The Stability of Conditional
Cooperation: Egoism Trumps
Reciprocity in Social Dilemmas

Luciano Andreozzi
Matteo Ploner
Ali Seyhun Saral

MAX PLANCK
SOCIETY

The Stability of Conditional Cooperation: Egoism Trumps Reciprocity in Social Dilemmas

Luciano Andreozzi / Matteo Ploner / Ali Seyhun Saral

October 2019

The Stability of Conditional Cooperation: Egoism Trumps Reciprocity in Social Dilemmas

Luciano Andreozzi [*]   Matteo Ploner [*]
Ali Seyhun Saral [†‡]

October 2019

Abstract

An often-replicated result in the experimental literature on social dilemmas is that a large share of subjects reveal conditionally cooperative preferences. Cooperation generated by this type of preferences is notoriously unstable, as individuals reduce their contributions to the public good in reaction to other subjects free-riding. This has led to the widely-shared conclusion that cooperation observed in experiments (and its collapse) is mostly driven by imperfect reciprocity. In this study, we explore the possibility that reciprocally cooperative preferences may themselves be unstable. We do so by observing the evolution of subjects' preferences in an anonymously repeated social dilemma. Our results show that a significant fraction of reciprocally cooperative subjects become selfish in the course of the experiment, while the reverse is rarely observed. We are thus driven to the conclusion that egoism is more resistant to exposure to social dilemmas than reciprocity.

Keywords: reciprocity, conditional cooperation, strategy method

JEL Classification: C72; C91

[*]Department of Economics, University of Trento, 38122 Trento, Italy.
[†]Max Planck Institute for Research on Collective Goods, 53113, Bonn, Germany
[‡]To whom correspondence should be addressed. E-mail: `saral@coll.mpg.de`

1 Introduction

The large experimental literature on social dilemmas, such as the Public Goods Game, the Prisoner's Dilemma and the Trust Game, has been remarkably coherent in revealing two stylized facts. First, subjects cooperate more than they would do if they were rational and purely egoistic. Second, when the game is repeated anonymously, cooperation declines over time, although it rarely disappears altogether. The early explanations for this phenomenon can be grouped under two headings: the *learning hypothesis* (Palfrey and Prisbrey, 1997; Cooper and Stockman, 2002; Burton-Chellew et al., 2015) and the *conditional cooperation hypothesis* (Croson, 1996; Fischbacher et al., 2001; Fischbacher and Gächter, 2010; Neugebauer et al., 2009).

According to the learning hypothesis, subjects are mostly self-interested, but they are boundedly rational and it takes time for them to learn the workings of the game they are playing. For example, they may have trouble understanding the logic of dominant strategies or backward induction. According to this view, what is observed in the early stages of any experiment involving a social dilemma is a mixture of noise and confusion (Gale et al., 1995; Palfrey and Prisbrey, 1997).

This explanation has been repeatedly found wanting. In a classic article, J. Andreoni concluded that bounded rationality and learning could explain no more than half of the observed deviations from self-interest. His verdict was that "[..] the decline in cooperation often observed in [...] public-goods experiments may not be due to learning, but instead may be due to frustrated attempts at kindness" (Andreoni, 1995). The obvious alternative to the learning hypothesis was that the cooperation observed in the lab (and its decline) could be the consequence of some form of pro-social behavior mostly driven by reciprocal altruism (Andreoni, 1995; Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002). After all, deviations from pure self-interest were observed in simpler games, like the sequential Prisoner's Dilemma (Clark and Sefton, 2001) and even in decision settings like the Dictator Game, in which they were harder to attribute to confusion (Andreoni and Miller, 2002; Fisman et al., 2007).

