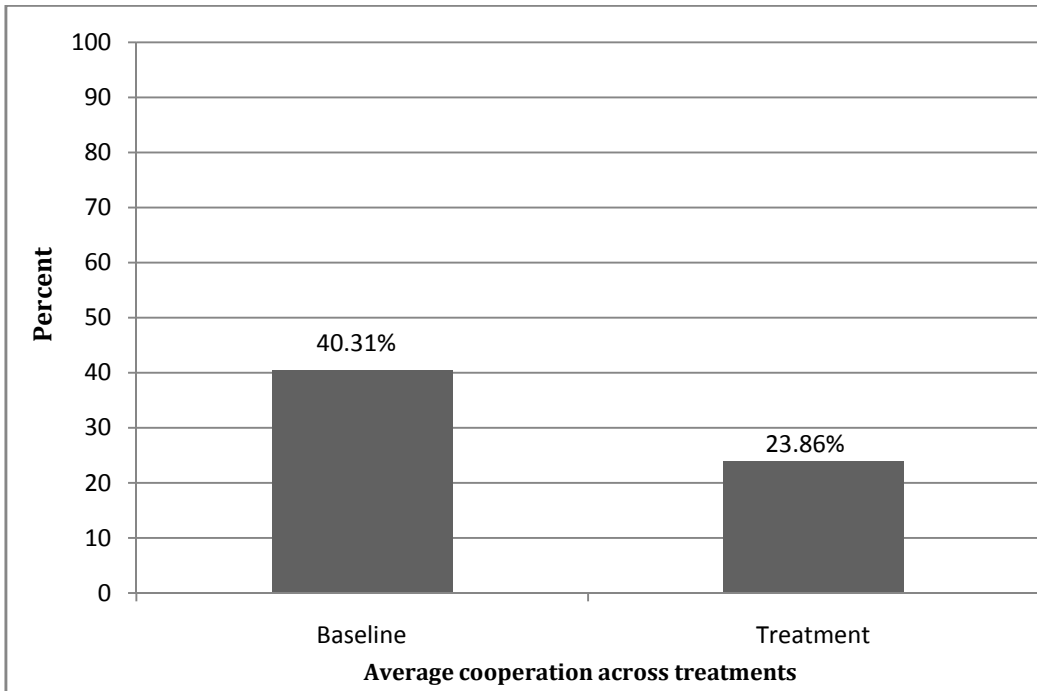
FIGURE 1: Average of Choice

40.31%

23.86%

Percent

Baseline

Treatment

Average cooperation across treatments

FIGURE 2: Average Cooperation
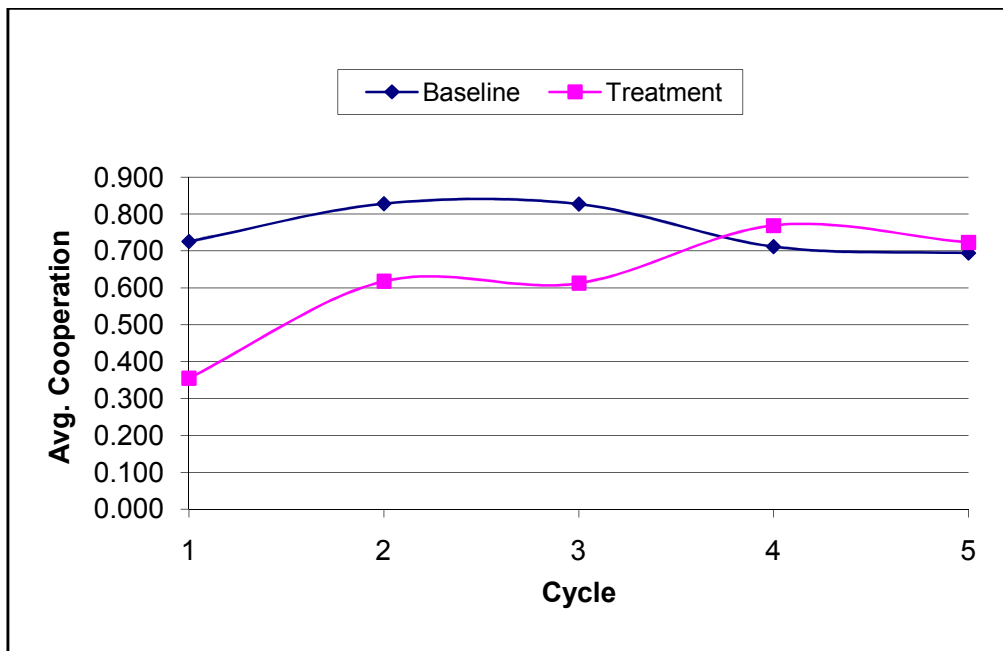

Baseline   Treatment

Avg. Cooperation

Cycle

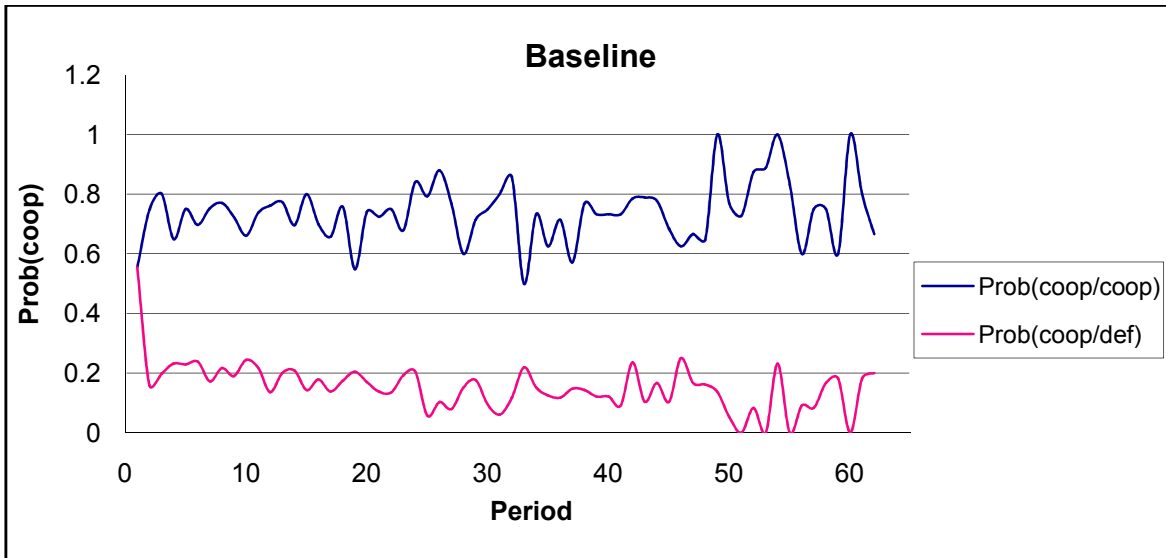FIGURE 3: Proportion of Defection

