

Bayer, Ralph-C; Renner, Elke; Sausgruber, Rupert

Working Paper

Confusion and Reinforcement Learning in Experimental Public Goods Games

NRN Working Paper, NRN: The Austrian Center for Labor Economics and the Analysis of the Welfare State, No. 0922

Provided in Cooperation with:

NRN: The Austrian Center for Labor Economics and the Analysis of the Welfare State, Johannes Kepler University Linz

Suggested Citation: Bayer, Ralph-C; Renner, Elke; Sausgruber, Rupert (2009) : Confusion and Reinforcement Learning in Experimental Public Goods Games, NRN Working Paper, NRN: The Austrian Center for Labor Economics and the Analysis of the Welfare State, No. 0922, Johannes Kepler University Linz, NRN - The Austrian Center for Labor Economics and the Analysis of the Welfare State, Linz

This Version is available at:

<https://hdl.handle.net/10419/115030>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**Confusion and Reinforcement Learning in
Experimental Public Goods Games**

by

Ralph-C BAYER, Elke RENNER, Rupert SAUSGRUBER

Working Paper No. 0922

October 2009

Supported by the
Austrian Science Funds

FWF

**The Austrian Center for Labor
Economics and the Analysis of
the Welfare State**

JKU Linz
Department of Economics
Altenberger Strasse 69
4040 Linz, Austria
www.laborrrn.at

Corresponding author: rupert.sausgruber@uibk.ac.at
phone +43 (0) 512 507 7176, - 2970 (fax)

Confusion and Reinforcement Learning in Experimental Public Goods Games*

Ralph-C Bayer[†], Elke Renner[‡] and Rupert Sausgruber[§]

September 2009

Abstract

We use a limited information environment to mimic the state of confusion in an experimental, repeated public goods game. The results show that reinforcement learning leads to dynamics similar to those observed in standard public goods games. However, closer inspection shows that individual decay of contributions in standard public goods games cannot be fully explained by reinforcement learning. According to our estimates, learning only accounts for 41 percent of the decay in contributions in standard public goods games. The contribution dynamics of subjects, who are identified as conditional cooperators, differ strongly from the learning dynamics, while a learning model estimated from the limited information treatment tracks behavior for subjects, who cannot be classified as conditional cooperators, reasonably well.

Keywords: public goods experiments, learning, limited information, confusion, conditional cooperation

JEL classification: C90, D83, H41.

*We are grateful for financial support by the Austrian Science Fund (FWF) under Projects No. P17029 and S10307-G14 as well as by the Faculty of Profession Research Grant Scheme of the University of Adelaide.

[†]School of Economics, University of Adelaide, Napier Building, SA 5005 Adelaide, Australia; E-mail: ralph.bayer@adelaide.edu.au

[‡]School of Economics, Sir Clive Granger Building, University Park, Nottingham NG7 2RD, United Kingdom; E-mail: elke.renner@nottingham.ac.uk

[§]Department of Economics, University of Copenhagen, DK-1455 Copenhagen K, Denmark; E-mail: rupert.sausgruber@uibk.ac.at

1 Introduction

The degree to which subjects, who participate in an experiment, understand the task involved is typically not observable by the researcher. Efforts to reduce confusion through detailed instructions and to control for confusion using control questions are valuable but can never fully rule out that some observed behavior was caused by a lack of understanding. Unlike naturally selected and experienced individuals in real-world situations, people who come to the laboratory, typically are unfamiliar and inexperienced with the choices they need to make. Additionally, experimental environments are often artificial and deliberately do not provide a real-world context to the experiment, which makes understanding difficult. Arising non-observable confusion poses a threat to the internal validity of experimental results. Some behavior caused by confusion might be wrongly attributed to some other cause. At best, confusion – if it does not lead to systematic changes of behavior – only causes an inference problem through increased noise. Generally, confusion may thus impair the researchers' confidence in internal and external validity of experimental results (for a general discussion see Levitt and List 2007).

A widely used alternative instrument intended to overcome the problem of confusion is repetition. Repetition enables subjects to gain familiarity with the new environment and to experience the consequences of their decisions. However, in games, where payoffs depend on the actions of all players in a group, a characteristic of repetition is that the subjects do not only learn about the game itself but can also use feedback to predict future play of others. The researcher then faces the complicated task of separating the strategic response to this information from learning behavior that results from overcoming confusion. This task is complicated further, as there might also be interaction effects, since a strategic response is likely to depend on a player's assessment of how confused the group members are. In this paper we tackle this problem for the domain of public goods games.

Experimental research on the private provision of public goods has an interest to discriminate between strategic behavior and confusion. The repeated public goods experiment is among the best known games in economics. It is meant to mimic behavior in a wide class of situations outside of laboratories and classrooms. Today, social scientists employ the experiment as tool to measure social norms and preferences (see, e.g., Camerer and Fehr 2003). The outcome of the typical experiment is that subjects contribute more to the public good than the standard model predicts, while contribution levels decrease with repetition (for a recent and comprehensive list of references reporting an increase of free riding over multiple decision rounds see Fischbacher and Gächter 2009). One early explanation for this pattern of behavior is that contribution levels fall due to the learning

of initially confused subjects (for a discussion see Ledyard 1995). Another explanation, which is favored by many researches, is that the decay of contributions stems from the heterogeneity of social preferences within a group, where conditional cooperators start out with high contributions but consecutively adjust their contributions downwards as a negative-reciprocal reaction to selfish group members (see Andreoni 1995; Kurzban and Houser 2005; Muller et al. 2008; Fischbacher and Gächter 2009).¹

In this paper we take a new avenue to assess the role of learning of confused subjects in explaining the dynamic pattern of choices in repeated public goods experiments. Our approach starts with the assumption that a confused subject neither understands the incentive structure, nor uses the contributions of other subjects as a basis for imitation learning. One might label this kind of confusion as “ignorant”. Admittedly, this view of confusion is somewhat extreme but has a considerable advantage over other definitions when it comes to experimental methodology. Confusion of this kind can be studied in an experiment that withholds information about the payoff structure and the choices of other group members from subjects. A treatment, where subjects do neither know the incentive structure nor the past choices of the group members, perfectly mimics this kind of “ignorant” confusion. We run such a treatment and compare the contribution dynamics to those in a standard public goods game.

The traditional approach to studying confusion in public goods games follows a different strategy. Confusion is identified as a residual instead of implementing conditions that allow for studying learning directly. Designing two treatments, a standard treatment and at least one other, where strategic and/or social motives are excluded by design, under certain assumptions allows for an isolation of the strategic and/or social component. The difference in contributions between the treatments is attributed to strategic and social factors, respectively. Then the residual contributions are attributed to confusion. In this tradition Andreoni (1988) compares public goods games with partner and stranger treatments, in order to isolate strategic repeated game effects. This study and the replications by Croson (1995) and Keser and van Winden (2000) produced outcomes that are consistent with strategic motives rather than learning being the explanation for an increase in free-riding.²

Andreoni (1995) introduced a treatment where subjects played a public goods game but were paid according to how they rank compared to group members with respect to their experimental earnings. In this design the dominant strategy of full free-riding

¹Gintis et al. (2003) explain similar dynamics with an evolutionary approach.

²Andreoni (1988) also observed that cooperation increases again after a restart of the game, which is at odds with the simple learning hypothesis. See also Cookson (2000).

is preserved while the imposed zero-sum nature is designed to remove any cooperation incentives. In this treatment, subjects' contributions fell rapidly already in early decision rounds, which suggests that the decline of cooperation in the standard game "may not be due to learning, but instead ... due to frustrated attempts at kindness" (p. 892).

Houser and Kurzban (2002) designed a study to tighten the bounds on the impact of confusion as identified by Andreoni (1995). They ran a "computer condition" in which all other group members were simulated by automata. Off-equilibrium play in the computer condition cannot be associated with social motives towards other players. Hence, a comparison to a standard public goods game was used to isolate the effect of confusion. The results were that approximately 50% of all contributions can – under certain assumptions – be attributed to confusion and that the contributions in the computer condition fell with repetition at a rate much higher than in the human condition.³ The second observation casts doubts on the previously identified role of reciprocity in causing the decay of contributions.

