



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS



The University of
Nottingham

Discussion Paper No. 2012-18

Ralph-C Bayer, Elke Renner
and Rupert Sausgruber
November 2012

Confusion and Learning in the
Voluntary Contributions Game

CeDEx Discussion Paper Series

ISSN 1749 - 3293



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/cedex> for more information about the Centre or contact

Sue MacCormick
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 15469
Fax: +44 (0) 115 95 14159
sue.maccormick@nottingham.ac.uk

The full list of CeDEX Discussion Papers is available at

<http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx>

Confusion and Learning in the Voluntary Contributions Game*

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

November 2012

Abstract

We use a limited information environment to assess the role of confusion in the repeated voluntary contributions game. A comparison with play in a standard version of the game suggests, that the common claim that decision errors due to confused subjects biases estimates of cooperation upwards, is not necessarily correct. Furthermore, we find that simple learning cannot generate the kind of contribution dynamics commonly attributed to the existence of conditional cooperators. We conclude that cooperative behavior and its decay observed in public goods games is not a pure artefact of confusion and learning.

Keywords: voluntary contribution mechanism, 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. We thank the editor Jordi Brandts, Simon Gächter, Martin Sefton and three anonymous referees for helpful comments and suggestions.

[†]School of Economics, University of Adelaide, Nexus 10, SA 5005 Adelaide, Australia; E-mail: ralph.bayer@adelaide.edu.au

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

[§]University of Innsbruck, Department of Economics & Statistics, Universitätsstrasse 15, A-6020 Innsbruck; and Department of Economics, University of Copenhagen, DK-1455 Copenhagen K, Denmark; E-mail: rupert.sausgruber@uibk.ac.at

1 Introduction

In experiments based on a simple Voluntary Contributions Mechanism (*VCM*) we observe that subjects contribute more than the conventional model predicts, while contribution levels decrease with repetition. Ever since this kind of behavior has been observed, there has been a debate in the literature whether this regularity is the consequence of intentional behavior, or if it is caused by confused subjects who learn (see Ledyard 1995).

In this paper, we propose a novel approach to assess the role of confusion and learning in explaining the stylized contribution patterns in repeated VCM experiments. Our approach involves studying behavior in a voluntary contributions game with limited information.¹ In our limited-information VCM game, which we refer to as the Learning Condition, subjects choose a number and are only informed about their payoff at the end of each period. Subjects do *not* get initial information about the specifics of the game and its payoff structure. Such a game allows us to observe behavior generated by subjects adjusting their initially arbitrary contribution according to simple adaptive learning. We use behavior of subjects from this Learning Condition as an approximation of the behavior of a confused subject in a standard VCM. Our approach adds to the debate on confusion because, rather than asking whether or not subjects are confused, we place the focus on which effects confusion can possibly have.

We observe, perhaps surprisingly, that subjects who lack essential information about the incentives of the game generate a contribution pattern that is qualitatively very similar to the stylized pattern in the standard game where subjects have full information. At first sight, these similarities seem to suggest that it might be possible that essentially everyone in the standard VCM game could be confused. However, closer inspection of the data also reveals significant differences between the two conditions. These differences indicate that a considerable part of the contributions in a standard game are motivated by social considerations, and confusion cannot solely explain the observed stylized contribution patterns.

The details of our results are as follows. We find that the distribution of initial choices in the Learning Condition is almost symmetric around the median at contributing half of the endowment. This outcome establishes that limiting the information for subjects generates merely random unsystematic noise. However, if the sort of noise we see in this treatment exists also in the Standard Condition and biases contributions upwards, we would predict lower contributions of subjects who are informed. This is not what

¹The study of behavior in limited information environments is common in experiments on learning; see, Mookherjee and Sopher (1994), Van Huyck et al. (2007), and Chen and Khoroshilov (2003).

we observe. In period one, subjects contribute 54% of their endowment in the Learning Condition as compared to 56% the Standard Condition. The difference is not statistically significant. We can therefore infer that, if some unobserved part of the informed subjects were confused and behave as if they were non-informed, eliminating the decisions of the concerned subjects would not readily lead to less cooperation. This finding suggests that the presence of confusion does not necessarily bias contributions upwards, as often claimed.

A second insight from our study is that the dynamics produced by learning with limited information do not appear to explain the decrease of cooperation over time in the standard game. Contributions drop off in both conditions. However, after a few periods of fast learning in the Learning Condition, the rate of decline becomes more than twice as steep in the standard VCM. While a simple stimulus response learning model is capable of explaining the dynamics of contributions in the Learning Condition, the same model fails to track the dynamics of contributions in the standard game. We also find that, in contrast to the standard game where subjects exhibit strong patterns of conditional cooperation, there is no correlation between own and group members' past contributions in the Learning Condition. As we will argue in detail below, these findings suggest that the decay in VCM experiments, which is typically attributed to conditional cooperation, is not merely an artefact of learning of confused subjects.

2 Related Literature

The literature usually describes confusion as subjects' inability to understand the incentives of the game or their incapability to deduce the dominant strategy. There are various potential sources of confusion and subjects may be confused to different degrees. The basic argument we typically find in the literature is that, unlike naturally selected and experienced individuals in real-world situations, subjects who come to the laboratory, are unfamiliar and inexperienced with the choices they need to make. Other standard concerns are that subjects may not understand the incentives because of limited cognitive abilities and/or due to the lack of salient rewards provided by the experimental environment. These forms of confusion may exist simultaneously and be interdependent. For example, when incentives are weak subjects may strive less to make better decisions, which may go as far as exerting little or no effort to understand the instructions.²

²The choices of subjects who are confused may also be more likely to be influenced by objectively irrelevant contextual cues in language and other merely procedural details of the experiment. Ferraro and

The literature has systematically responded to these concerns. Already early contributions by Marwell and Ames (1980), Isaac, Walker and Thomas (1984), and Isaac and Walker (1988) studied whether free-riding increases with experience.³ The results were mixed. Subsequently, a series of experiments by Andreoni (1988), Croson (1996), and Cookson (2000) documented large “restart effects” for experienced subjects, which is at odds with a simple learning hypothesis. In response to the the “lack-of-incentives” argument, Marwell and Ames (1980) and later Kocher, Martinsson and Visser (2008) report conditions where they imposed a fivefold increase in the incentives; these studies found no effect. Goetze and Orbell (1988) asked subjects survey questions about their understanding of the incentives. They found no relationship between subjects’ responses and the level of contributions to the public good. A similar result was later reported by Fosgaard, et al. (2011).

