Revisiting Kindness and Confusion in Public Goods Experiments: Comment*

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

February 2008

Abstract

We present a novel experiment to evaluate the role of confusion in explaining the dynamics of contributions in public goods experiments. In a learning condition we keep the subjects confused by withholding the information that would be necessary for them to understand the game. The comparison of the results with those of a standard public goods treatment provides a lower bound for the influence of confusion on contribution dynamics. We find that learning in the state of confusion explains 41 percent of the contribution dynamics in the standard public goods game. This result complements that of Houser and Kurzban (2002) who found that all of the decrease in contributions can be attributed to the reduction of confusion. We argue that their findings can be seen as an upper bound for the influence of confusion.

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

JEL classification: C90, D83, H41.

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

†School of Economics, University of Adelaide, Napier Building, SA 5005 Adelaide; 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

§Faculty of Economics and Statistics, University of Innsbruck, Universitätsstr. 15, A-6020 Innsbruck. E-mail: rupert.sausgruber@uibk.ac.at

1 Introduction

"A natural explanation for the large rate of contribution in many voluntary contribution experiments can be found in the inexperience of the subjects. Perhaps, a 40 to 60 percent contribution rate occurs simply because if one must contribute a number between 0 and Z and does not understand the implication of the act, then a natural choice is somewhere in the middle. Clearly it is important to be able to discover whether the data are simply the result of confusion and inexperience or the result of some more purposeful behavior."
John O. Ledyard (1995: 146)

In a 2002 paper in this *Review*, Daniel Houser and Robert Kurzban (henceforth HK) report the results from an experiment aimed at discriminating between strategic behavior and confusion in public goods experiments. In their design, subjects participated in a "computer condition" in which all other group members were simulated by automata. A comparison to a standard public goods game was used to isolate the effect of confusion.[1] HK's findings were the following. The subjects on average contributed 28.6 percent of the endowment in the computer condition but 52.8 percent in a standard human condition. Off-equilibrium play in the computer condition cannot be associated with social motives towards other players and therefore has to be caused by confusion. So HK conclude that confusion accounts for 54 percent of all contributions to the public good. The finding that confusion accounts for a substantial amount of contributions is consistent with the existing literature (e.g. James Andreoni, 1995; Thomas R. Palfrey and Jeffrey E. Prisbey, 1996).

More interestingly, contributions in their computer condition fell with repetition at a rate much higher than in the human condition. HK conclude from this observation that the decay in contributions is entirely due to the reduction of confusion. This conclusion has very important implications. Given that HK find that the initial contributions in the human treatment are above the contributions in the computer treatment, a corollary of their conclusion would be that in the absence of confusion cooperation would be stable at an initial level. This view contradicts the conventional wisdom that the decay of contributions stems from the heterogeneity of preferences within a group, where the conditional cooperators start out with high contributions but consecutively adjust their contributions downwards as a reaction to selfish group members (see Andreoni, 1995; Kurzban and Houser, 2005; Laurent Muller et al., forthcoming; Urs Fischbacher and Simon Gächter,

[1] Paul J. Ferraro and Christian A. Vossler (2006) use a similar design for the same purpose and complement it with econometric modeling.

2006).[2] In contrast, HK's findings imply that cooperation unravels due to some confused subjects, while most other authors claim that cooperation unravels because of some selfish individuals.

HK state their underlying assumption, which is sufficient for their conclusion, to be that "cooperation due to confusion is similar in the human and computer conditions" (HK: 1066). In what follows, we analyze the implicit impact of this assumption and argue that HK's attribution of the decay in contributions to confusion can be regarded as an upper bound. We then propose an experimental approach that provides a lower bound for the impact of confusion on the decay of contributions to complement HK's findings.