FIGURE 4: Probability of Cooperation conditional on own behavior

Treatment

Probability of Cooperation conditional on own behavior: Baseline

Difference between treatment and baseline: conditional on own cooperation
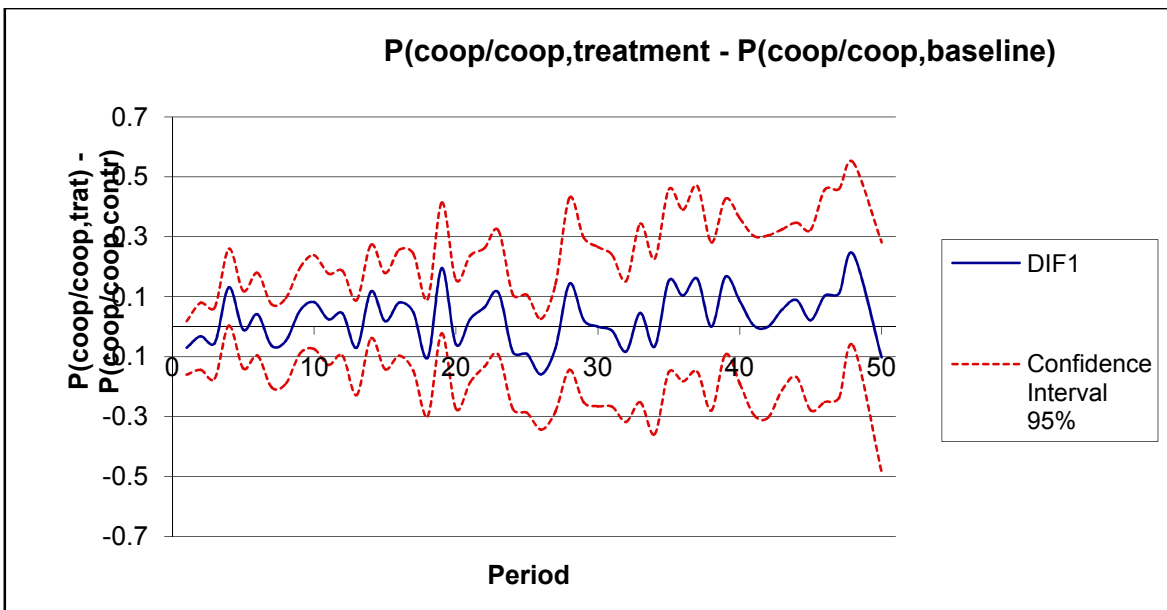
FIGURE 5: Cooperation over periods