We contribute to this literature in that we analyze the contribution behavior under a condition of limited information. The study of behavior in limited information environments is common in experiments on learning. Mookherjee and Sopher (1994) varied the information available to the subjects about past choices and payoffs of their opponents in a repeated matching pennies game. They found that experience with choices in the past affected behavior under limited information. Van Huyck et al. (1996) asked whether a simple reinforcement learning algorithm can predict behavior in a coordination game. To do so, they limited the information available to subjects to their feasible choices and their experienced payoffs. The results were that the median play converges to the interior equilibrium and that convergence is much faster than reinforcement learning would predict. Chen and Khoroshilov (2003) evaluate alternative payoff-based learning models to explain behavior in a cost-sharing and coordination game under limited information. They found that a simple reinforcement learning model tracks the data not as well as the payoff-assessment learning model (Sarin and Vahid 1999) or the experience-weighted attraction learning model (Camerer and Ho 1999).

We believe that the method of limiting information can generate novel insights that add to the hitherto incomplete picture of learning in the repeated public goods experiment. In particular, the extreme assumption of ignorant confusion assures that subjects can only learn by reinforcement. More sophisticated learning models such as belief learning, experience-weighted attraction learning or rule learning are ruled out (for a nice

³Ferraro and Vossler (2006) used a similar design for the same purpose and complemented it with econometric modeling.

overview of alternative learning models see Camerer 2003, chapter 6, and the references therein). In fact, all studies that try to isolate the dynamic effects of confusion have to make assumptions about the nature of confusion. For example, in the study by Houser and Kurzban (2002) contributions fell faster in the computer than in the human condition. This observation might indicate that cooperation is stable at an initial level in the absence of confusion. The authors note that such a conclusion would be warranted only if “cooperation due to confusion is similar in the human and computer conditions” (p. 1066). Such an assumption may not hold empirically (see Duersch et al. 2009, who explore how subjects learn to play a Cournot duopoly game against computers that are programmed to follow one of various learning algorithms).

The merit of studying confusion as ignorance is that it gives sufficient structure to model behavior by simple learning dynamics and at the same time avoids a great deal of other necessary implicit assumptions. While measured confusion in traditional experiments is a residual and might contain other impact factors not removed by the design, a limited information approach provides additional control. On the flip side, the assumption of ignorant confusion comes at a price. Subjects that are more sophisticated but still confused are not properly represented. Confused subjects, who are cleverer than the assumed “ignorant” subjects, are expected to learn and reduce their contributions more quickly (if they are self interested).⁴ Hence, the dynamics observed in a treatment inducing confusion defined as ignorance can be seen as a lower bound to the true impacts of confusion on contribution dynamics. We believe that complementing the traditional approach, which can be seen as providing an upper bound (in particular Houser and Kurzban, 2002), with a methodology that provides a lower bound is valuable in that it helps to better understand the impact of confusion on free-riding.

Our results are as follows. At the aggregate level we find that contributions drop off in both our confusion condition and the standard public goods game. This observation supports the claim that reinforcement learning leads to dynamics similar to those readily interpreted as conditional cooperation in standard public goods games. However, we also find that the contributions in the standard treatment decrease at a significantly higher rate than in the confusion condition. According to our estimates learning only accounts

⁴Janssen and Ahn (2006) apply an agent-based modeling approach to estimate the dynamic pattern of contributions in the repeated public goods experiment. They conclude that the stylized facts about the contribution dynamics are explicable only by a complicated mixture of social preferences, learning, and signaling. Within the learning explanation, they found reinforcement learning to play a prominent role. However, their results depend very much on the specific functional forms they used. Their study also differs from ours in that it does not employ limited information to reduce the complexity of the choice environment.

for 41 percent of the contribution dynamics. We also analyzed whether the reduction of confusion (due to reinforcement learning) can cause correlated behavior at the individual level that can be mistaken for conditional cooperation. A dynamic panel analysis provides direct evidence that the correlation of contributions with average past contributions of the other group members cannot be explained by the reduction of confusion. Moreover, we used our estimation results to simulate contributions of those classified as “conditional cooperators” and of “others”. While the learning dynamics estimated from our limited information treatment does a poor job at explaining the behavior of “conditional cooperators”, it tracks that of the “others” reasonably well. Finally, we use our estimation results to investigate if a hypothetical group of four conditional cooperators would be able to prevent contributions from declining. We find that they would not be able to achieve stable cooperation.

The next section describes the experimental design. Section 3 presents our findings and section 4 concludes.

2 Experimental design

Our design deliberately withholds information about the game from the subjects. Participants in the Learning Condition only know the admissible action space and that the environment may change over time. After each stage subjects are informed about their payoff (see Appendix A).

We employ a within-subject comparison across two phases, which we complement by a control treatment that ensures that we can test for the effect of the same subject participating in two treatments. The within-subject comparison enables us to observe the same subject’s behavior both in a state of confusion (Learning Condition) and in an otherwise identical public goods game, where subjects have all the information (Standard Condition). Each condition consists of 20 periods. In the Learning Condition, the subjects choose a number between 0 and 20 in each period. The subjects do not know that this number is a contribution choice. The instructions tell them that the aim of the experiment is the study of learning behavior. We inform the subjects that their payoff is determined by their choice and “other factors that might change across periods”.

At the end of this first phase, the subjects in the within-subject treatment are informed that a new experiment (the Standard Condition) will start. Only at this stage, they are given instructions for a standard linear public goods game, where they are assigned to

groups of four.⁵ In the control treatment subjects only play the standard public goods game with instructions. This design enables us to test if the fact that the subjects played phase one first altered the behavior in the following standard public goods game. We did not find any differences in behavior and therefore are confident that our within-subject analysis is valid.

The structure of the experimental public goods game was as follows. Every period the subjects receive 20 points as their initial endowment. Every point invested into a public good pays 0.4 Australian Cents to each subject in the group, while every point kept for private investment pays 1 Cent to only the subject who kept it. This underlying structure was the same in both Learning and Standard Condition. The sole difference was the information the subjects received. In the standard condition full free-riding is a dominant strategy and therefore in the subgame-perfect Nash Equilibrium no one ever contributes anything. It would be socially optimal however, if everyone contributed their full endowment.

3 Results

We ran five sessions with 16 or 20 subjects each. Out of 96 subjects in total, 60 were participating in the within-subject treatment (Learning Condition in phase 1 followed by the Standard Condition in phase 2) and 36 in control treatment (the Standard Condition in phase 1, no phase 2). The subjects were first-year students at the University of Adelaide from a variety of fields, who had never before been in an experiment. The experiments were conducted with the software package *z-Tree* (Fischbacher 2007). The experiment lasted between 15 (control treatment) and 25 minutes (both phases), and the average subject earned the equivalent of US\$ 10.1 (in Australian Dollars) within this time.

Figure 1 shows the time series of the average contributions in the learning condition as a percentage of the endowment. The black line shows the average observed contribution behavior. As one would expect for a situation where subjects cannot understand the implications of their behavior, the contributions start out around the midpoint of the admissible action space. With repeated play, however, contributions drop off from 53.4 percent of the total endowment in period one to 35.7 percent in period 20. On average (using a linear trend) contributions drop by 0.18 units per period. This negative time trend is significant at the 1-percent level.⁶

⁵We told the subjects in the within-subject condition that they will now participate in a new experiment, and also used very different styles for the instructions and screen layouts.

⁶We use robust standard errors adjusted for clustering on groups throughout this paper.

3.1 A learning model

The observation that chosen numbers decrease with repetition in the Learning Condition just as contributions do in the standard public goods game provides support for the claim that learning can be mistaken for conditional cooperation. To gain some more confidence that our Learning Condition accurately picks up learning dynamics – and nothing else – we compare the actual behavior to the simulated outcomes of a simple learning model.