In HK's computer treatment, the reduction of confusion can only lead to a reduction in contributions, as social motives are excluded. So their assumption of equivalent behavior of confused players in both the human and computer treatment implies that confused people in the human treatment learn selfishness independently of conditional cooperation. If individuals are allowed to exhibit some form of conditional cooperation, then the true decay of contributions by confused subjects is likely to be different in the human treatment than the decay recovered from the computer treatment. Further, their assumption implies that imitation learning has to be identical regardless of the identity (computer or human) of the imitated and whether the behavior to be imitated is explicitly given before a decision or just as information after a round.[3]

HK's assumption that confused subjects behave and learn exactly in the same way regardless of the treatment cannot be tested empirically. Hence, their finding that the decay in public goods games can be fully attributed to the reduction of confusion is to be regarded as an upper bound. Alternative designs based on other plausible assumptions might lead to different results.

As demonstrated, studies that use computer treatments for isolating the effects of confusion suffer from the need for auxiliary assumptions on how the behavior in the computer treatment translates into a human situation. In fact, all studies that try to isolate the dynamic effects of confusion have to make some assumptions. Given that the results and the assumptions made by HK provide a conclusion that gives an upper bound on the impact of confusion on the decay of contributions, we choose our design such that our results can be seen as a lower bound. Suppose confused subjects are much less sophisticated than assumed by HK. Assume that a confused subject neither understands

[2]Herbert Gintis et al. (2003) explain similar dynamics with an evolutionary approach.

[3]In HK's computer treatment the contributions of the computers were given before individuals had to decide.

the incentive structure, nor uses the contributions of other subjects (or computers) as a means of imitation learning.

Such an extreme initial assumption has the advantage that no further implicit assumptions about the reaction to computer contributions has to be made. Additionally, the dynamics of contributions of individuals, who are confused and ignorant in the sense explained above, do not have to be imputed from a treatment where social motives are excluded. Instead, confusion of this kind can be induced by withholding information about the payoff structure and the choices of other group members. So a treatment where subjects do not know either the incentive structure or the choices of the group members induces this kind of "ignorant" confusion. Subjects can only learn by reinforcement in such an environment. We run such a treatment and compare the contribution dynamics with those in a standard public goods game.

Inducing confusion with such a procedure instead of recovering it from a computer treatment 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). 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. The merit of a study inducing confusion as ignorance is that it complements HK's upper bound with a lower bound. As such, finding that the decay of contributions in our confusion treatment is still stronger than in the traditional public goods game would put HK's claim that confusion is responsible for all decay beyond doubt. The converse finding would imply that not all decay is necessarily caused by confusion. Such a finding would leave room for the conventional belief that (at least some of) the decay in contributions is due to the reaction of conditional cooperators to low contributions of selfish subjects.

Our results are as follows. At the aggregate level we find that contributions drop off in both the confusion condition and the standard public goods game. This observation supports the claim that learning leads to dynamics similar to those which are readily interpreted as conditional cooperation in standard public goods games. However, in contrast to HK, we 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 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. Our experiment provides direct evidence that the correlation of contributions with average past contributions of the other group members cannot be explained by the reduction of induced confusion.

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

2 Experimental design

Our design induces confusion by withholding information about the game from the subjects. They only know the admissible action space and that the environment may change over time. Over the repetitions of the game, the subjects are informed only about the payoff resulting from their last choice (see Appendix II for instructions).

We employ a within-subject comparison between two phases, which we complement by a treatment to control 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 and in the standard condition. A phase consists of 20 periods. In phase one, 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 between periods".

At the end of phase one, the subjects in the within-subject treatment are informed that a new experiment (phase two) will start. Only at this stage are they given instructions for a standard linear public goods game, where they are assigned to groups of four. 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. We will provide the details in the next section.

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 the Group Exchange pays 0.4 cents to every subject in the group, while every point kept pays 1 cent only to the subject. The underlying structure of phase two is identical to that of phase one except that the subjects did not know that they played a public goods game in phase one, while they did know in phase two. In what follows we refer to phase one as the learning condition and to phase two (and the control treatment) as the standard condition.