However, the most compelling evidence against the learning hypothesis came directly from Public Goods Game experiments: Fischbacher et al. (Fischbacher et al., 2001), and later Fischbacher and Gächter (Fischbacher and Gächter, 2010), used the strategy method to elicit subjects' preferences

for conditional contribution to the public good. Their data revealed that around a half of the subjects could be classified as conditional cooperators: They were willing to contribute to the public good only if other subjects were contributing as well. Unconditional defectors where slightly more than 20 percent. Cooperation sustained by such conditionally cooperative preferences is inherently fragile as most subjects display a "selfish bias": They are willing to give less than what they expect others to give. Thus, the cooperation would unravel over time, even in the absence of free-riders (Fischbacher and Gächter, 2010). These experimental findings proved to be robust (Kurzban and Houser, 2005; Kocher et al., 2008; Neugebauer et al., 2009; Herrmann and Thöni, 2009; Volk et al., 2012; Thöni and Volk, 2018), and the conditional cooperation hypothesis became the dominant explanation in the experimental literature.

Despite its success, the conditional cooperation hypothesis had a weak spot that passed mostly unnoticed. One of its implicit assumptions is that the conditionally cooperative preferences revealed at the beginning of the experiment are the subjects' "true" preferences. However, this begs the original question, because proposers of the learning hypothesis may retort that *all* preferences are subject to change as the game unfolds and experience accumulates. The conditionally cooperative preferences revealed at the beginning of the experiment should not be an exception, as they too may be due to confusion.

Indeed, there is evidence that at least a share of the reciprocity observed in the lab may be due to an imperfect understanding of the game. Experiments show that a fraction of the subjects involved in social dilemmas display similar pro-social behavior in interactions with computers and human beings alike (Houser and Kurzban, 2002), and conditionally cooperative subjects are precisely those who seem unable to distinguish computers from human beings (Burton-Chellew et al., 2016).

It follows that confused subjects may be of two different kinds. Some are of the familiar type, whose conditional contribution follows a non-discernible pattern. Others may display reciprocally cooperative preferences that, however, are determined by a faulty understanding of the game. For example, some subjects may fail to see that they are playing a repeated game which is anonymous; thus, non-cooperative choices will not have negative consequences on future interactions. This kind of mistake is clearly more difficult to detect, but it may be much more pervasive. To see why, consider the frequently-made observation that choices made by inexperienced
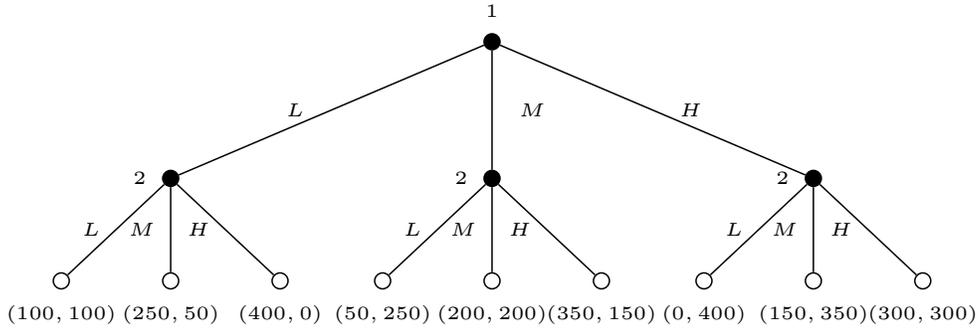
Figure 1: The Three-Actions Sequential Prisoner's Dilemma (3SPD)

subjects reflect their day-by-day experience with situations they perceive as similar to the ones they encounter in the lab (Henrich et al., 2005; Gächter and Herrmann, 2009). It has been argued that when subjects face an unfamiliar game, they tend to respond intuitively with the strategies that are successful in their daily social interactions, while deliberate processes shift their behavior towards the strategies that are most advantageous in the context (Rand et al., 2012). The relevance of this approach for the theory of pro-social behavior stems from the fact that, in everyday life, repeated interactions are rarely anonymous, and in such settings neither unconditional defection nor unconditional cooperation are good strategies. (Axelrod, 2009; Nowak, 2006; Gintis et al., 2005). Hence, it is just natural to expect that, in dealing with unfamiliar situations that are repeated, but anonymous, a non-negligible fraction of subjects initially reveal reciprocally cooperative preferences.