A second branch of the literature takes a more sophisticated angle on the confusion debate, by trying to disentangle contributions that are due to confusion from those that are motivated by social norms or social preferences. To identify confusion amidst these motives, Andreoni (1995) modifies the standard VCM design by paying subjects according to their rank in their group, rather than paying them their income from the VCM. This feature removes the social gains from contributing and thus eliminates altruism as a contribution motive. Houser and Kurzban (2002) and Ferraro and Vossler (2010) eliminate other regarding motives for contributions by replacing other group members with pre-programmed computers. In both studies subjects were aware that they were playing against automata and therefore off-equilibrium play should not be associated with social motives towards other players. The results provide evidence that many subjects fail to maximize their own payoff and that the optimization error is sizeable. The average contributions in the treatments where social motives are removed sometimes also markedly decline with time, which suggests that contribution dynamics in the standard linear VCM may indeed be at least partly a result of learning. These results potentially bear strong implications: if we cannot assume that rationality is common knowledge, then observed cooperation patterns in the VCM might be a result of strategic play and reputational concerns (Kreps et al., 1982).

Vossler (2010), for example, report that “many subjects believe they are playing a sort of stock market game” (p. 24) and they conjecture that this might be caused by the “investment language” used in many voluntary contributions experiments. Fosgaard, Hansen and Wengström (2011) find that more subjects are able to identify the dominant strategy when the game is presented in a “take” as compared to a “give frame.”

³In contrast to repetition, which represents a sequence of decision rounds within the same group of subjects, the term ‘experience’ means that subjects play the game again with a different group.

Today, the bulk of evidence suggests that reputation effects are not sufficient to explain the observed contribution pattern (see, for example, Andreoni 1988, Palfrey and Prisbey 1996, and Keser and van Winden 2000). It also seems widely accepted that observed patterns are most likely driven by some combination of confusion and social preferences. Palfrey and Prisbey (1996, 1997) use a statistical model which differentiates between errors and individual heterogeneity in preferences to explain positive VCM contributions. They conclude that most of the cooperation in the VCM is because of errors and ‘warm-glow’ utility derived from contributing per se, and that altruism is of minor importance in explaining the data.⁴ On the other hand, building on a similar specification, Anderson, Goeree, and Holt (1998) and Goeree, Holt, and Laury (2002) find that some kind of other-regarding motivation is needed in addition to errors to explain empirical facts like that a greater number of players and a higher marginal return from the public good increases the contributions.

Despite such advances in the literature, the problem remains that if subjects can be both confused and motivated by some form of social preferences, then we cannot be sure which assumption (i.e. selfishness or rationality) is at fault when we observe positive contributions. Consequently, possible explanations are confounded. Ultimately, to properly assess the role of confusion, the researcher would need to know the outcome of the VCM in the absence of confusion. As confusion in subjects cannot be simply turned off in a treatment, any attempt to answer this question empirically is necessarily incomplete. A problem for generalizing from the above mentioned results is that the structure of an econometric model or the conditions used to study confusion may miss key elements that are present in the standard VCM. To see why this matters, consider the nowadays widely accepted explanation that contributions in the VCM stem from the heterogeneity of social preferences within a group, where conditional cooperators start out with high contributions. Consecutively they adjust their contributions downward as a negative-reciprocal reaction to selfish group members (see Andreoni 1995, Kurzban and Houser 2005, Muller et al. 2008, Fischbacher and Gächter 2010).⁵ Such motives are not accounted for in the noisy decision models by Palfrey and Prisbey (1997) and Anderson, et al. (1998) as there the behavior is independent of the behavior of others (see Brandts and Schram 2001; for a survey of theoretical explanations for contributions in the linear VCM see Holt and Laury 2008). In the same vein, with all other-regarding behavior removed,

⁴The term ‘warm glow’ was introduced by Andreoni (1993). Based on a similar thought, some authors have proposed games with an interior equilibrium prediction to test economic theory. This literature typically reports a lower level of excess cooperation. See, e.g., Willinger and Ziegelmeyer (2001) and the references they cite.

⁵Gintis et al. (2003) explain similar dynamics with an evolutionary approach.

as for example in Houser and Kurzban (2002), non-confused subjects would choose zero contributions irrespective of social preferences. Hence, drawing inference from such a condition tends to overstate the aggregate effect of confusion, because it only identifies errors that bias contributions upwards. Our approach is agnostic to assumptions about social preferences and does not suffer from this bias.

A similar argument applies with respect to the dynamics of contributions. The process of learning to optimize according to one's (potentially social) preferences is not necessarily equivalent to that of learning to play in a materially self-interested manner.⁶ An advantage of our approach is that it enables us to infer about the effects of confusion and learning independently of the distribution of preference types, as it assumes that confusion is so strong that confused subjects' social preferences are not relevant for behavior.