3 Results

We ran 5 sessions with between 16 and 20 subjects each. In total we had 15 groups (60 subjects) participating in the within-subject condition and 9 groups (36 subjects) in the control. 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 (Urs Fischbacher 2007). The experiment lasted between 25 (control treatment) and 35 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 (see the introductory quote by Ledyard 1995), the average contributions start out around the midpoint of the admissible action space. With repeated interaction, however, contributions show a 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 regression) contributions drop by -0.18 units per period. This negative time trend is significant at the 1-percent level.[4]

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. The model is in the tradition of reinforcement learning. Individuals decide over their choices by comparing the payoffs resulting from their last two choices.[5]

The choice rule 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

[4]We use robust standard errors adjusted for clustering on groups throughout this paper.

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

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. A more detailed description of the model can be found in Appendix I.

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. We see that a model as simple as ours already does very well at tracking the behavior in the learning condition. Hence, we feel confident to conclude that our learning condition isolated learning dynamics from the dynamics generated by strategic behavior of any kind.

Figure 1

Time series of average contributions in the learning condition (phase one)

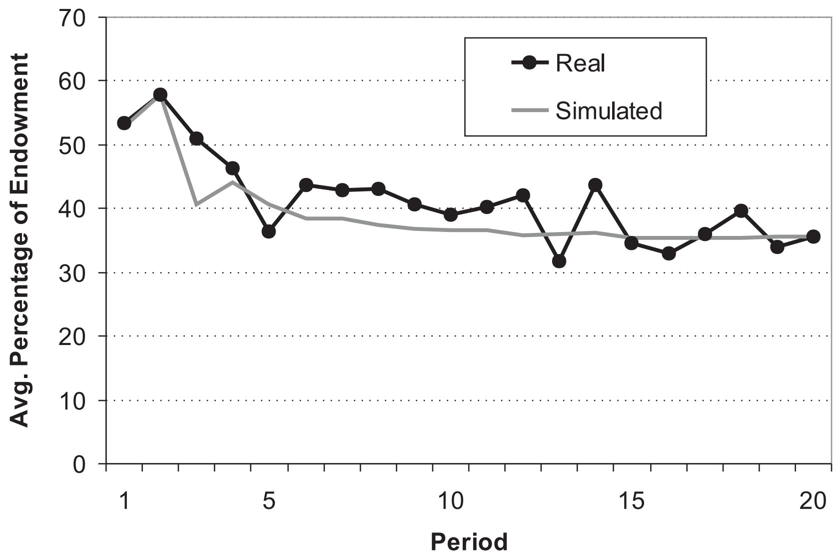

Our next step is to quantify how much of the downward trend of contributions in the standard public goods game (phase two and the control treatment) is due to the reduction of confusion.[6] In Figure 2, the black line plots the average contributions of the subjects in the standard condition. Obviously, the rate at which subjects reduce their contributions is greater than in phase one. In the standard condition, contributions drop from 57.7 percent of the total endowment in period 1 to 16.6 percent in period 20.

[6]We pool the data from phase two and the control treatment, as we could not find significant differences. We shortly discuss this when we present our main regression below.

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 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 ($p < 0.01, N = 15$).[7]

The analysis above shows that a reinforcement-learning model explains the dynamics in the learning condition, while it fails to explain all 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.[8] 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 phase one).

To summarize: Assuming that confusion in the human condition 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. This result is in contrast to the 100 percent contribution of learning found by HK, which we consider an upper bound.

Given our results, the lower bound approach of the influence of confusion on contribution dynamics does not rule out the conventional belief that (at least some of) the decay observed in public goods games is due to conditional cooperation. So we conduct further tests and compare the correlation between current contributions and past contributions of other group members across treatments. This will show if reinforcement learning can generate correlation patterns, which are usually regarded as evidence for conditional cooperation. For this purpose, we run the following regression:

$$c_t^i = \beta_0 + \beta_1 c_{t-1}^i + \beta_2 \overline{c}_{t-1}^{-i} + \beta_4 t + \varepsilon_t^i,$$

where c_t^i is the subject's current contribution, $\overline{c}_{t-1}^{-i}$ is the average declaration of the other

[7]The average mean square error of the simulation is more than four times larger in phase two (3.66 vs. 0.82 points).