These considerations suggest that the existing evidence on the Public Goods Game is stacked in favor of the conditional cooperation hypothesis. Since in ordinary life it usually pays to be reciprocally cooperative, one should expect this to be the instinctive, immediate choice in the early stages of any experiment involving social dilemmas. However, standard experiments in which conditionally cooperative preferences are only elicited in the early phases of the sessions provide no information about their evolution as experience accumulates and learning takes place.

To address this issue, we run a simple social dilemma experiment in which conditionally cooperative preferences are elicited. However, at variance with

4

the existing studies, our design allows subjects to change their conditionally cooperative choices after every round of the game. This allows us to observe how subjects' stated preferences evolve over time.

We use a three-action version of a sequential Prisoner's Dilemma. Each player receives 100 tokens and is given the opportunity to transfer nothing (low transfer, L), 50 tokens (medium transfer, M), or 100 tokens (high transfer, H) to the other player. The second player chooses after having observed the first player's choice. The Pareto-optimal choice is thus to transfer 100 tokens, although the dominant choice is to transfer nothing. The extensive form of the game is represented in Figure 1. We will refer to this game as the Three-Actions Sequential Prisoner's Dilemma (3SPD). We use the strategy method to elicit subjects' preferences as first and second movers. When playing as second movers, subjects have to state which level of cooperation they would choose in response to any of the three levels of cooperation the first mover may have chosen. A player's final payoff is the sum of the tokens he did not transfer and the tokens he received by the other player, multiplied by three. After each period, roles are randomly assigned, first and second movers are informed about each other's respective action, and the payoffs are obtained. We use the stranger matching procedure: in each period, each subject is matched randomly to another subject to have a distinct one-shot game in each period.

To test the extent to which subjects may respond to the strategies employed by the other subjects, we manipulate the information they receive about the conditional strategies of the others. In the baseline treatment, which we call *NoCondInfo*, subjects are given no information about the other subjects' conditional choices. In the treatment we call *CondInfo*, each first player is informed about the way in which the subject he has been matched with has filled the strategy method questionnaire.

5
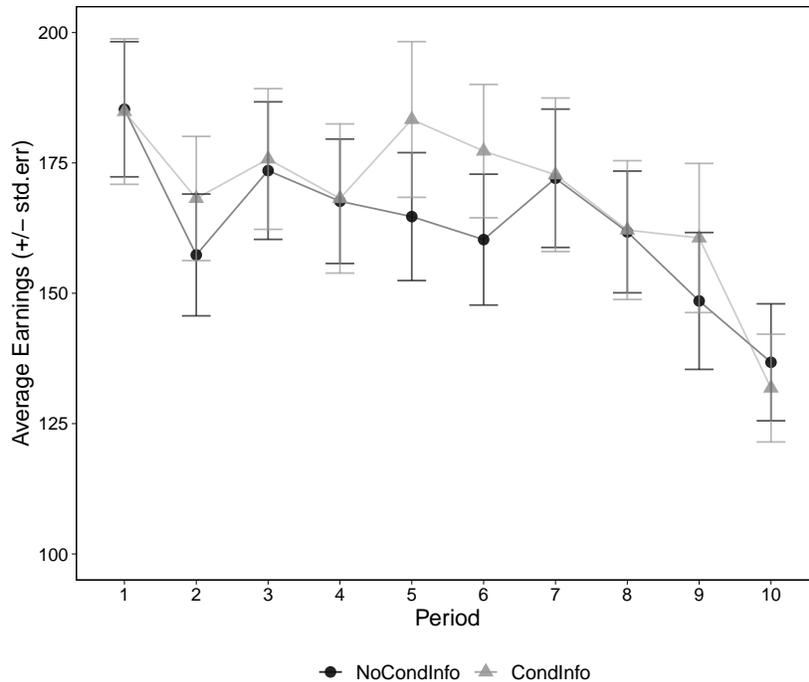
Figure 2: The decline of cooperation over ten periods: payoffs