Implicitly, our approach assumes that confused subjects are naïve and learn according to stimulus-response or reinforcement rules. To the extent that the observed behavior in our limited information setting generalizes to that in a VCM, we can evaluate the claim that all or part of the commonly observed contribution pattern typically attributed to conditional cooperation can also be explained by confusion and simple learning. Our assumption on the nature of confusion and on how people learn are appropriate for a number of reasons. On a general note, research shows that stimulus-response learning rules often provide the best fit of learning models.⁷ Originally developed in biology and psychology (e.g., Bush and Mosteller 1955, Herrnstein 1970), models known as stimulus-response or reinforcement learning have gained considerable importance to explain behavior observed in experimental games (for surveys see Fudenberg and Levine 1998, and Camerer 2003). Moreover, the connection between simple reinforcement learning and confusion finds support from recent surveys. Ferraro and Vossler (2010) asked subjects to state the contribution which would maximize own income in a condition where they play the VCM with automata. They find that almost one third of subjects give an incorrect answer to this question even if they are experienced with playing the game. Fosgaard, et al. (2011) asked a similar question and report that 47% of subjects get this question wrong.⁸ In

⁶Besides potential interaction effects between heterogenous social preferences and decision errors, there are other arguments why subjects may learn differently between strategic and individual choice situations. See, for example, Duersch et al. 2010, who explore how subjects learn to play a Cournot-Duopoly game against computers that are programmed to follow one of various learning algorithms.

⁷For studies on the comparative power of alternative learning models see, among others, Gale, Binmore, and Samuelson (1995), Erev and Roth (1998), Chen and Tang (1998), Feltovich (2000), Tang (2001), Janssen and Ahn (2006). Camerer and Ho (1999) propose a weighted model with choice reinforcement and belief-based (fictitious play) learning as two special cases. Using data from a large class of experimental games, they show that learning is best explained by a combination of both.

⁸These authors use a large sample from the general population in Denmark. They also run a follow-up

both studies, subjects were paid a reward for answering correctly, which implies that this outcome is not likely to be caused by a lack of incentives. Ferraro and Fossler (2010) also provide evidence for naïve learning in the VCM context. In response to a question of how they reached their decision, 55% of subjects chose the option “I invested different amounts and watched how my payoff changed.” When being asked for oral feedback, many subjects frankly admitted that they had no idea what was going on and that they varied contribution levels and tried to infer a pattern to payoffs. Based on such evidence, the assumption of strong naïveté appears correct for modeling behavior of subjects who are confused, and naïve learning might not only be an adequate description of behavior of subjects in our Learning Condition but also of subjects that are informed about the incentives but are confused.

3 Experimental design

The structure of the experimental voluntary contributions game is as follows. The subjects are assigned to groups of four. Every period the subjects receive 20 points as their initial endowment. Every point invested into a public good pays 0.4 points (0.5 Australian Cents) to each subject in the group, while every point kept for private investment pays one point (1.25 Cents) to only the subject who keeps it. Thus, free riding is a dominant strategy but full contributions to the public good are socially efficient.

We implement two conditions: the first condition is the Standard Condition which replicates a standard VCM. The second condition is the Learning Condition where we deliberately withhold information about the incentive structure of the game from the subjects. We implemented two versions of the Learning Condition – one with minimal information and one with limited information (for instructions see the online supplement to this paper).

In the Minimum Information Condition subjects learn nothing about the game and its incentive structure. Subjects are asked to choose a number between 0 and 20 in each period. Participants do not know that this number is a contribution choice nor do they know that they are a member of a group. The instructions tell them that the aim of the experiment is the study of learning behavior and that their payoff is determined by the number they have chosen and “other factors” and that “these factors might change from period to period.” After each period subjects are informed about their payoff. The only

experiment with standard student subjects in which even fewer subjects correctly answer this question.

additional feedback is a reminder of their own choice.⁹

In the Limited Information Condition¹⁰ subjects still do not learn the incentive structure but the instructions contain information about group membership and use some of the terminology that is also used in the instructions of the standard VCM like “project” or “contribution.” In this condition subjects are told that they are randomly assigned into groups of four in a partner matching and that each participant is asked in each period to decide how much of their 20 point endowment they “want to contribute to a project”. They still do not know how payoffs exactly are determined, but they are told that their income depends on their own decisions and “the decisions of the other group members.” Thus, subjects know that their payoffs are interdependent but they do not know in what way. The feedback in each period is the same as in the Minimum Information Condition. Subjects learn their own payoff and are reminded about their own choice.

Both Learning Conditions have in common that, due to the lack of information about the incentives of the game, subjects are in a state of “ignorant” confusion. In both conditions subjects can only learn through their own past choices and payoffs. Other contextual clues that are also present in a standard VCM are kept to a minimum in the Minimum Information Condition, but are present in the Limited Information Condition. We run both variations of the Learning Condition to test if our results are robust with respect to these contextual cues.

Table 1: Experimental design

	<i>Phase 1 (Periods 1-20)</i>	<i>Phase 2 (Periods 21-40)</i>
Standard Condition (36 subjects)	Standard VCM with full instructions	n/a
Learning Condition		
– Minimum Information (60 subjects)	Instructions in “learning frame”, no information about incentive structure	Standard VCM with full instructions
– Limited Information (68 subjects)		

In all conditions the game was repeated 20 times in a partner-matching protocol. To allow for a within-subject comparison, the experiments in both Learning Conditions consisted of two phases: in the first phase subjects made decisions in the 20-times repeated


⁹Ferraro and Vossler (2010) suggest that subjects may use the actions of others as an indication of profit-maximizing behavior. Note that such “herding” is ruled out by our design.