A stochastic adaptive decision rule maps the observed history – in our case past contributions and past profits – into a probability distribution from which the actual contribution is drawn. In the past a large variety of different learning models requiring different levels of sophistication of subjects have been proposed.⁷ The rule we use is in the tradition of reinforcement learning. This is governed by the sparse information environment in our Learning Condition. “Confusion as ignorance” as observed in this treatment, means that individuals do not know the underlying structure of the public goods game. They also do not observe any action of the other group members. Therefore, more sophisticated learning models such as belief learning, experience-weighted attraction learning or rule learning are ruled out. We also restrict the memory of the learning rule. Individuals decide over their choices by comparing the payoffs resulting from their last two choices.⁸

So we arrive at an extremely simple learning model. Define the set of players as $I = \{1, 2, 3, 4\}$ and the action space as $C = \{0, 1, \dots, 20\}$. Denote the contribution of person $i \in I$ in period $t \in \{1, 2, \dots, 20\}$ as $c_t^i \in C$. The player uses the payoffs p and the own choices of the last two periods to determine the contribution in the current period (if possible). The attraction of choosing a certain contribution $A(c_t^i)$ is therefore a function of the two past contributions and the payoffs in the two last periods:

$$A(c_t^i) = f(c_{t-1}^i, c_{t-2}^i, p_{t-1}^i, p_{t-2}^i). \quad (1)$$

After having observed the two last outcomes given the choices made, for the next round individuals only consider choices which are closer to the choice that resulted in a higher payoff. Suppose c_{t-1}^i was greater than c_{t-2}^i and the payoff in period $t-1$ was greater than in period $t-2$, then the individual only chooses values in the interval from the midpoint between the two previous choices to the maximum choice (20). For equal profits in periods $t-1$ and $t-2$ the support is $[0, 20]$, as then the history contains no information about in

⁷See Camerer (2003, chapter 6) and the references therein for a nice overview.

⁸We use such a short memory for two reasons: i) in the instructions we inform subjects that the environment might change over time and ii) Sarin and Vahid (2004) have shown that the use of rapidly decaying past attractions improves the fit of reinforcement-learning models.

which direction to go. Moreover, the support will also be the whole spectrum of possible choices if the previous two choices were identical.

To find the region of choices (the support) that satisfies these conditions given the history, define the changes in choices and payoffs between periods $t - 1$ and $t - 2$ as

$$\Delta p_t^i \equiv p_{t-2}^i - p_{t-1}^i \quad (2)$$

$$\Delta c_t^i \equiv c_{t-2}^i - c_{t-1}^i \quad (3)$$

Then we can introduce a variable d_t^i that tells us whether the player wants to choose a number closer to the higher ($d_t^i = 1$) or the lower of the previous choices ($d_t^i = -1$):

$$d_t^i = \text{sign}(\Delta p_t^i \cdot \Delta c_t^i). \quad (4)$$

Note that if either the profits or the previous choices have not changed between periods $t - 2$ and $t - 1$ then we have $d_t^i = 0$. Denoting the admissible support for period t as C_t^i we have:

$$C_t^i = \begin{cases} \{c \in C : c \leq (c_{t-1}^i + c_{t-2}^i)/2\} & \text{if } d_t^i = -1 \\ \{c \in C : c \geq (c_{t-1}^i + c_{t-2}^i)/2\} & \text{if } d_t^i = 1 \\ \{c \in C\} & \text{if } d_t^i = 0 \end{cases} \quad (5)$$

Next, we have to specify which point within the admissible range will be chosen. The simplest assumption is that subjects are equally likely to choose any element of C_t^i .⁹ To implement this we set the attraction for a choice in C_t^i equal to one, while the attraction of a contribution outside of C_t^i is set to zero:

$$A(c_t^i) = \begin{cases} 1 & \text{if } c_t^i \in C_t^i \\ 0 & \text{if } c_t^i \notin C_t^i \end{cases} \quad (6)$$

To arrive at the desired uniform distribution over the support C_t^i we transform attractions into probabilities using the following rule:

$$g(c_t^i) = \frac{A(c_t^i)}{\sum_{c_t^i \in C_t^i} A(c_t^i)} \quad (7)$$

Analyzing the data, we found that the median of choices for both experimental conditions is approximately in the middle of the support, which is consistent with assuming

⁹This assumption differs slightly from a traditional reinforcement learning model in that it allows for “strategy similarity”. Sarin and Vahid (2004) show that this modification helps to explain behavior in minimal information games like ours.

choices with equal probabilities. However, we observed quite a few focal points (bottom or top of the range), which cannot be modeled with the uniform distribution. This is not problematic, since this clustering around boundary values typically occurred in the Standard Condition, where we do not expect the learning model to fit well.

The remaining question is how to deal with the choice behavior of the individuals in periods one and two. In these early periods there is not enough information available to use reinforcement-learning. We follow the widespread approach and use the observed choice distribution in those first two periods. The first two choices are assumed to be driven by some factors exogenous to our learning model, such as focal points.

In summary, the choice rule in our learning model is simple. A subject only chooses numbers that are closer to the number that led to the higher payoff in the previous two periods. For simplicity we assume that subjects randomize over all choices that are in the remaining domain with equal probability. In cases where a subject is not able to learn anything from her last choices – either the last two choices or the last two payoffs are equal – a subject randomizes over the unrestricted domain.


This simple model produces quite tricky dynamics. The grey line in Figure 1 shows the result of simulated behavior from the model. We simulated 5000 groups. As the starting values are not determined endogenously in the model, we drew them from the empirical distributions of the real contributions observed in the two first periods.¹⁰ We see that a model as simple as ours does very well at tracking the behavior in the Learning Condition. Hence, we feel confident to conclude that our Learning Condition can be used to isolate learning dynamics from the dynamics generated by strategic behavior of any kind.

Our next step is to quantify how much of the downward trend of contributions in the standard public goods game (Standard Condition and control treatment) is due to the reduction of confusion. In Figure 2, the black line plots the average contributions of the subjects in the Standard Condition. Obviously, the rate at which subjects reduce their contributions is greater than in the Learning Condition. In the Standard Condition, contributions drop from 57.7 percent of the total endowment in period one to 16.6 percent in period 20.

As before, the grey line in Figure 2 shows the choices simulated using the learning model with the starting values drawn from the empirical distributions of the first two periods. The learning model does not fit the data well. The dynamics in the standard public

¹⁰We also simulated the learning model with different initial choices. Even starting with extreme values (only 0 or 20) simulated behavior quickly converges to that following starting values drawn from the empirical distribution.

Figure 1: Time series of average contributions in the Learning Condition



goods game appear to be different than the simulated reinforcement-learning dynamics, which performed so well at explaining behavior in the Learning Condition. A Wilcoxon matched-pairs test (for the subjects that played both phases) confirms that the deviations of the average group contributions from the simulated contribution averages summed over the 20 periods are significantly smaller in the Learning Condition ($p < 0.01$, $N = 15$).¹¹

The analysis above shows that a reinforcement-learning model explains the dynamics in the Learning Condition, while it fails to explain the dynamics in the standard public goods game. The learning speed predicted by the model is insufficient to explain actual behavior in the Standard Condition. Therefore confusion defined as ignorance cannot explain all of the decrease in cooperation in the public goods game. Even after controlling for learning dynamics some decay in contributions still remains.¹² In the Standard Condition the linear time trend is -0.44 , which is significantly different both from zero and -0.18 (the time trend estimated for the Learning Condition).

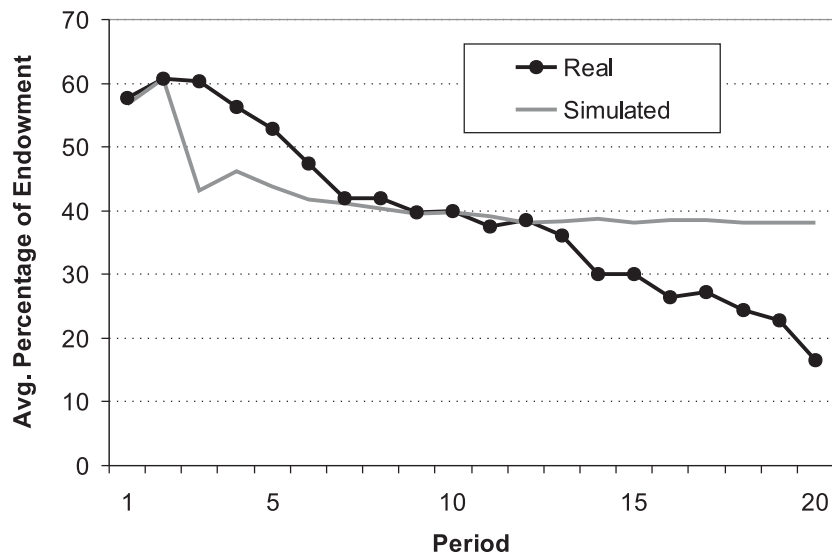
To summarize: Assuming that confusion in a standard public goods game takes the form of ignorance, we conclude that the reduction of confusion accounts for $-0.18 / -0.44 = 41$ percent of the total decrease in cooperation. From our discussion above, this figure gives a lower bound for the true impact of learning on contribution dynamics, as

¹¹The average mean square error of the simulation is more than four times larger in the Standard condition (3.66 vs. 0.82 points).