Average of Choice over Time
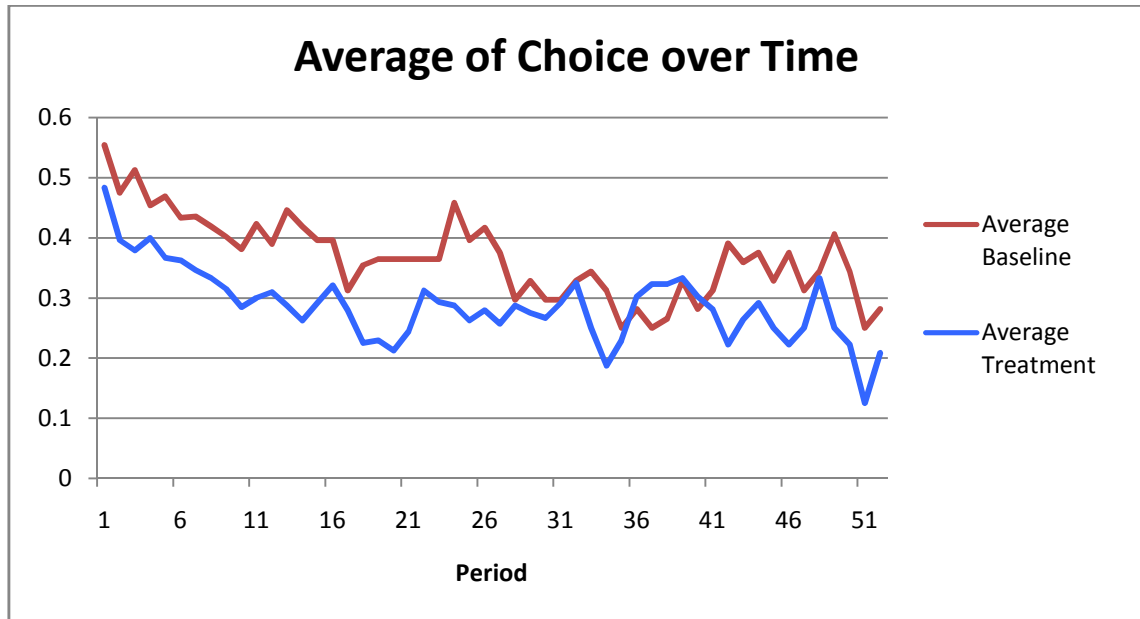
TABLE 3: Maximum Likelihood Strategy Identification

	Treatment	Baseline
Always Defect	0.4749***	0.4501***
	(0.1433)	(0.1127)
Always Cooperate	0.2018**	0.2577***
	(0.087)	(0.0863)
Grim	0.064	0.091
	(0.0603)	(0.0578)
Tit for Tat	0.2441**	0.201
	(0.1061)	(0.1336)
Win Stay Loose Shift	0.010	0.000
	(0.0201)	(0.0135)
Trigger (2 period punishment)	0.005	0.000
Gamma	0.551	0.696
	(218.3857)	(238.1455)