¹⁰We are grateful to an anonymous referee who suggested this treatment as a robustness check.

game under limited or minimum information. Then they were informed that there will be a second part of the experiment and they received instructions and made decisions in the Standard Condition, which consisted of a standard VCM with full instructions. We also ran some sessions with only the Standard Condition, which allows us to check for order effects. We ran eight sessions with 16, 20, or 24 subjects each and a total of 164 subjects (36 in the Standard Condition, 60 in the Learning Condition with minimum and 68 with limited information). Table 1 summarizes our design. Subjects were recruited with ORSEE (Greiner, 2003) among first-year students at the University of Adelaide from a variety of fields, who had not participated in an experiment before. The sessions were conducted with the software package z-Tree (Fischbacher 2007) and lasted between 35 (Standard Condition only) and 50 minutes (both phases). The average subject earned the equivalent of US\$ 11.9 (in Australian Dollars) within this time.

4 Results

The left panel in Figure 1 shows the time series of average contributions, as a percentage of the endowment, in both Learning Conditions. The black line shows the average observed contribution behavior in the case of Limited Information, while the grey line depicts the contributions in the Minimum Information Condition. As one would expect for a situation where subjects cannot fully understand the implications of their behavior, the contributions on average start out around the midpoint of the admissible action space. With repeated play, however, contributions in both treatments drop off from 53.4 (53.7) percent of the total endowment in period one to 35.7 (27.6) percent in the Minimum (Limited) Information Condition in the last period.



Overall, the dynamics in the two Learning Conditions do not only look very similar but

can also not be distinguished between statistically. Using period-specific Mann-Whitney tests, we fail to reject the null hypothesis of equal contributions in the two Learning Conditions in all cases. For this reason we pool the data from the Minimum and Limited Information Conditions and refer to them as the data from the Learning Condition.

In the right panel of Figure 1 we compare the average contributions of the Learning Condition (pooled) with those in the Standard Condition.¹¹ Inspection shows two things: contributions start out slightly higher in the Standard Condition (56.1 vs. 53.6), and the decay in the Standard Condition is stronger. This is confirmed by a simple linear regression, where the group contribution (in percent of the endowment) is regressed on a time trend and a constant. We conduct this regression separately for the Learning and for the Standard Condition and allow for error clustering on the group level. Using a Chow test to compare the time trend coefficients and the constants confirms the results from our visual inspection. The estimated decay of contributions per period is significantly larger in the the Standard Condition (-2.15 percentage points vs. -1.08 percentage points, $p < 0.001$).¹²

Figure 2 shows the distributions of contributions in the Standard and in the Learning Condition in the first period. In both conditions there are mass points at focal contribution levels (0, 25, 50, 75 and 100 percent) and about 65 percent of subjects in each condition choose one of these focal contributions. There is one remarkable and significant difference between the two conditions: in the Learning Condition only 12.5% of the contributions involve full contributions, while in the Standard condition in 25% of the decisions subjects contribute their entire endowment (significant at $p = .0075$, two tailed, according to a Fisher's exact test). Apart from this difference, the distributions of contributions look remarkably similar in both conditions.¹³

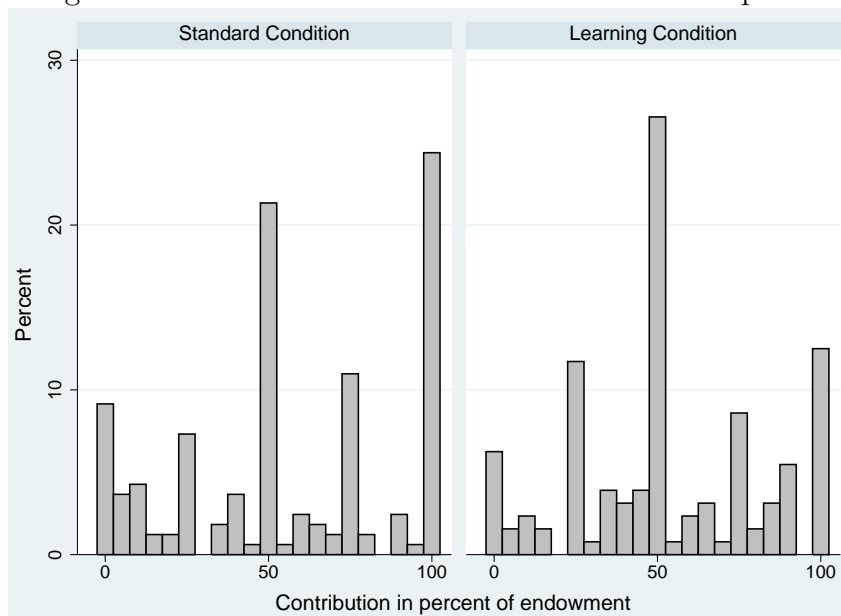
The observations above lead to the following result. The common claim that the

¹¹Contributions in the Standard Condition are very similar to contributions in the experiments where the standard VCM was played in a second phase of the Learning Condition (the averages are 7.24 in the Standard Condition and 8.25 and 7.73 after the minimum and limited information condition, respectively). A Mann-Whitney U-test (two tailed, applied on group averages) is insignificant at $p > 0.28$ ($p > 0.80$) if we compare the contributions in the Standard Condition with the results in the standard VCM played in a second phase after the Minimum (Limited) Information Condition. We therefore pool the data of all experiments that involve the standard VCM. Note that this result is consistent with our finding, that a reduction of confusion (which may happen in this case because of learning with limited information), does not necessarily lead to lower contributions in the Standard Condition.

¹²This findings is in line with Ferraro and Vossler (2010), who also observe that contributions decrease faster in the standard VCM game as compared to their learning condition involving virtual players.

¹³A Kolmogorov-Smirnov test to test the null hypothesis of identical distributions is insignificant ($p > 0.19$).