¹²Taking the difference between the real and simulated contributions in Figure 2 reveals that the remaining dynamics still point downwards.

confusion as ignorance is an extreme assumption.

Figure 2: Time series of average contributions in the standard condition



3.2 Econometric Analysis

In what follows, we investigate the dynamics in the two treatments econometrically. We run two separate dynamic panel regressions – one per treatment – where we try to explain c_t^i , the contribution in period t , with lagged own contributions (c_{t-1}^i, c_{t-2}^i) and the past average contributions of the other group members ($\bar{c}_{t-1}^{-i}, \bar{c}_{t-2}^{-i}$). We also include period dummies. This econometric model is designed to test for reciprocity in the Standard Condition and to check if the Learning Condition can produce data that could be misinterpreted as being generated by reciprocity. We are mainly interested in the coefficient on the lagged average contribution of the other group members. A positive significant coefficient in the Standard Condition is evidence for reciprocity, as it measures how subjects react to past kindness (or unkindness) of the other group members. However, a positive coefficient in the Learning Condition would mean that the simple learning environment can produce the same dynamics. This would support the hypothesis that the positive coefficient in the Standard Condition might not be due to reciprocity but learning.

Note that simple random-effect panel-data estimators yield biased estimates if lagged dependent variables are included, since by construction, the unobserved panel-level effects are correlated with the lagged dependent variables. For this reason, we use an Arellano-Bond GMM estimator (Arellano and Bond 1991) with additional moment conditions

(Arellano and Bover 1995, Blundell and Bond 1998), which improve the performance of the estimator.¹³

3.2.1 The standard condition

Table 1 reports the estimation results for the Standard Condition. The standard errors for the coefficients are reported in parentheses. Two stars indicate significance on the one-percent level, one star on the five percent level. We find that the evolution of the contributions of an individual is governed by two main influences. On the one hand, there is a feedback structure. Contributions react to the past own and past average contributions of others. On the other hand there are some period effects, which are independent from the feedback structure.

The regression shows that the contribution in the previous period has a positive influence on the actual contribution. Contributions are somewhat persistent. More importantly, actual contributions c_t^i increase with the average contribution of the other group members from the previous period \bar{c}_{t-1}^{-i} . Previously generous group members who contribute, induce the subjects to reciprocate with a higher contribution. An increase in the average last period contribution of the other group members by one unit induces an increase of a subject's contribution by about one third of a unit. Through the persistence of the own contribution this increase will have a positive but decaying influence on future contributions. In the public goods game this kind of reciprocity is responsible for a large part of the decay. As some subjects reduce their contributions, other subjects will respond with negative reciprocity and will reduce theirs in the future.

If we ignore the insignificant second lags and all time effects (but include the intercept) then contributions within a group have a single stable steady state where all subjects contribute somewhat less than 12 monetary units or 60 percent.¹⁴ If subjects kept behaving like this from the beginning of the game (here period 3) they would soon reach this steady state. However, the period effects push the steady state downwards. If we take the average of all time dummies in the model then the steady state the contributions converge to is at about 4.8 monetary units (24 percent). For periods with large negative time effects (such as periods 14, 18 and 20) the steady state shifts to full free riding. The downward

¹³We use the two-step estimation technique with robust standard errors. All our models passes tests for the validity of the over-identification restrictions (Sargan Test) and autocorrelation in the first-differenced errors (Arellano-Bond Test).

¹⁴To see this take $c_t^i = \alpha + \beta_1 c_{t-1}^i + \beta_2 \bar{c}_{t-1}^{-i}$, replace the lagged own and average other contributions by the own actual contributions, plug in the estimated coefficients and solve.

Table 1: Dynamic Panel Estimation of contributions in the standard condition

Variable	Coefficient	(Std. Err.)
<i>Lagged own contribution</i>		
c_{t-1}^i	0.299**	(0.088)
c_{t-2}^i	0.055	(0.078)
<i>Lagged average contribution of others</i>		
\bar{c}_{t-1}^{-i}	0.324**	(0.107)
\bar{c}_{t-2}^{-i}	-0.037	(0.123)
<i>Period dummies, t = 3 omitted</i>		
$t = 4$	-0.807*	(0.382)
$t = 5$	-1.334**	(0.307)
$t = 6$	-1.646**	(0.361)
$t = 7$	-2.167**	(0.602)
$t = 8$	-1.731*	(0.755)
$t = 9$	-1.772	(0.944)
$t = 10$	-1.535	(0.796)
$t = 11$	-2.157*	(0.926)
$t = 12$	-1.809	(0.994)
$t = 13$	-2.318*	(0.968)
$t = 14$	-3.175**	(1.048)
$t = 15$	-2.301*	(1.175)
$t = 16$	-3.061*	(1.419)
$t = 17$	-2.599	(1.553)
$t = 18$	-3.134	(1.610)
$t = 19$	-3.243*	(1.638)
$t = 20$	-4.262*	(1.777)
<i>Intercept</i>	4.580	(2.788)
N		1728
$prob > \chi_{(21)}^2$		0.000

shift caused by the period effects tends to be larger in the later periods of the game. We also find a strong end effect in the final period.

Note that the reciprocity resulting from the estimation is quite weak compared to

strong reciprocity or tit-for-tat behavior.¹⁵ A large part of the contribution is not determined by the lagged contributions of the other group members. There is quite some persistence. *Ceteris paribus*, subjects who start with higher contributions keep contributing more. In addition we observe otherwise not explained time dependent level effects. These level effects tend to reduce contributions as the game progresses. Consequently, reciprocity explains some but not all decay of contributions. Compared to period three (the first period for which we can estimate period effects), *ceteris paribus* a subject contributes about 4.8 (out of a maximum of 20) monetary units less in the last period. However, it is important to keep in mind that this is average behavior and does not take into account that there might be different types of subjects.

Figure 3 shows the contribution dynamics under the assumptions that the players 2,3,4 initially choose identical contributions. This reduces the dimensionality of the dynamics to two and allows plotting. In the graph we include an average of the period effects (β , which is -2.17). The blue (bold) dot represents the unique steady state the contributions converge to. Note that this steady state and the fact that it is globally and asymptotically stable are independent of our assumptions that three players start with the same initial contribution.¹⁶

3.2.2 The Learning Condition

We have seen that the dynamic panel estimation for the Standard Condition is consistent with negative reciprocity being responsible for (at least some of) the decay of contributions. We now ask if this could be an artifact of learning dynamics. In other words, does the Learning Condition produce similar dynamics as the Standard Condition? As we have seen in previous sections, both treatments show decay of contributions over time. The decay is stronger in the the standard condition though. This alone does not prove that the processes driving the decay are different. For this reason we run the same regression we used for the Standard Condition also for the learning data. Table 2 shows the results.


The regression results clearly show that the dynamics in the Learning Condition are quite different to those in the Standard Condition. There is no significant feedback. There is no significant influence of past own or past average contributions of other group members.¹⁷ The regression does a bad job explaining the data in the Learning Condition.

¹⁵Strong reciprocity would imply a coefficient of one, while tit-for-tat behaviour would also require that all other coefficients are equal to zero.

¹⁶The coefficient matrix of the dynamic system has eigenvalues which are below unity.

¹⁷This is established by testing if the coefficients on the four variables capturing past contributions (own or by others) are jointly different from zero, which they are not ($p > 0.87$).

Figure 3: Contribution dynamics in the standard condition (with average period effects)



The dynamics are driven by factors not included in the model, as only the intercept and some period effects are significant.