Results

We focus on the conditional choices made by our subjects. We represent a strategy for the second player as a triple ABC, where A is the action chosen in response to L, B is the action chosen in response to M, and C is the action chosen in response to H. For example, LMH is the perfect conditionally cooperative strategy that always matches the first mover's choice. Given our simple strategic setting, the classification of subjects is immediate. We distinguish *selfish* individuals (LLL) from *perfect conditional cooperators* (LMH) and *imperfect conditional cooperators*, that is, all the subjects whose contribution is a monotonic function of the first-mover contribution other than perfect conditional cooperators . In our sample, there is a small group of hump-shaped subjects, whose contribution is higher in response to intermediate contributions by the other subjects (MHM, LHM, LML). All other subjects are labeled as "other patterns". For more information about

	Cond.Coop.	Selfish	Hump Shaped	Other
This Study (Period I)	67.0	21.6	6.8	4.6
Fischbacher et al. (2001)	50	29.6	13.6	6.8
Kocher et al. I (2008)	80.6	8.3	0	11.1
Kocher et al. II (2008)	44.4	22.2	11.1	22.2
Kocher et al. III (2008)	41.7	36.1	11.1	11.1
Herrmann & Thöni (2009)	52.5	6.3	6.2	35
Fischbacher and Gächter (2010)	50.7	22.9	12.1	14.3

Table 1: Comparison of conditional strategies in different studies

the classification, see *Supplementary Information.*

First-round choices and the evolution of cooperation. Our first results are in line with the previous experiments on repeated social dilemmas: the majority of subjects can be classified as conditional cooperators and cooperation declines over time. Table 1 compares the distribution of alternative strategies in the first round of our experiment with similar results in other well-known studies. Like in most other experiments, in our sample more than half of the subjects (67%) can initially be classified as conditional cooperators, while only a minority is selfish (21.6%). Figure 2 shows the average payoff subjects receive in each period, in the two treatments. As in most other experiments with similar games, payoffs decline over time as subjects switch to less cooperative strategies.

Conditional preferences are unstable. Figure 3 represents the evolution over time of the choice made by each subject as second mover, in both treatments. The red and dark green rectangles represent choices of subjects who are selfish (LLL) and perfect reciprocators (LMH), respectively. The light green rectangles pool together all imperfectly cooperative choices, while the yellow rectangles correspond to hump-shaped choices. All other patterns are represented by gray rectangles. This picture reveals that only a small minority of subjects keep the same strategy throughout the game (26%). Some of them repeatedly switch between several strategies, while others switch only once.

Figure 4 shows the evolution over time of the number of strategies played by second movers, in the two treatments. It reveals that the change in the composition of the population is mostly due to the decline of the number

of subjects who choose the perfect reciprocating strategy LMH and to a corresponding increase in the number of subjects who choose the selfish strategy LLL. The frequency of all other strategies remains fairly constant across the periods.

Perfect conditional cooperators have unstable preferences. To investigate the stability of preferences at individual level, we retain the customary classification based on the choices subjects make in the first round of the experiment. In this setting, to say that a subject is, for example, a perfect conditional cooperator only means that this is the choice he makes in the first round of a social dilemma experiment. We then define a measure of the *stability* of preferences. We say that the strategy chosen in the first round reveals a subject's stable preference if it is also his modal choice in the last five rounds. Otherwise, a subject is labeled as unstable.

Table 2 shows the fractions of stable subjects for different types according to this measure. The most stable subjects are those who turn out to be *selfish* (choosing LLL) in the first round. More than 3/4 of subjects falling into this category display stable preferences. The least stable subjects are those who are initially *perfect conditional cooperators* (LMH). Less than 1/3 of subjects in this category are classified as stable. A series of chi-squared tests show that both *perfect conditional cooperators* and *imperfect conditional cooperators* differ from *selfish* subjects (p-value<0.001 and p-value=0.028, respectively). A significant difference in stability relative to the selfish type is detected for the *hump-shaped* (p-value=0.016) subjects as well. All other comparisons do not detect any significant difference (all p-values > 0.123). See Table 2 for comparison.