Figure 2: Distributions of contributions in the first period



presence of confusion implies upwards biased estimates of cooperation is not necessarily correct.¹⁴ If confusion takes the form of ignorance, as mimicked in our limited information environment, then a population with only confused subjects produces the same over-all level of contributions as a mixed population of confused and non-confused subjects (i.e. our Standard Condition).¹⁵ To see why the existence of confusion does not necessarily lead to more contributions, imagine two subjects with identical social preferences. Suppose one subject is confused, while the other subject is not. Then it is plausible that the confused subject contributes less than the non-confused subject. Social preferences might induce a non-confused subject to make positive contributions, while they might not have this impact if a subject is confused.

4.1 Learning versus conditional cooperation

The observation that chosen numbers decrease with repetition in the Learning Condition just as contributions do in the standard VCM can potentially provide support for the claim that learning can be mistaken for conditional cooperation. To gain some more confidence that our Learning Condition really accurately picks up learning dynamics –

¹⁴The typical argument claims that, since a rational selfish player chooses the lowest possible contribution (i.e. zero), any confusion would lead to a positive and therefore higher contribution.

¹⁵Comparing the over-all average group contributions across the conditions reveals that there are no significant differences ($p > 0.35$, Mann-Whitney U-test). The average group contribution per period is 39.1 percent in the Standard Condition and 40.0 percent in the Learning Condition.

and nothing else – we compare the actual behavior to the simulated outcomes of a simple learning model. We opted for an extremely simple learning model without any parameters that require estimation. For our purpose this is sufficient, since it is not our aim to find the learning model that fits our data best among a large class of models. We rather want to demonstrate that our data are consistent with a learning model that combines some features from a variety of different classes of models that have been shown to perform well in different games.¹⁶

In what follows we briefly describe our model. The model’s simplicity allows for a verbal description. A precise mathematical treatment is relegated to Appendix A.1. Some key features of the model are dictated by the sparse information environment subjects face. “Confusion as ignorance” as mimicked by this treatment means that individuals do not know the underlying structure of the VCM 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. Since a subject only knows her own strategy choice and the corresponding payoff, a model in either the hill-climbing or reinforcement tradition is called for. As both hill-climbing (e.g. Huck et al., 1999) and reinforcement models that account for similarity of strategies (e.g. Sarin and Fahid, 2004; Chen and Khoroshilov, 2003) have been shown to fit well in low-information environments, we use a model that has elements of both. Moreover, our model is related to directional learning (Selten and Stöcker, 1986), as subjects move within the strategy space into the direction where higher payoffs are utilized.

In our model a subject decides on which strategy to play by comparing the payoff-choice pairs of the last two periods. A subject only considers strategies that are closer to the strategy (out of the last two chosen) that yielded the higher payoff. Among these strategies a subject chooses randomly with equal probability.¹⁷ Closeness to the better performing strategy is defined as the Euclidean distance. For example, a subject who has chosen contributions of 5 and 12 in the last two periods, where the choice of 5 yielded the better payoff, will randomize over contributions that are less distant to 5 than to 12 (i.e. $\{0, 1, \dots, 8\}$) in the next period.


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 in order to

¹⁶See Camerer (2003, chapter 6) and the references therein for a nice overview. Our approach is similar to the reinforcement learning model by Ferraro and Vossler (2010) and Mason and Philips (1997).

¹⁷Initially, we planed to refine the distribution allowing for more mass on the past choice. 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 choices with equal probabilities. So we decided to stick with this simple formulation.

apply the learning model. We follow the widespread approach in the literature and use the observed choice distribution from the first two periods as a starting point. The first two choices are assumed to be driven by some factors exogenous to our learning model, such as focal points.

Figure 3: Time series of model simulation (grey) vs. average contributions (black) in the Learning Condition (left panel) and the Standard Condition (right panel)



The left panel of Figure 3 shows the result of simulated behavior from the model (grey line) together with the real choices in the Learning Condition (pooled data from the Limited and Minimal Information Conditions). We simulated 10,000 groups. As the starting values are not determined endogenously in the model, they were drawn from the empirical distributions of the observed contributions in the two first periods.¹⁸ 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.

The right panel of Figure 3 investigates if the data generated in the Standard Condition could be explained by the same learning model, which was designed to capture learning behavior of ignorant subjects. As before, the grey line 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 VCM appear to be different than the simulated 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

¹⁸We also simulated the learning model with different initial choices. Even when starting with extreme values (only 0 or 20) simulated behavior quickly converges to that generated by starting values drawn from the empirical distributions.

20 periods are significantly smaller in the Learning Condition ($p < 0.01$ pooled, $N = 32$, and separately by Learning Condition, $N = 15$ and $N = 17$).¹⁹

The analysis above shows that a learning model explains the dynamics in the Learning Condition, while it fails to explain the dynamics in the standard VCM. 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 VCM. Even after controlling for learning dynamics some decay in contributions still remains. We can conclude that the typical decay observed in voluntary contributions games cannot be an artefact of learning of ignorant subjects alone.

4.2 Within-subject comparison

While the analysis so far has been on an aggregate (or average) level, we now look at individual behavior. In what follows, we identify subjects that behave like conditional cooperators in the sense that their contribution is positively correlated with the past average contribution of their group members. We then compare the number of subjects that exhibit this positive correlation in the Standard and Learning Condition. Similar numbers of subjects with positive correlations in the two conditions would cast doubt on the common interpretation that a positive correlation is due to conditional cooperation. Hence, similar frequencies of subjects with a positive correlation would also indicate that concern withdrawal by disappointed conditional cooperators is not necessarily the driving force behind the decay of contributions in standard VCM games. This is the case as then simple learning dynamics could produce the same correlation structure and similar decay as observed in VCM games.