Summarizing our results so far, we can say that the dynamics from the two different treatments are clearly distinct. A simple learning model explains the data in the Learning Condition reasonably well, while it fails to explain the data from the Standard Condition. Providing further evidence for this claim, a dynamic panel-data regression designed to identify the occurrence of reciprocity fits well in the case of the standard public goods game. The parameter on the influence of the past average contribution of other group members on the own contribution hints at reciprocity, as it is positive. The same regression model does not fit well for the learning data and the coefficient of interest is not significantly different from zero. These findings provide strong evidence against the claim that learning of subjects initial in the state of ignorant confusion can explain the typical dynamics in public goods games. In other words, dynamics typically interpreted as stemming (at least partly) from the breakdown of cooperation due to conditional cooperators withdrawing their concern for free-riders is not an artefact of simple reinforcement

Table 2: Dynamic Panel Estimation of contributions in the Learning Condition

Variable	Coefficient	(Std. Err.)
<i>Lagged own contribution</i>		
c_{t-1}^i	-0.032	(0.107)
c_{t-2}^i	-0.032	(0.085)
<i>Lagged average contribution of others</i>		
\bar{c}_{t-1}^{-i}	-0.039	(0.067)
\bar{c}_{t-2}^{-i}	-0.029	(0.078)
<i>Period dummies, t = 3 omitted</i>		
$t = 4$	-1.156	(0.844)
$t = 5$	-4.200**	(1.130)
$t = 6$	-2.718	(1.622)
$t = 7$	-2.161	(1.348)
$t = 8$	-1.946	(1.270)
$t = 9$	-3.026*	(1.209)
$t = 10$	-2.783*	(1.378)
$t = 11$	-3.067*	(1.226)
$t = 12$	-3.152**	(1.176)
$t = 13$	-5.145**	(1.381)
$t = 14$	-3.690*	(1.734)
$t = 15$	-4.082*	(1.757)
$t = 16$	-4.048*	(1.686)
$t = 17$	-3.737*	(1.691)
$t = 18$	-3.188	(1.939)
$t = 19$	-4.816**	(1.752)
$t = 20$	-3.884*	(1.639)
<i>Intercept</i>	12.328**	(3.811)
N		1080
$prob > \chi_{(21)}^2$		0.000

learning by confused subjects.

3.3 Individual analysis and different types of subjects

The analysis so far has been on an aggregate (or average) level. In this section we want to dig a bit deeper and look at individual behavior. A last test of whether one could wrongly attribute the contribution dynamics to conditional cooperation explicitly takes into account individual heterogeneity with respect to social preferences and learning by exploiting the within-subject variation between phase one (Learning Condition) and phase two (Standard Condition). Appendix B contains a table indicating whether a subject’s contributions are significantly correlated with past contributions of others within a treatment (Spearman rank-correlation coefficient, 5-percent level).

In the Learning Condition there are only four out of 60 subjects (6.6 percent) who exhibit a significantly positive correlation, whereas in the Standard Condition 32 of 60 (53.3 percent) do so. The subjects identified as conditional cooperators due to positive correlation in the standard condition are therefore unlikely to be just confused subjects who learn. Only 12.5 percent of the subjects (four out of 32) showing positive correlation in the standard condition also show a positive correlation in the Learning Condition. An individual, who behaves in a way consistent with conditional cooperation in the standard condition, typically does not show that same behavior in the Learning Condition. This is more evidence that learning behavior (in the sense of reduced ignorance) does not fully explain the decay in public goods games, which leaves room for conditional cooperation as an important factor.

The fact that in the standard condition not all but only about 53.3 percent of subjects exhibit a positive correlation of their contributions with the past contributions made by their group members shows that there is quite some heterogeneity. While these subjects conform with the criterion for conditional cooperation, the others do not. In the previous section we estimated average contribution dynamics. It is instructive to also separately estimate the contribution dynamics of the two sub-populations.

Table 3 shows the regression results for the sub-populations of the “conditional cooperators” and the “others” in the standard public goods game. The dynamics of the conditional cooperators are as expected. These subjects react strongly to the average lagged contribution of the other group members. By design, picking out only the subjects, who exhibit positive correlation between c_t^i and \bar{c}_{t-1}^i , increases the estimated coefficient (here from 0.346 to 0.496). Conditional cooperation is still not “perfect” though. Three group members, who increase their contribution by one monetary unit on average only induce a conditional cooperator (the fourth group member) to increase her contribution by half a unit. Interestingly, the conditional cooperators do not exhibit significant persistence.

Table 3: Estimation of contributions in the standard condition by subject type

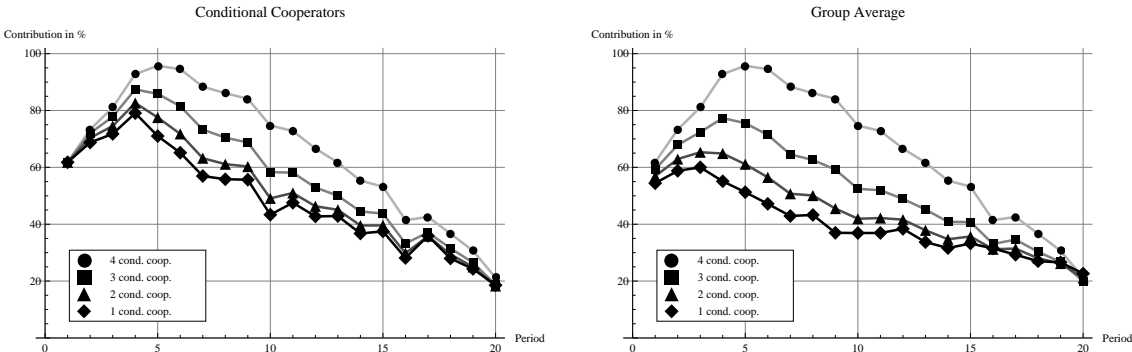
Variable	Cond. Coop.		Others	
	Coef.	(S. E.)	Coef.	(S. E.)
<i>Lagged own contribution</i>				
c_{t-1}^i	0.185	(0.198)	0.181	(0.096)
c_{t-2}^i	-0.127	(0.202)	-0.021	(0.086)
<i>Lagged average contribution of others</i>				
\bar{c}_{t-1}^{-i}	0.496**	(0.129)	0.074	(0.128)
\bar{c}_{t-2}^{-i}	-0.052	(0.226)	-0.068	(0.106)
<i>Period dummies, t=3 omitted</i>				
$t = 4$	1.290	(1.291)	-1.819*	(0.776)
$t = 5$	1.292	(1.330)	-1.979**	(0.688)
$t = 6$	-0.339	(1.122)	-2.592**	(0.859)
$t = 7$	-1.413	(1.567)	-3.062**	(0.684)
$t = 8$	-1.041	(1.950)	-2.783*	(1.092)
$t = 9$	-1.116	(2.470)	-4.478**	(1.343)
$t = 10$	-2.745	(2.863)	-3.372*	(1.417)
$t = 11$	-1.846	(3.336)	-3.795*	(1.615)
$t = 12$	-2.824	(3.399)	-3.038*	(1.530)
$t = 13$	-2.979	(3.287)	-4.416**	(1.451)
$t = 14$	-3.583	(3.788)	-4.342**	(1.600)
$t = 15$	-3.139	(4.244)	-3.933**	(1.404)
$t = 16$	-5.221	(4.322)	-3.844*	(1.874)
$t = 17$	-3.396	(4.814)	-4.982*	(2.008)
$t = 18$	-4.726	(4.821)	-4.857*	(2.004)
$t = 19$	-5.134	(4.669)	-4.703*	(2.140)
$t = 20$	-6.211	(5.278)	-5.403*	(2.121)
<i>Intercept</i>	6.307	(7.040)	9.212**	(3.347)
N	738		990	
$prob > \chi_{(21)}^2$	0.000		0.000	

Furthermore, a group which consists entirely of “conditional cooperators” is not able to sustain full cooperation. To see this we used the estimated dynamics (including the period effects) and simulated contributions for a group of four identical “conditional co-

operators”. As starting values we used the empirical average first-period contributions of the conditional cooperators, which were around 61.5 percent. The time series identified by circles in Figure 4 depicts the resulting contributions. In early periods groups are able to build considerable cooperation. In periods five to eleven cooperation moves close to 100 percent. Then negative exogenous period effects start to destroy the achieved level of cooperation. Contributions fall by about seven to eight percentage points per period. The observation that even subjects who are deemed to be “conditional cooperators” are not able to sustain full cooperation is consistent with the recent study by Fischbacher and Gaechter (2009). The existence of free-riders is not necessary for contributions to decline eventually.

Figure 4 also shows the simulated contributions of “conditional cooperators” depending on how many conditional cooperators are contained in a group. The simulation is again based on the dynamics previously estimated. For the subjects classified as “others” the starting contribution is taken from the empirical distribution (around 52 percent of the endowment). It is apparent that the first non-cooperator in a group has a very strong influence. A single non-cooperator prevents a group from reaching full cooperation temporarily. Facing one non-cooperator, cooperators reduce their contributions already after period five. The existence of a free-rider accelerates the decline of contributions from “conditional cooperators”. Adding additional non-cooperators does not change the dynamics much but slightly shifts the cooperators’ contributions downwards.