	Num. of obs.			Pearson's chi-squared test p-value			
	Stable	Unstable	Percentage	Selfish	Perf. Cond.	Imp. Cond.	Hump-Shp.
Selfish	27	8	77.1%				
Perf. Cond.	16	38	29.6%	< **.001**			
Imp. Cond.	16	17	48.5%	**.028**	.123		
Hump-Shp.	3	7	30%	**.016**	1.000	.504	

Table 2: Stability according to the last five periods for each type and proportion comparisons.
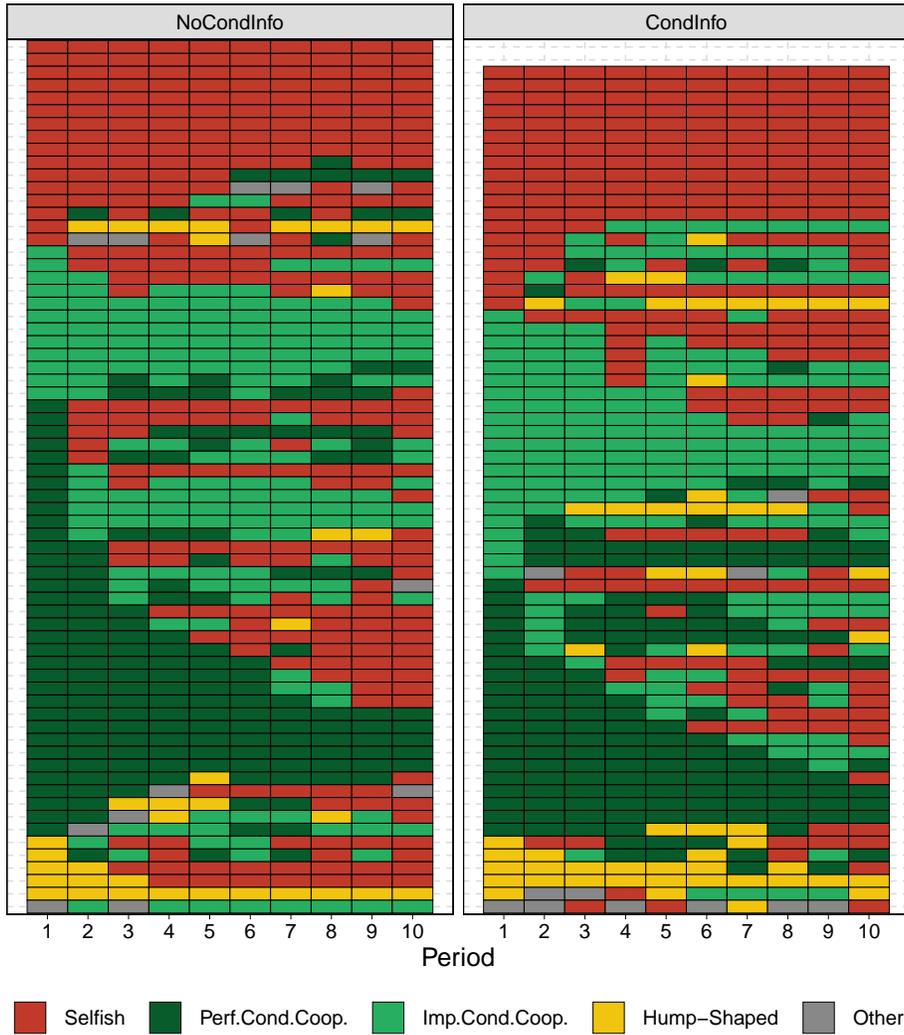
Figure 3: Subjects' conditional strategies over ten periods for each treatment. Every row represents a subject and every colored rectangle captures his conditional strategy in the period denoted on the x axis.
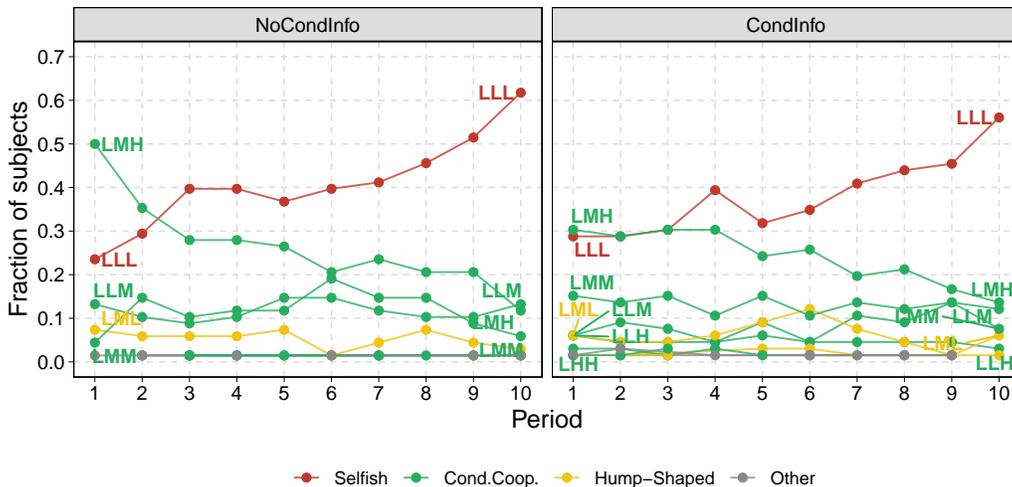
Figure 4: Fraction of subjects choosing each distinct strategy in each treatment.