Table 2 reports absolute (and relative) frequencies of subjects exhibiting positive, negative, or no correlation of contributions with past average group contribution (not including their own) in the different conditions (Spearman rank-correlation coefficient: positive and significant (+); negative and significant (-); insignificant (0); $\alpha = 0.05$). In the Learning Condition only seven out of 128 subjects (5.5 percent) exhibit a significantly positive correlation, whereas in the Standard Condition 62 of 128 (48 percent) do.²⁰ 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 4.7 percent of the subjects (six out of 128) show positive correlation in the Standard Condition and

¹⁹The average mean square error of the simulation is almost five times larger in the Standard Condition (3.21 vs. 0.66 points).

²⁰Fischbacher, Gächter and Fehr (2001) report 50 percent conditional cooperators.

Table 2: Correlation between contributions and past average group contribution

		$corr(c_t^i, \bar{c}_{t-1}^{-i})$ Learning Conditions			
		-	0	+	Total
$corr(c_t^i, \bar{c}_{t-1}^{-i})$ Standard Condition	-	0 (0.00)	2 (1.56)	0 (0.00)	2 (1.56)
	0	2 (1.56)	61 (47.66)	1 (0.78)	64 (50.00)
	+	1 (0.78)	55 (42.97)	6 (4.69)	62 (48.44)
Total		3 (2.34)	118 (92.19)	7 (5.47)	128 (100.00)

also 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. Less than 10 percent (six out of 62) of the subjects with a positive correlation in the Standard Condition also exhibit a positive correlation in the Learning Condition. This is more evidence that learning behavior (in the sense of reduced ignorance) does not fully explain the decay in VCM games, which leaves room for conditional cooperation as an important factor. We also estimated the dynamics using a dynamic panel regression for the Learning Condition and for conditional cooperators and separately for others in the Standard Conditions. Such a regression can deal with potential serial correlation in subjects' contributions. The results confirm the findings of the simple analysis presented above. Only the conditional cooperators in the Standard Condition vary their contributions with past contributions of their group members.²¹ In the Learning Condition past contributions of group members have no significant impact on contributions.


4.3 Discussion

In this subsection we discuss a possible concern according to which naïve subjects may face different possibilities to learn between the Learning and the Standard Condition. If this were the case, then a comparison of the dynamics between the two conditions may have little meaning. To deal with this concern, we searched for patterns in our data which

²¹For the full regression results see Appendix A.2.

may be informative of the learning of those subjects who are confused. In Section 4.2 on the within-subject comparison, we report that 5.5% of subjects in the Learning Condition exhibit a significantly positive correlation with past group contributions, whereas in the Standard Condition 48% do. Existing research proposes conditional cooperation based on social preferences as likely motive to explain the correlation pattern in the Standard Condition. Moreover, the details of our study suggest that the 48% are more likely to be aware of the incentives than the rest of the subjects. To see this, notice that the subjects in our experiment are informed only about their own payoff at the end of each period. This design feature differs from most other studies, where subjects typically receive feedback also about the average or individual contributions of other subjects and where it is therefore difficult to discriminate between conditional cooperation, herding or imitation as a decision motive. Without feedback on others' contributions, a conditional cooperator needs to be able to infer this information from their own payoff. Hence, if a subject is classified as conditional cooperator we can be reasonably confident that this subject has understood the incentives. Clearly, these subjects were not confused and also could not have learned by just imitating.

Figure 4: Estimated contribution dynamics of “others” in the Standard Condition and of all subjects in the Learning Condition



The remaining subjects in the Standard Condition, which are *not classified as conditional cooperators*, exhibit contribution dynamics indistinguishable from those of uninformed subjects in the Learning Condition. Figure 4 illustrates this fact by plotting the average contribution dynamics for “others” in the Standard Condition from our panel regression against the observed average contributions in the Learning Condition. Un-

der the assumption that a sizable portion of the subjects that were not categorized as conditional-cooperators were confused (just like the 55% of subjects in Ferraro and Fossler (2010) who stated “I invested different amounts and watched how my payoff changed.”), the absence of differences in the contribution dynamics provides strong evidence against the hypothesis that the learning opportunities differed across our experimental conditions.

5 Conclusion

In this paper we report on a novel experiment designed to identify the influence of confusion on the dynamics in repeated VCM games. In contrast to previous studies, we study confusion in a benchmark condition 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 VCM.

From our experiments we draw two main conclusions. First, the existence of confusion in the VCM does not necessarily lead to an upward biased estimate of cooperation levels, as often conjectured. Secondly, the decay in the VCM that is typically attributed to conditional cooperation is not an artefact of learning of confused subjects. The decay from learning is less strong, and learning does not produce the positive correlation between contributions and past group member contributions observed in the standard VCM.

On a more general note, we believe that the limited information approach can also be used in contexts other than voluntary contributions games in order to test if results obtained from full information treatments are likely to be caused by confused subjects. Similar to the concept of “Zero Intelligence” (e.g., Gode and Sunder 1993), the approach generates insights into the dynamics of outcomes without having to make strong behavioral assumptions.²² The strength of this method is that one can study the effects of the mechanism in isolation from behavior governed by heterogeneous social preference and the interaction of strategies.

It remains to be mentioned that our design deals with a very specific kind of confusion. There are many other elements to confusion that are not captured by the limited-information approach. In particular, cases not captured by our approach are those, where subjects misinterpret all or parts of the instructions such that the misinterpretation leads to a specific bias in behavior.

²²An important difference of our approach to agent-based modeling is that we study the behavior of human subjects rather than the outcomes generated by computerized agents.