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

Figure 2

Time series of average contributions in the standard condition

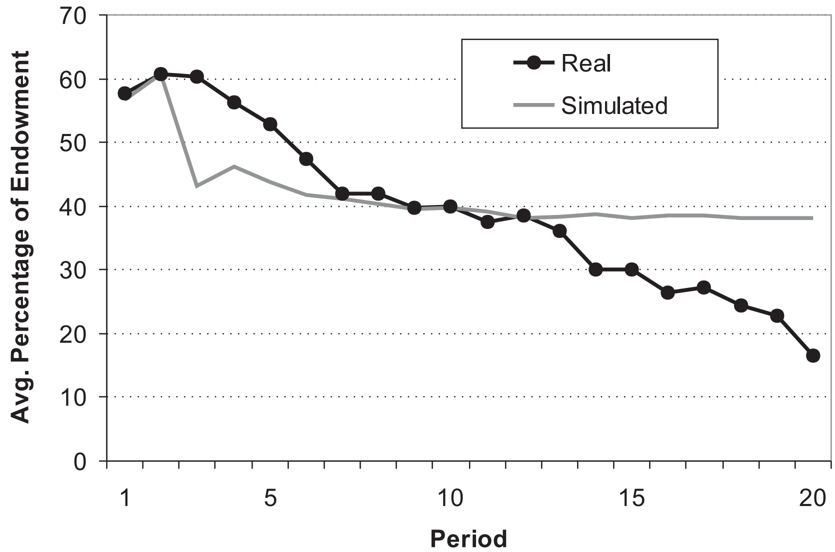

group members in the past period, t is a time trend, and ε_t^i denotes the error term. Previous research has taken a positively significant β_2 in the standard public goods game as evidence for social preferences in the guise of conditional cooperation, as it indicates that variation in the past contribution of the other group members influences the own contribution in the same direction (e.g., Rachel Croson 1998). Initially, we ran this regression separately for phase one, phase two and the control treatment. Then we combined the variance-covariance matrices using seemingly unrelated regression, allowing for clustering on groups and followed by a Chow test establishing if the coefficients are jointly different among the three treatments. We found that the coefficients from the learning condition are significantly different to those from both of the treatments where a standard public goods game was played ($p < 0.01$), while we could not find a significant difference between the two public goods treatments (phase two and the control treatment, $p > 0.23$). We conclude that a) behavior in the learning condition differs from that in the public goods games, while b) behavior in the public goods games are not influenced by the experience of having played the learning condition before. Below we present regressions where we pool the data from the standard public goods games and test if the estimate of β_2 differs across the learning and standard condition.

Table 1 shows the results. The p-values in parentheses are calculated from robust standard errors accounting for clustering on groups. For the standard condition the esti-

Table 1 OLS regression with robust standard errors. Dependent variable: c_t^i

	Standard	Learning
const	3.05	7.35
	(< 0.01)	(< 0.01)
c_{t-1}^i	0.58	0.21
	(< 0.01)	(< 0.01)
$\overline{c}_{t-1}^{-i}$	0.16	0.04
	(< 0.01)	(> 0.43)
t	-0.14	-0.13
	(< 0.01)	(< 0.01)
	$N = 1824$	$N = 1140$
	F(3,23)=458.21	$F(3, 14) = 21.13$
	$r^2 = 0.48$	$r^2 = 0.07$

mated coefficient for β_2 is significantly positive, which is consistent with previous studies. Without further control, it is not possible to decide if the positive β_2 is due to conditional cooperation, learning, or a compound of both. We employ the same regression for the data generated in the learning condition in order to isolate the impact of confusion. We find that there the estimate of β_2 is insignificant. Given our assumption relating to the nature of confusion, we can conclude that the serial correlation in the standard public goods game cannot be explained by the reduction of confusion, as reinforcement learning cannot reproduce the correlation structure usually taken as evidence for conditional cooperation.