Information has a limited effect on conditional strategies. A comparison between the treatments *NoCondInfo* and *CondInfo* reveals that it makes no difference whether subjects are shown the conditional choices of the other subjects. Table 3 shows the results of mixed effects logistic regressions where the dependent variable is a dummy capturing conditionally cooperative strategies. The results of the analysis indicate that the likelihood of using such a strategy decreases sharply over time (*Period*). However, the additional information on the counterpart's conditional strategy does not have a significant effect on the choice of playing a conditionally cooperative strategy, as shown by the coefficients of *Treatment CondInfo* and *Period:Treatment CondInfo*.

Discussion

There are two plausible explanations for the data we collected. First, reciprocal cooperators may be inclined to reciprocate not only the *actions* chosen by their partners when playing as first movers, but also the type of conditionally cooperative preferences they reveal when playing as second movers. This is a subtle, but crucial difference. Upon learning that

	Model 1	Model 2
(Intercept)	1.41 (0.44)**	1.56 (0.47)***
Period	−0.24 (0.03)***	−0.27 (0.04)***
Treatment *CondInfo*	0.06 (0.58)	−0.22 (0.66)
Period:Treatment *CondInfo*		0.05 (0.06)
AIC	1223.76	1225.04
BIC	1244.56	1251.04
Log Likelihood	-607.88	-607.52
Num. obs.	1340	1340
Num. groups: subject	134	134
Var: subject (Intercept)	9.81	9.84

***$p < 0.001$, **$p < 0.01$, *$p < 0.05$

Table 3: Mixed-Effects Logistic Regression: Dependent Variable: *isCondCoop*: The binary variable is equal to 1 when a subject is a conditional cooperator and equal to 0 otherwise.

the first mover has defected, a reciprocator playing as second mover will defect. This is the familiar consequence of reciprocity which is captured by the strategy method and is deemed to be responsible for the decay of cooperation. On the other hand, upon learning that many of the other subjects are unconditional defectors (i.e., they would respond with defection to cooperation), a reciprocator may be induced to switch to unconditional defection as well. We shall call these two types of reaction level-one and level-two reciprocity, respectively. The decay in the number of conditional cooperators can obviously be due to level-two reciprocity.

A second explanation is that preference instability is mostly due to a combination of selfishness and learning. When the interaction is anonymous, unconditional defection is the optimal strategy, as a reciprocally cooperative subject fails to exploit the other subjects who play cooperatively as first movers. Selfish subjects who learn from experience will soon discover that when the setting is anonymous there is no point in reciprocating the cooperative behavior of the first mover, and they will eventually become unconditional defectors.