Figure 4: Contribution dynamics in mixed groups (conditional cooperators and group averages)



Finally, we investigate if the subjects we called “others” exhibit contribution dynamics that could stem from the reduction of ignorant confusion. Recall that subjects that were

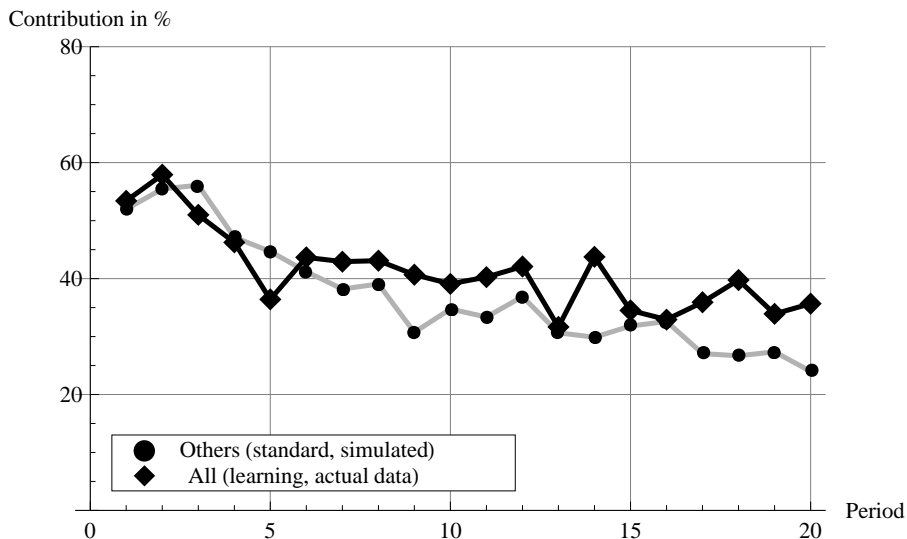
not identified as conditional cooperators do not react to the past contributions of the other group members. The same is true for the subjects in the Learning Condition (see Table 2). Under some (very) strong assumptions we could hypothesize that the contributions of the two groups should follow the same dynamics. If (a) the dynamics in our learning treatment capture all learning dynamics (even though the information environment is quite different to the standard treatment), (b) the subjects in the category “others” have no other motivation than to maximize their payoff and if (c) the “others” do not expect to be able to strategically manipulate the contributions of potentially existing conditional cooperators then the dynamics should be the same. We do not really expect all these assumptions to hold. However, not too dissimilar dynamics would give us some confidence that our Learning Condition actually captures something comparable to learning dynamics in the standard public goods game.

Figure 5 plots the two contribution time series (the estimated dynamics for “others” and the actual average contributions in the Learning Condition). Contributions of “others” in the Standard Condition fall below the average contributions in the Learning Condition, after starting off almost identically. Both curves are relatively flat if one compares them to the contributions of the “conditional cooperators” from Figure 4. So on average the time series show similarities. However, recall that on an individual level we have found some differences: in the Learning Condition contributions show no significant persistence (the coefficient on the own lagged contribution is positive but not significant), while the “others” in the standard condition exhibit slight persistence (only significant at the 10 percent level though) in their contributions. We conclude that our Learning Condition (as expected) cannot capture all the learning and other potentially relevant behavioral factors beyond conditional cooperation at work in standard public goods games. Nevertheless, the similarities of the time series (as described above) give us some confidence that the dynamics in the Learning Condition are related to the unobserved learning dynamics in standard public goods games.

4 Conclusion

In this paper we report on a novel experiment designed to identify the influence of confusion on the dynamics in repeated public goods games. In contrast to previous studies, we study confusion in a benchmark treatment by withholding information on the structure of the game (instead of treating confusion as a residual) and compare the resulting contributions to those in a standard public goods game. We argue that this approach can provide a lower bound for the role confusion plays in reducing contributions over

Figure 5: Contribution dynamics of “others” in the standard condition vs. all subjects in the Learning Condition



time. We find that the reduction of confusion causes a decay in contributions over time. However, learning only accounts for 41 percent of the total decay. Furthermore, we could not find evidence for the reduction of confusion (due to reinforcement learning) producing correlation patterns of contributions within groups that could be wrongly attributed to conditional-cooperation behavior. Consequently, we conclude that part of the dynamics in public goods games is due to conditional cooperators reacting to low contributions of other group members. This claim is backed by both within-subject comparisons of behavior across treatments and dynamic panel analysis allowing for heterogeneity with respect to types of subjects. In line with findings from other studies, conditional cooperation is not sufficient to prevent the decline of contributions even in groups where all subjects can be classified as conditional cooperators. Finally, we observe similar contribution time series for subjects who could not be classified as conditional cooperators in the Standard Condition and for subjects in the Learning Condition. This provides some support for our underlying assumption that the sort of confusion we have considered in our experiment generates learning dynamics that are also relevant in standard public goods games.

References

- [1] **Andreoni, James**, 1988. “Why Free Ride? Strategies and Learning in Public Goods Experiments.” *Journal of Public Economics* 37(3): 291-304.
- [2] **Andreoni, James**, 1995. “Cooperation in Public-Goods Experiments: Kindness or Confusion?” *American Economic Review* 85(4): 891-904.
- [3] **Arellano, Manuel and Steven Bond**, 1991. “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations.” *Review of Economic Studies* 58: 277-297.
- [4] **Arellano, Manuel and Olympia Bover**, 1995. “Another Look at the Instrumental Variable Estimation of Error-Components Models.” *Journal of Econometrics* 68: 29-51.
- [5] **Blundell, Richard and Steven Bond**, 1998. “Initial Conditions and Moment Restrictions in Dynamic Panel-Data Models.” *Journal of Econometrics* 87: 115-143.
- [6] **Camerer, Colin F.**, 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- [7] **Camerer, Colin F. and Ernst Fehr**, 2003. “Measuring Social Norms and Preferences Using Experimental Games: A Guide for Social Scientists.” In: J. Henrich, R. Boyd, S. Bowles, H. Gintis, E. Fehr, R. McElreath (Eds.): *Foundations of Human Sociality: Experimental and Ethnographic Evidence from 15 Small-scale Societies*. Oxford University Press, in press.
- [8] **Camerer, Colin F. and Teck-Hua Ho**, 1999. “Experience-Weighted Attraction Learning in Normal Form Games.” *Econometrica* 67: 827–874.
- [9] **Chen, Yan and Yuri Khoroshilov**, 2003. “Learning under Limited Information”. *Games and Economic Behavior* 44: 1–25.
- [10] **Cookson, Richard**, 2000. “Framing Effects in Public Goods Experiments.” *Experimental Economics* 3: 55–79.
- [11] **Croson, Rachel**, 1996. “Partners and Strangers Revisited.” *Economics Letters* 53: 25-32.

- [12] **Duersch, Peter; Albert Kolb; Jörg Oechssler; and Burkhard C. Schipper**, 2009. “Rage Against the Machines: How Subjects Play against Learning Algorithms.” *Economic Theory*: DOI 10.1007/s00199-009-0446-0.
- [13] **Ferraro, Paul J. and Christian A. Vossler**, 2006. “Stylized Facts and Identification in Public Goods Experiments: The Confusion Confound.” Working Paper, Georgia State University.
- [14] **Fischbacher, Urs**, 2007. “Z-Tree: Zurich Toolbox for Readymade Economic Experiments.” *Experimental Economics* 10(2): 171-178.
- [15] **Fischbacher, Urs and Simon Gächter**, 2009. “Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Good Experiments.” *American Economic Review*, in press.
- [16] **Gintis, Herbert; Samuel Bowles; Richard Boyd; and Ersnt Fehr**, 2003. “Explaining Altruistic Behavior in Humans.” *Evolution and Human Behavior* 24(3): 153–172.
- [17] **Houser, Daniel and Robert O. Kurzban**, 2002. “Revisiting Kindness and Confusion in Public Goods Experiments.” *American Economic Review* 92(4): 1062-1069.
- [18] **Janssen, Marco A. and Toh-Kyeong Ahn**, 2006. “Learning, Signaling, and Social Preferences in Public-Good Games.” *Ecology and Society* 11(2): Art. 21.
- [19] **Keser, Claudia and Frans van Winden**, 2000. “Conditional Cooperation and Voluntary Contributions to Public Goods.” *Scandinavian Journal of Economics* 102: 23-39.
- [20] **Kurzban, Robert O. and Daniel Houser**, 2005. “An Experimental Investigation of Cooperative Types in Human Groups: A Complement to Evolutionary Theory and Simulations.” *Proceedings of the National Academy of Sciences* 102(5): 1803-1807.
- [21] **Ledyard, John O.**, 1995. “Public Goods: A Survey of Experimental Research.” In: J.H. Kagel and A.E. Roth (Eds.): *The Handbook of Experimental Economics*. Princeton Univ Press, Princeton: 111-194.
- [22] **Levitt, Steven D. and John A. List**, 2007: “What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?” *Journal of Economic Perspectives* 21(2): 153-174.