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] **Andreoni, James** 1993. “An Experimental Test of the Public Goods Crowding-Out Hypothesis.” *American Economic Review* 83(5): 1317-1327.
- [4] **Anderson, Simon P., Jacob K. Goeree, and Charles A. Hol**, 1998. “A Theoretical Analysis of Altruism and Decision Error in Public Goods Games.” *Journal of Public Economics* 70: 297–323.
- [5] **Blundell, Richard and Steven Bond**, 1998. “Initial Conditions and Moment Restrictions in Dynamic Panel-Data Models.” *Journal of Econometrics* 87: 115-143.
- [6] **Brandts, Jordi and Arthur Schram**, 2001: “Cooperation and Noise in Public Goods Experiments: Applying the Contribution Function Approach.” *Journal of Public Economics* 79: 399-427.
- [7] **Bush, Robert R. and Frederik Mosteller**, 1955. *Stochastic Models for Learning*. John Wiley & Sons, Inc.
- [8] **Camerer, Colin F.**, 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- [9] **Camerer, Colin F. and Teck-Hua Ho**, 1999. “Experience-Weighted Attraction Learning in Normal Form Games.” *Econometrica* 67: 827-874.
- [10] **Chen Y. and Tang, F.**, 1998. “Learning and Incentive Compatible Mechanisms for Public Goods Provision: An Experimental Study.”, *Journal of Political Economy* 106: 633-662.
- [11] **Chen, Yan and Yuri Khoroshilov**, 2003. “Learning under Limited Information.” *Games and Economic Behavior* 44: 1-25.
- [12] **Cookson, Richard**, 2000. “Framing Effects in Public Goods Experiments.” *Experimental Economics* 3: 55-79.
- [13] **Croson, Rachel**, 1996. “Partners and Strangers Revisited.” *Economics Letters* 53: 25-32.

- [14] **Duersch, Peter; Albert Kolb; Jörg Oechssler and Burkhard C. Schipper**, 2010. “Rage Against the Machines: How Subjects Play against Learning Algorithms.” *Economic Theory* 43(3): 407-430.
- [15] **Erev, I., and Roth, A. E.**,1998. “Predicting How People Play Games: Reinforcement Learning in Experimental Games With Unique, Mixed Strategy Equilibria.” *American Economic Review* 88: 848-881.
- [16] **Feltovich, Nick**, 2000. “Reinforcement-Based vs. Belief-Based Learning Models in Experimental Asymmetric-Information Games. *Econometrica* 68: 605-641.
- [17] **Ferraro, Paul J. and Christian A. Vossler**, 2010. “The Source and Significance of Confusion in Public Goods Experiments.” *BE Journal of Economic Analysis & Policy* 10(1). Article 53.
- [18] **Fischbacher, Urs**, 2007. “Z-Tree: Zurich Toolbox for Readymade Economic Experiments.” *Experimental Economics* 10(2): 171-178.
- [19] **Fischbacher, Urs and Simon Gächter**, 2010. “Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Good Experiments.” *American Economic Review* 100(1): 541-556.
- [20] **Fosgaard, Toke; Lars Gårn Hansen and Erik Wengström**, 2011. “Framing and Misperceptions in a Public Good Experiment” FOI Working Paper 2011/11, University of Copenhagen, Institute of Food and Resource Economics.
- [21] **Fudenberg, D. and Levine, D.**, 1998. “Learning in Games.” *European Economic Review* 42: 631-639.
- [22] **Gale, J., Binmore K. and Samuelson L.**, 1995. “Learning to Be Imperfect: The Ultimatum Game.” *Games and Economic Behavior* 8(1): 56 -90.
- [23] **Gintis, Herbert; Samuel Bowles; Richard Boyd and Ernst Fehr**, 2003. “Explaining Altruistic Behavior in Humans.” *Evolution and Human Behavior* 24(3): 153-172.
- [24] **Gode, Dhananjay K. and Shyam Sunder**, 1993: Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy* 101: 119–137.

- [25] **Goeree Jacob K., Charles A. and Susan K. Laury**, 2002. Private Costs and Public Benefits: Unraveling the Effects of Altruism and Noisy Behavior. *Journal of Public Economics* 83: 255-276.
- [26] **Goetze, David and John M. Orbell**, 1988. "Understanding and Cooperation in Social Dilemmas." *Public Choice* 57:275-279.
- [27] **Greiner, Ben**, 2003. "An Online Recruitment System for Economic Experiments." in: Kremer, K. and Macho, V. (Eds.): *"Forschung und wissenschaftliches Rechnen"*.
- [28] **Herrnstein, Richard J.**, 1970: "On the Law of Effect." *Journal of the Experimental Analysis of Behavior* 13(2): 243-266.
- [29] **Holt, Charles A., and Susan K. Laury**, 2008: "Theoretical Explanations of Treatment Effects in Voluntary Contributions Experiments." In: Charles R. Plott and Vernon S. Smith (Eds.): *Handbook of Experimental Economic Results*, Elsevier: New York. Vol. 1, Part 6: 846-855.
- [30] **Houser, Daniel and Robert O. Kurzban**, 2002. "Revisiting Kindness and Confusion in Public Goods Experiments." *American Economic Review* 92(4): 1062-1069.
- [31] **Huck, Steffen; Normann, Hans-Theo and Jörg Öchssler**, 1999. "Learning in Cournot Oligopoly-An Experiment." *Economic Journal* 109: 80-96
- [32] **Isaac, R. Mark, James Walker, and Susan Thomas**, (1984). Divergent Evidence on Free Riding: An Experimental Examination of Possible Explanations. *Public Choice* 43(1), 113-149.
- [33] **Isaac, R. Mark and James Walker**, (1988). Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism. *Quarterly Journal of Economics* 103, 179-199.
- [34] **Janssen, Marco A. and Toh-Kyeong Ahn**, 2006. "Learning, Signaling, and Social Preferences in Public-Good Games." *Ecology and Society* 11(2): Art. 21.
- [35] **Keser, Claudia and Frans van Winden**, 2000. "Conditional Cooperation and Voluntary Contributions to Public Goods'." *Scandinavian Journal of Economics* 102: 23-39.
- [36] **Kocher, Martin G., Peter Martinsson, and Martine Visser**, 2008. Does Stake Size Matter for Cooperation and Punishment? *Economic Letters* 99: 508-511.