Naturally, our results may be driven by a combination of these two factors.

However, the fact that the patterns in *CondInfo* e *NoCondInfo* treatments are similar suggest to us that learning plays a prominent role in this decline. To see why this is relevant, notice that in the *NoCondInfo* treatment a subject must *guess* what kind of conditional strategy the other subjects are choosing. For example, if he cooperates when in the first-mover position, he may discover that there are subjects who reply with defection to cooperation. If he always defects he will never discover that. By contrast, in the *CondInfo* treatment, there is no need to guess, as the information about the conditional strategy chosen by the opponent is provided by the experimenter. If this piece of information were determinant in the decision to switch to unconditional selfishness, one would expect the decline of reciprocity to be sharper in the *CondInfo* than in the *NoInfo* treatment. The fact that the two treatments are indistinguishable reveals that information about the conditional choice made by the other subjects plays a minor role, if any.

Our data lend support to the thesis that preferences revealed by inexperienced subjects at the beginning of any experiment are rarely stable, as they are prone to change as experience with the game accumulates. This delivers two main messages. First, the literature on social preferences has somewhat downplayed the role of learning. Not only subjects participating to social dilemma experiments need to learn about other subjects' choices. They also need to learn their own true preferences (Plott, 1996; Cooper and Stockman, 2011). This has consequences that go well beyond the study of the decay of cooperation in anonymous social dilemmas. Studies on social norms and social preferences (Krupka and Weber, 2013; Kimbrough and Vostroknutov, 2016) show that a large fraction of the subjects believe reciprocally cooperative behavior to be the normatively appropriate behavior in a Public Goods Game. Our results suggest that these normative preferences may unreliable, as they may be prone to change as subjects gain experience.

The second message is that the subjects' faulty understanding of the game in the early stages of the experiment does not manifest itself as mere noise, and hence is particularly hard to detect. If all subjects who needed to learn how to play the game were initially choosing randomly, the increase of selfish subjects would have been accompanied by a parallel decrease of individuals whose strategy method follows a random pattern. However, this is not what we observe. Rather, we see that most of the change in the composition of the population is due to subjects who are initially labeled as perfect reciprocators and eventually switch to selfishness. In other words, the subjects who are

most likely to change their mind along the experiment are precisely the ones whose initial choices are easier to explain in terms of reciprocity, and hence would not be labeled as "confused".

It is left to future research to study how this learning process takes place, whether there are individual differences, and how the game used in the experiment may affect learning. Most importantly, more research is needed to ascertain whether this type of learning may eventually erode any form of pro-social behavior in experimental settings.

Methods

Participants and Sessions

The experimental sessions were conducted at CEEL, University of Trento. In total, 134 subjects participated in six experimental sessions. A copy of the instructions that were handed out in printed form is available in *Supplementary Information*. All subjects were able to answer the control questions correctly. No subject or session was excluded from the data. Experiments were programmed and conducted in z-Tree (Fischbacher, 2007). The game consisted of ten periods and one period was randomly selected for payment. Participants were given 3 EUR each as a fixed payment and earned between 0 to 15 EUR in additional to that amount, according to the selected period in the session and their payoff in that period.

The Game

The *Three-Actions Prisoners Dilemma Game* (3SPD) we used is a two-player variant of the Public Goods Game which is used by Fischbacher & Gächter (Fischbacher and Gächter, 2010). To see this, consider a four-player Public Goods Game in which conditionally cooperative preferences are elicited. This requires each subject to decide, before the game is played, how much to contribute to the public good (unconditional decision) and how much to contribute conditionally on the average contributions of the other players (conditional decision). To make the elicitation process incentive-compatible, it is necessary to impose a rule whereby the contribution of three of the four members of the group will be their unconditional contribution, while the fourth member's choice will be based on his conditional contribution. This

amounts to assuming that the game has a sequential structure in which one of the players will be able to decide his contribution after having observed the other subjects' choices. The 3SSPD is just a version of this game with only two players.