A further interesting result of our regression is that the residual time trend in the standard condition after controlling for serial correlation is very close to that in the learning condition. This suggests that the decay that cannot be attributed to conditional cooperation is similar in both treatments. If our assumptions on confusion are correct then this decay represents the common influence of the reduction of confusion in both treatments.[9]

A last test explicitly takes into account individual heterogeneity with respect to social preferences and learning by exploiting the within-subject variation between phase one and

[9] Note also that for the learning condition the regression does not explain much of the variance ($r^2 = 0.07$), while it fits the data well ($r^2 = 0.45$) for the standard condition.

phase two. Appendix III contains a table indicating whether a subject's contributions are significantly correlated with past contributions of other group members for both treatments (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.

4 Conclusion

In this paper we presented a novel experiment to identify the influence of confusion on the dynamics in public goods experiments. Our study complements the work of Houser and Kurzban (2002) who suggested that 100 percent of the decay in contributions is due to the reduction of confusion. They used a learning friendly full-information individual-choice problem to isolate the effects of confusion by framing it as a public goods game. We argue that their findings can be considered an upper bound. In contrast, we induce confusion by withholding information, which does not remove the interactive structure of a public goods game but makes it much harder to learn. Our results can therefore be interpreted as a lower bound of the impact of confusion. We find that the reduction of confusion causes a decay in contributions over time, as did Houser and Kurzban. However, learning only accounts for 41 percent of the total decay. Furthermore, we could not find evidence for the reduction of confusion producing correlation patterns of contributions within groups that could be wrongly attributed to conditional-cooperation behavior. Consequently, we believe part of the dynamics in public goods games is due to conditional cooperators reacting to the low contributions of other group members.

References

Andreoni, James. 1988. "Why Free Ride? Strategies and Learning in Public Goods Experiments." *Journal of Public Economics*, December 1988, 37(3), pp. 291-304.

Andreoni, James. 1995."Cooperation in Public-Goods Experiments: Kindness or Confusion?" *American Economic Review*, September 85(4), pp. 891-904.

Croson, Rachel T. A. 1998. “Theories of Altruism and Reciprocity: Evidence from Linear Public Goods Games.” Working Paper, The Wharton School, University of Pennsylvania.

Ferraro, Paul J. and Christian A. Vossler 2006. “Stylized Facts and Identification in Public Goods Experiments: The Confusion Confound.” Working Paper, Georgia State University.

Fischbacher, Urs. 2007. “Z-Tree: Zurich Toolbox for Readymade Economic Experiments.” Experimental Economics 10(2), pp. 171-8.

Fischbacher, Urs, and Simon Gächter. 2006. “Heterogenous Social Preferences and the Dynamics of Free Riding in Public Goods.” CeDex Discussion Paper No. 2006-01.

Gintis, Herbert, Samuel Bowles, Robert Boyd, and Ernst Fehr. 2003. “Explaining Altruistic Behavior in Humans.” *Evolution and Human Behavior*, 24(3), pp. 153–72.

Houser, Daniel, and Robert Kurzban. 2002. “Revisiting Kindness and Confusion in Public Goods Experiments.” *American Economic Review*, September, 92(4), pp. 1062-9.

Kurzban, Robert, 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), pp. 1803-7.

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, pp. 111-94.

Muller, Laurent, Martin Sefton, Richard Steinberg, and Lise Vesterlund. forthcoming. “Strategic Behavior and Learning in Repeated Voluntary-Contribution Experiments.” *Journal of Economic Behavior and Organization.*

Palfrey, Thomas R., and Jeffery E. Prosbey. 1996. “Anomalous Behavior in Public Goods Experiments: How Much and Why?” *American Economic Review*, August, 61(2), pp. 829-46.

Sarin, Rajiv, and Farshid Vahid. 2004. “Strategy Similarity and Coordination.” *Economic Journal* 114, pp. 506-27.