- [23] **Mookherjee, Dilip and Barry Sopher**, 1994. “Learning Behavior in an Experimental Matching Pennies Game.” *Games and Economic Behavior* 7: 62–91.
- [24] **Muller, Laurent; Martin Sefton; Richard Steinberg; and Lise Vesterlund**, 2008. “Strategic Behavior and Learning in Repeated Voluntary Contribution Experiments.” *Journal of Economic Behavior and Organization* 67(3-4): 782-793.
- [25] **Sarin, Rajiv and Farshid Vahid**, 2004. “Strategy Similarity and Coordination.” *Economic Journal* 114: 506-527.
- [26] **Sarin, Rajiv and Farshid Vahid**, 1999. “Payoff Assessments Without Probabilities: A Simple Dynamic Model of Choice.” *Games and Economic Behavior* 28: 294–309.
- [27] **Van Huyck, John B.; Raymond C. Battalio; and Frederik W. Rankin**, 2007. “Selection Dynamics and Adaptive Behavior Without Much Information.” *Economic Theory* 33(1): 53-65.

A Instructions

A.1 Instructions to subjects for the Learning Condition

General Information

You are participating in an experiment on learning behaviour

Money

In the experiment you earn points, which will be converted into real money at the end of the experiment with the following exchange rate:

100 points = AUD 1.25.

Your earnings are paid in cash at the end of the experiment

Please note

It is strictly forbidden to communicate with other participants during the experiment. You are not allowed to speak with other participants.

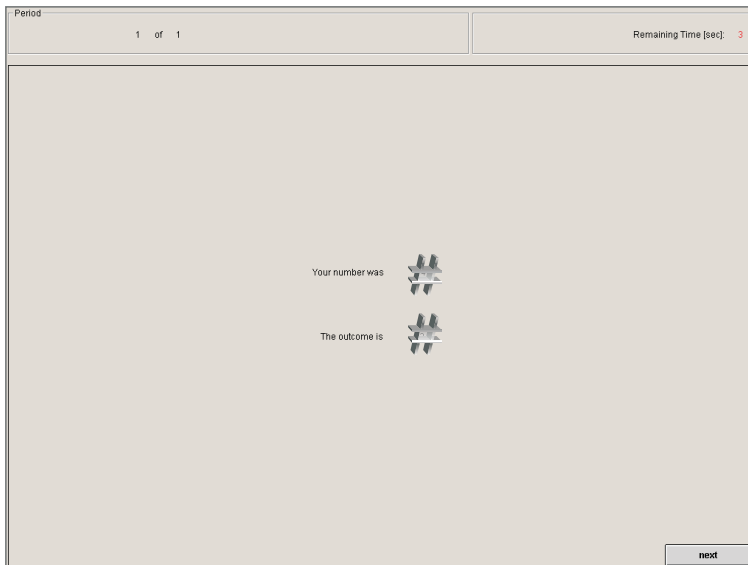
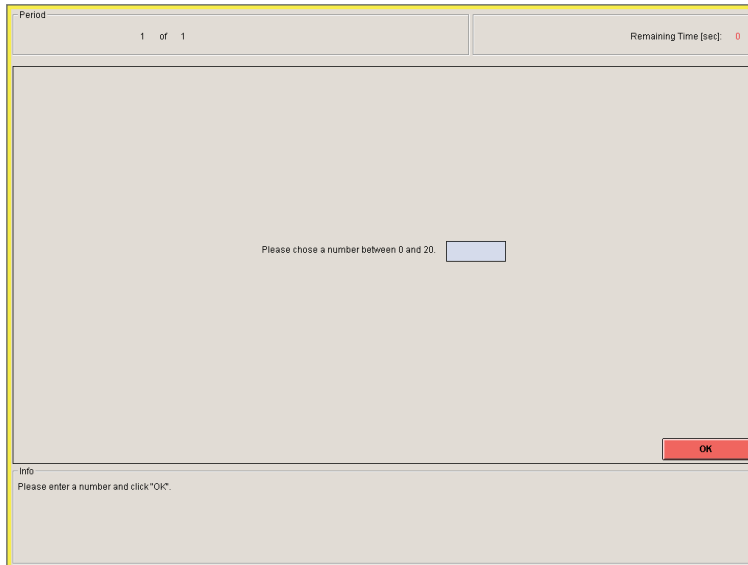
If you have questions during the experiment please raise your arm and somebody will come and help you.

The timing of the experiment

The proceedings for the 20 periods are:

At the beginning of each period you see the following screen:

In every period you just have to make a single decision. You simply have to choose a number between 0 and 20. After you have entered your number and have clicked “OK” you will see the following result screen:



Here you can see your previous decision and the number of points you receive.

How is the outcome determined?

Your income depends on the number you have entered. However, other factors may influence your income. These factors may change from period to period. This means that a certain number you choose does not lead to the same outcome all the time.

Are there any further questions?

A.2 Instructions to subjects for the standard condition and for the control treatment.

Experimental Instructions

Thank you for participating in the experiment. If you read these instructions carefully and follow all rules, you can earn money. The money will be paid out to you in cash immediately after the experiment. During the experiment we shall not speak of Dollars but rather of points. Points are converted to Dollars at the following exchange rate:

$$100 \text{ Points} = \text{AUD } 1.25$$

It is forbidden to speak to other participants during the experiment. If you have any question, please ask us. We will gladly answer your questions individually. It is very important that you follow this rule. Otherwise, we shall have to exclude you from the experiment and from all payments.

Participants of this experiments are randomly assigned into groups of 4 members, i.e., there are three more persons forming a group together with you. The composition of groups will remain the same during the whole experiment, i.e. there will always be the same persons in your group. The identity of your group members will not be revealed to you at any time. At the start of each period, each participant gets 20 points. We will refer to these points as your endowment. Your task is it to decide, how many of your 20 points you want to contribute to a project or to keep for yourself.

Your income consists of two parts:

1. Points that you keep
2. Your “income from the project”. This income is calculated as follows:

Your income from the project = $0.4 \times$ Sum of contributions of all group members to the project

The income of the other members of your group is determined in the same way, i.e. each group member receives the same income from the project. Suppose, for example, that the total contributions to the project by all members in your group sum up to 60. In this case you and every other member of your group receives $0.4 \times 60 = 24$ points as income from the project. Suppose that you

and the other 3 members of your group in total contribute only 10 points to the project. In this case every group member receives $0.4 \times 10 = 4$ points as income from the project.

For each point that you keep for yourself you earn an income of one point. If you contribute that point to the project, instead, the sum of contributions to the project would rise by one point, and your income from the project would rise by $0.4 \times 1 = 0.4$ points. However, the income of the other group members would also rise by 0.4 points, such that the total income of the group would rise by $4 \times 0.4 = 1.6$ points. Your contribution to the project, therefore, raises the income of the other members of your group. On the other hand, you earn from each point that other members of your group contribute to the project. For each point that another group member contributes, you earn $0.4 \times 1 = 0.4$ points.

You take your decision via the computer. After all participants have made their contributions a new period starts, in which you decide again how many of your 20 points you want to contribute to the project. In total there will be 20 periods.

B Within-subject comparison

Individual correlation between own and others' contributions (learning and standard condition): Spearman rank-correlation coefficient; positive and significant (+); negative and significant (-); insignificant (0); $\alpha = 0.05$.