TABLE 4: Memory one Strategies

| | Random transitions (p = 0.05) | | | | Deterministic (p = 0.00) | | | |
| | Baseline | | Treatment | | Baseline | | Treatment | |
	N	Avg. Profit	N	Avg. Profit	N	Avg. Profit	N	Avg. Profit
All observations	240	16.54	240	12.65	240	16.54	240	12.65
C in period 1	133	15.51	116	12.28	133	15.51	116	12.28
D in period 1	107	17.82	124	13.00	107	17.82	124	13.00
Classified	168	17.15	168	12.34	114	17.94	92	12.45
C in period 1	90	16.15	69	11.82	57	16.70	28	12.27
systematically coop.	59	16.30	51	11.51	38	17.03	22	12.21
forgiving	40	17.34	17	10.71	32	18.10	4	11.76
grim trigger	43	17.46	19	11.38	33	17.94	4	11.81
unconventional	11	19.50	2	13.54	8	19.44	0	0.00
D in period 1	78	18.32	99	12.70	57	19.17	64	12.52
systematically def.	68	18.93	91	12.48	51	19.86	64	12.52
unconventional	10	14.14	8	15.12	6	13.35	0	0.00
Unclassified	72	15.10	72	13.39	126	15.28	148	12.78
C in period 1	43	14.16	47	12.96	76	14.61	88	12.29
D in period 1	29	16.50	25	14.20	50	16.29	60	13.51
Opponents play C and D	191	16.36	192	13.51	191	16.36	192	13.51
Classified	123	16.80	125	13.35	70	17.94	70	13.30
C in period 1	62	15.47	51	12.89	30	16.12	25	12.78
systematically coop.	36	14.60	38	12.55	16	14.83	21	12.56
forgiving	16	15.10	6	12.72	9	17.04	2	13.88
grim trigger	19	15.74	8	13.81	10	16.63	2	13.98
unconventional	10	18.70	2	13.54	7	18.29	0	0.00
D in period 1	61	18.16	74	13.68	40	19.30	45	13.58
systematically def.	55	18.50	66	13.50	38	19.55	45	13.58
unconventional	6	15.09	8	15.12	2	14.64	0	0.00
Unclassified	68	15.57	67	13.80	121	15.45	122	13.63
C in period 1	40	14.71	44	13.35	72	14.78	70	13.22
D in period 1	28	16.79	23	14.67	49	16.45	52	14.20
Opponents always play C	25	26.64	0	0.00	25	26.64	0	0.00
Classified	25	26.64	0	0.00	24	26.67	0	0.00
C in period 1	17	25.21	0	0.00	16	25.16	0	0.00
systematically coop.	16	25.06	0	0.00	15	25.00	0	0.00
forgiving	16	25.06	0	0.00	15	25.00	0	0.00
grim trigger	16	25.06	0	0.00	15	25.00	0	0.00
unconventional	1	27.50	0	0.00	1	27.50	0	0.00
D in period 1	8	29.69	0	0.00	8	29.69	0	0.00
systematically def.	7	30.00	0	0.00	7	30.00	0	0.00
unconventional	1	27.50	0	0.00	1	27.50	0	0.00
Unclassified	0	0.00	0	0.00	1	26.00	0	0.00
C in period 1	0	0.00	0	0.00	1	26.00	0	0.00
D in period 1	0	0.00	0	0.00	0	0.00	0	0.00

TABLE 4 (*continued*): Memory one Strategies

| | Random transitions (p = 0.05) | | | | Deterministic (p = 0.00) | | | |
| | Baseline | | Treatment | | Baseline | | Treatment | |
	N	Avg. Profit	N	Avg. Profit	N	Avg. Profit	N	Avg. Profit
Opponents always play D	24	7.41	48	9.22	24	7.41	48	9.22
Classified	20	7.46	43	9.38	20	7.46	22	9.74
C in period 1	11	5.98	18	8.80	11	5.98	3	8.10
systematically coop.	7	5.00	13	8.49	7	5.00	1	5.00
forgiving	8	6.35	11	9.61	8	6.35	2	9.64
grim trigger	8	6.35	11	9.61	8	6.35	2	9.64
unconventional	0	0.00	0	0.00	0	0.00	0	0.00
D in period 1	9	9.26	25	9.80	9	9.26	19	10.00
systematically def.	6	10.00	25	9.80	6	10.00	19	10.00
unconventional	3	7.78	0	0.00	3	7.78	0	0.00
Unclassified	4	7.19	5	7.90	4	7.19	26	8.79
C in period 1	3	6.81	3	7.32	3	6.81	18	8.67
D in period 1	1	8.33	2	8.77	1	8.33	8	9.05

FIGURE 6: Classification by transition error

Percentage of classified observations