Appendix I: The Learning Model

Denote the contribution of person $i \in \{1,2,3,4\}$ in period $t \in \{1,2,..,20\}$ as c_t^i. The individual uses the payoff p and the own choices of the last two periods to determine the contribution in the current period (if possible). So

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

After having observed the two last outcomes given their choices made, 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 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

$$\begin{aligned} \Delta p_t^i &\equiv p_{t-2}^i - p_{t-1}^i \\ \Delta c_t^i &\equiv c_{t-2}^i - c_{t-1}^i \end{aligned}$$

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

$$d_t^i = sign(\Delta p_t^i \cdot \Delta c_t^i).$$

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 then we have:

$$C_t^i = \begin{cases} \left[0, \frac{c_{t-1}^i + c_{t-2}^i}{2}\right] & if \quad d_t^i = -1 \\ \left[\frac{c_{t-1}^i + c_{t-2}^i}{2}, 20\right] & if \quad d_t^i = 1 \\ [0,20] & if \quad d_t^i = 0 \end{cases}$$

Next we have to specify which point within the admissible range will be chosen. The simplest assumption is that the individual draws from a uniform distribution with support C_t^i:

$$c_t^i \sim U \text{ on } C_t^i$$

Analyzing the data we found that the median of choices for both experimental conditions is approximately in the middle of the support, which is the case under a uniform distribution. However, we observed quite a few focal points (bottom or top of the range), which cannot be modeled with the uniform distribution. The differences mainly occurred in the full-information treatment.

The remaining question is the choice behavior of the individual in periods 1 and 2. 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.

Basically, the model assumes that individuals believe in a "linear" world. This is implemented by the assumption that individuals choose values closer to the values that have been successful in the past. The belief that the world is linear is correct here as the underlying public goods game is indeed linear in the own choice. Another reason why we chose this learning rule is that it is consistent with the world the subjects know in the treatment with full information. Hence, such a learning rule should be most suited to explain choice behavior in the full information treatment if reciprocity motives are absent. Other more sophisticated beliefs about the underlying structure (like a quadratic relationship between choice and payoff) could lead to experimenting, where people choose numbers which are the extreme opposite of the previously more successful choice.

Appendix II:

1. Instructions to subjects for phase one.

General Information

You are participating in an experiment on learning behaviour. 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:

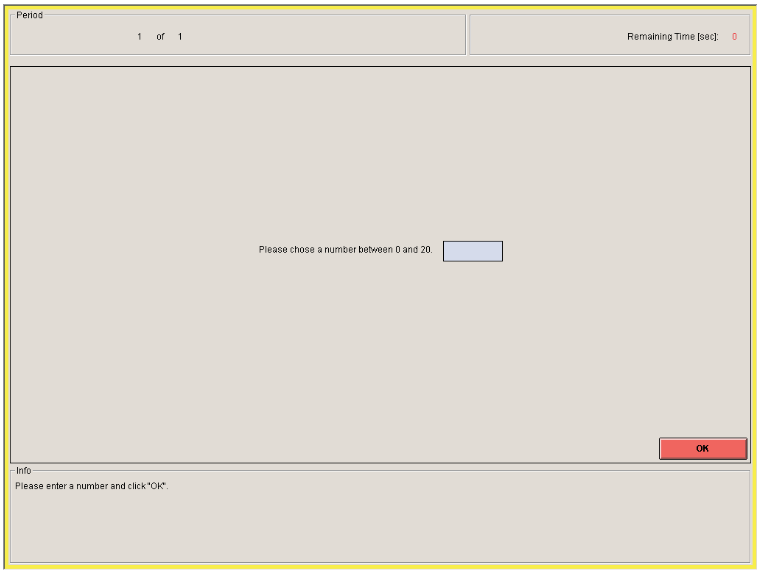

In every period you just have to make a single decision. You simply have to choose a number between 0 and 20. After you have entered your number and have clicked "OK" you will see the following result screen:

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

How the outcome is 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.

2. Instructions to subjects for phase two 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 Points = AUD 1.25

Please note: 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.

Appendix III: Individual correlation between own and other group members' contributions in the learning (phase 1) and standard condition (phase 2): 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