- [37] **Kreps, David M., Paul Milgrom, John Roberts and Robert B. Wilson**, 1982. “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma.” *Journal of Economic Theory* 27(2): 245-252.
- [38] **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.
- [39] **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.
- [40] **Marwell, Gerald and Ruth E. Ames**, (1980). “Experiments on the Provision of Public Goods II: Provision Points, Stakes, Experience, and the Free-Rider Problem. *American Journal of Sociology* 85(4): 926–937.
- [41] **Mookherjee, Dilip and Barry Sopher**, 1994. “Learning Behavior in an Experimental Matching Pennies Game.” *Games and Economic Behavior* 7: 62-91.
- [42] **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.
- [43] **Palfrey, Thomas R. and Jeffrey E. Prisbrey**, 1996. “Altruism, Reputation and Noise in Linear Public Goods Experiments. *Journal of Public Economics* 61: 409–427.
- [44] **Palfrey, Thomas R. and Jeffrey E. Prisbrey**, 1997. “Anomalous Behavior in Public Goods Experiments: How Much and Why? *American Economic Review* 87: 829–846.
- [45] **Sarin, Rajiv and Farshid Vahid**, 2004. “Strategy Similarity and Coordination.” *Economic Journal* 114: 506-527.
- [46] **Selten, Reinhard and Rolf Stöcker**, 1986. “End Behavior in Sequences of Finite Prisoner’s Dilemma Games: A Learning Theory Approach.” *Journal of Economic Behavior and Organization* 7: 47-70.
- [47] **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.

- [48] **Zieglmeyer, A. and Willinger, M.**, 2001. "Strength of the Social Dilemma in a Public Goods Experiment: an Exploration of the Error Hypothesis," *Experimental Economics* 4(2): 131-144.

A Appendix

A.1 Details of the learning model

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.²³ 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

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

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

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:

²³This specific restriction could be interpreted as an extreme similarity function in a similarity augmented learning model (Sarin and Vahid, 2004) or as an element taken from directional learning (Selten and Stöcker, 1986).

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

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 .²⁴ 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:

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

A.2 Dynamic-panel estimation

We estimated a dynamic panel, which allows for contributions to depend on past own contributions and on past contributions of other group members. By design any unobserved panel-level effects are correlated with the lagged own contributions. For this reason we used the Arellano-Bond-Bover GMM estimator with additional moment conditions developed in Blundell and Bond (1998) that can handle this endogeneity problem. Table A.1 reports the results.

²⁴This assumption differs slightly from a traditional reinforcement-learning model in that it allows for “strategy similarity” (Sarin and Vahid, 2004). In our formulation admissible strategies are not only seen as similar but even identical by the subjects.

Table A.1: Estimation of contributions in the Standard Condition (by subject type) and in the Learning Condition

Variable	Cond. Coop.		Others		Learning	
	Coef.	(S. E.)	Coef.	(S. E.)	Coef.	(S. E.)
<i>Lagged own contribution</i>						
c_{t-1}^i	0.185	(0.198)	0.181	(0.096)	-0.032	(0.107)
c_{t-2}^i	-0.127	(0.202)	-0.021	(0.086)	-0.029	(0.078)
<i>Lagged average contribution of others</i>						
\bar{c}_{t-1}^i	0.496**	(0.129)	0.074	(0.128)	-0.039	(0.067)
\bar{c}_{t-2}^i	-0.052	(0.226)	-0.068	(0.106)	-0.029	(0.078)
<i>Period dummies, t=3 omitted</i>						
$t = 4$	1.290	(1.291)	-1.819*	(0.776)	-1.156	(0.844)
$t = 5$	1.292	(1.330)	-1.979**	(0.688)	-4.200**	(1.130)
$t = 6$	-0.339	(1.122)	-2.592**	(0.859)	-2.718	(1.622)
$t = 7$	-1.413	(1.567)	-3.062**	(0.684)	-2.161	(1.348)
$t = 8$	-1.041	(1.950)	-2.783*	(1.092)	-1.946	(1.270)
$t = 9$	-1.116	(2.470)	-4.478**	(1.343)	-3.026*	(1.209)
$t = 10$	-2.745	(2.863)	-3.372*	(1.417)	-2.783*	(1.378)
$t = 11$	-1.846	(3.336)	-3.795*	(1.615)	-3.067*	(1.226)
$t = 12$	-2.824	(3.399)	-3.038*	(1.530)	-3.152**	(1.176)
$t = 13$	-2.979	(3.287)	-4.416**	(1.451)	-5.145**	(1.381)
$t = 14$	-3.583	(3.788)	-4.342**	(1.600)	-3.690*	(1.734)
$t = 15$	-3.139	(4.244)	-3.933**	(1.404)	-4.082*	(1.757)
$t = 16$	-5.221	(4.322)	-3.844*	(1.874)	-4.048*	(1.686)
$t = 17$	-3.396	(4.814)	-4.982*	(2.008)	-3.737*	(1.691)
$t = 18$	-4.726	(4.821)	-4.857*	(2.004)	-3.188	(1.939)
$t = 19$	-5.134	(4.669)	-4.703*	(2.140)	-4.816**	(1.752)
$t = 20$	-6.211	(5.278)	-5.403*	(2.121)	-3.884*	(1.639)
<i>Intercept</i>	6.307	(7.040)	9.212**	(3.347)	12.328**	(3.811)
N	738		990		1080	
$prob > \chi_{(21)}^2$	0.000		0.000		0.000	