Subject	Learning	Standard	Subject	Learning	Standard
1	0	0	31	0	+
2	0	-	32	0	+
3	0	0	33	0	+
4	0	0	34	0	+
5	+	+	35	0	0
6	+	+	36	0	+
7	+	+	37	0	+
8	0	0	38	0	+
9	0	0	39	0	0
10	0	0	40	0	+
11	0	0	41	0	+
12	0	0	42	0	0
13	0	0	43	0	0
14	0	+	44	0	+
15	0	+	45	0	+
16	0	+	46	0	0
17	0	+	47	0	0
18	0	+	48	0	0
19	0	+	49	0	+
20	0	+	50	0	0
21	0	+	51	0	0
22	+	+	52	0	0
23	0	+	53	0	+
24	0	0	54	0	+
25	0	0	55	0	0
26	0	0	56	0	+
27	0	0	57	0	0
28	0	+	58	0	0
29	0	+	59	0	+
30	0	+	60	0	0

C Contribution dynamics in groups

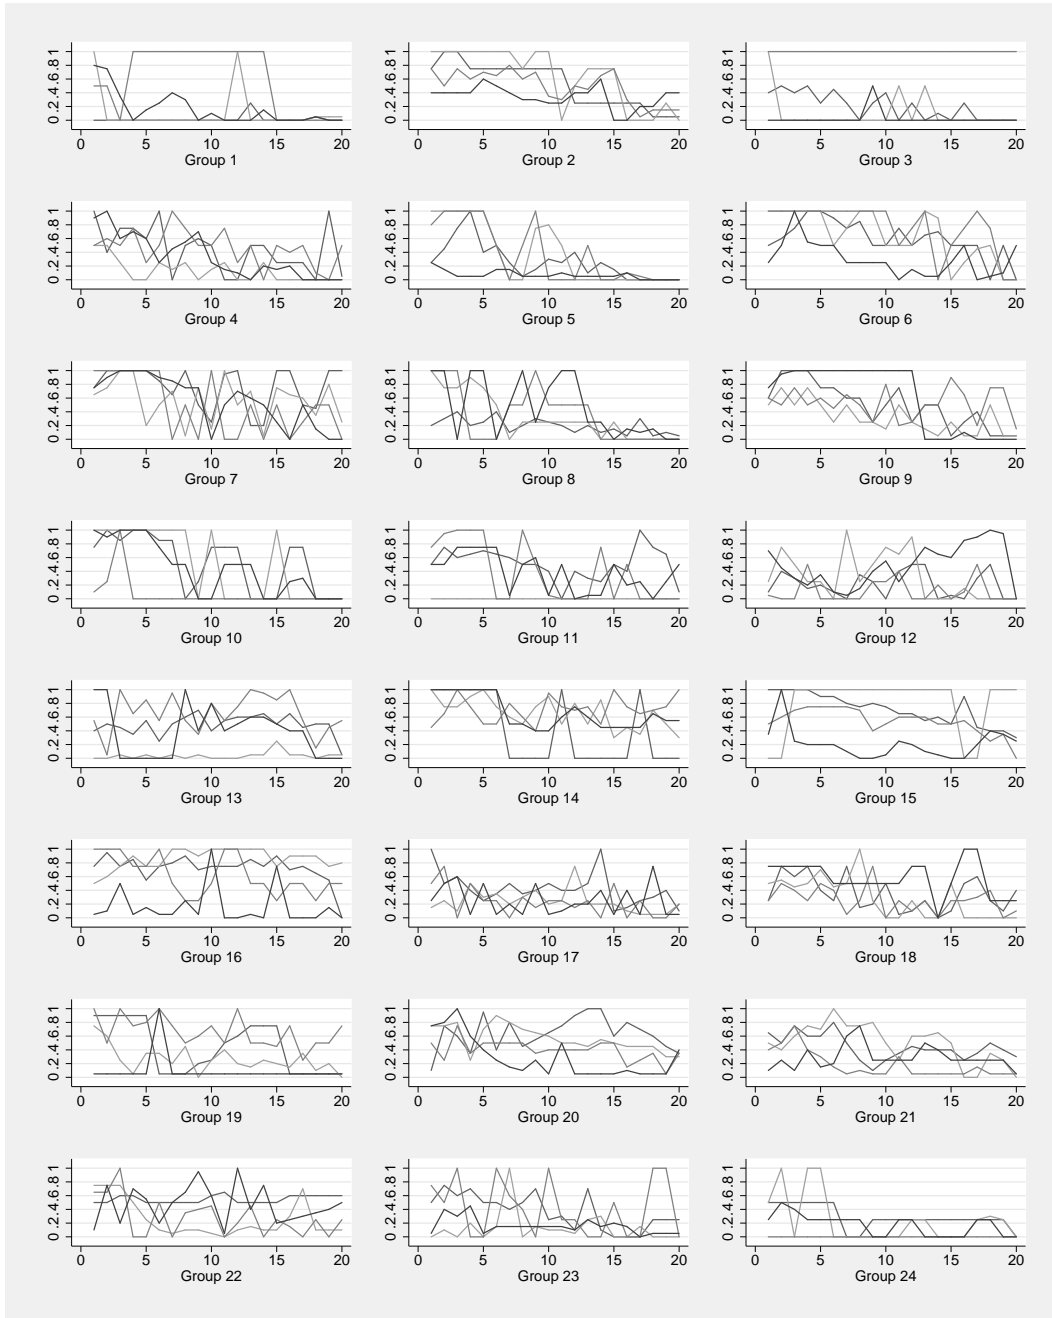


Figure 6: Contributions in groups over time in the standard condition

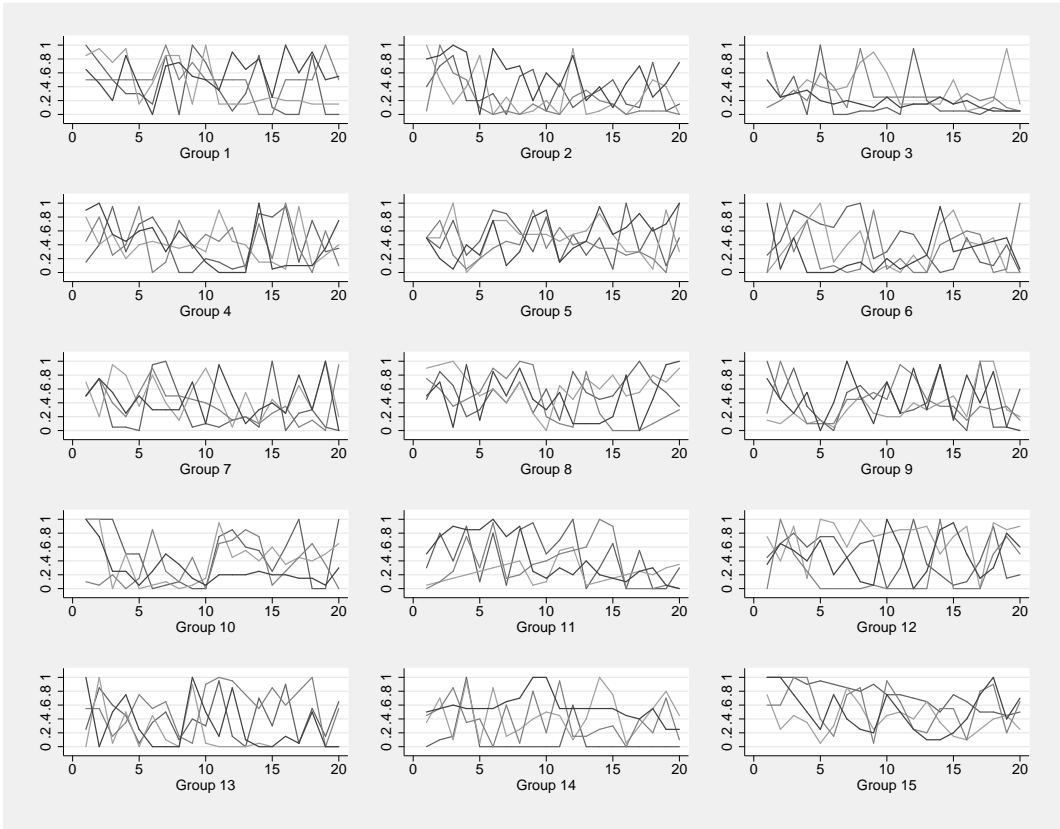


Figure 7: Contributions in groups over time in the Learning Condition