The 3SSPD game strikes a balance between two competing necessities. On the one hand, the game needs to be rich enough to discriminate subjects who have conditionally cooperative preferences from other types that are frequently observed in experiments of this kind. The simpler two-strategy Prisoner's Dilemma, for example, does not allow us to discriminate conditionally cooperative subjects from the so-called "hump-shaped" types. On the other hand, the game needs to be simple enough to avoid confusion. The main concern is that subjects must choose their conditional strategies several times during the course of the experiment. The conditional contribution table designed by Fischbacher et al. (2001), for example, requires a subject to make 21 distinct choices. It follows that, if the game is played repeatedly for ten periods, a subject has to enter his desired contribution levels up to 210 times. Our setup, by contrast, requires only 30 decisions for the conditional choices in ten periods.

There are also reasons to prefer a two-player game to the more common games in which four or more subjects interact. It has been argued that a two-person sequential Prisoner's Dilemma lends itself more naturally to elicit the conditionally cooperative preferences of the second mover (Andreoni and Samuelson, 2006). To see this, consider that when more than two players interact, a subject will have to take a decision in situations in which part of his opponents cooperate while others defect. It is not obvious what reciprocity dictates in such setting. Finally, a key advantage of the 3SSPD is that it does not require players to understand the logic of dominant strategies (as in the simultaneous Public Goods Game) or of backward induction (as in the Ultimatum Game and any other game that has Nash equilibria that are not subgame-perfect.)

Experimental Design

In all treatments, we used the strategy method to elicit subjects' strategies as both first and second movers. Subjects are informed that they will be selected in either role with equal probability. When choosing their strategy as the second mover, each subject has to fill in a questionnaire in which he states his preferred strategy, conditional upon the strategy chosen by the first mover.

14

The view of the subjects' decision screen can be found in *Supplementary Information.*

Subjects played the 3SSPD game for 10 periods. To minimize the effect of strategic considerations, we follow the standard practice of using a random protocol matching. In each period, subjects are paired randomly to play a match of the game and they disclose their strategies both as first and as second movers. Before each period, we elicited each subject's beliefs about their counterpart's choice. In order to preserve incentive compatibility, we awarded correct guesses by using a quadratic scoring rule (Brier, 1950; Artinger et al., 2010). Additional information on belief elicitation can be found in *Supplementary Information.*

Before the experiment, subjects were required to answer questions aimed at ascertaining their comprehension of the 3SSPD game. Payoffs were created randomly by the computer. Subjects could not proceed to the next step if they failed to provide correct answers to all questions. After completing this stage, four training periods took place, in which subjects played against randomly responding computer players. Following that, subjects were trained on the belief elicitation procedure. At this point the actual experiment took place: the subjects played ten periods of the 3SSPD game against each other. Before the end of the experiment, we gave participants a short questionnaire which contained demographic questions.

Treatments

In order to evaluate the effect of information, we presented subjects with the following information scheme, in a between-subjects fashion: in one treatment (*NoCondInfo*), subjects were informed after each period whether they were the first or the second player, the action they played in the assigned role, the action chosen by their counterpart, and their payoff. In the second treatment (*CondInfo*), we also informed subjects whose roles had been determined as first movers about the conditional strategy chosen by their counterpart who played as the second mover. Every other step is identical for the two treatments.

Data And Materials

Our data, the programs used in the experiment, and the scripts to reproduce graphs and tables are publicly available on the GitHub repository of the

15

study.[1]

References

Andreoni, J. (1995). Cooperation in public-goods experiments: Kindness or confusion? *The American Economic Review 85*(4), 891–904.

Andreoni, J. and J. Miller (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica 70*(2), 737–753.

Andreoni, J. and L. Samuelson (2006). Building Rational Cooperation. *Journal of Economic Theory 127*(1), 117–154.

Artinger, F., F. Exadaktylos, H. Koppel, and L. Sääksvuori (2010). Applying quadratic scoring rule transparently in multiple choice settings: A note. Technical report, Jena Economic Research Papers.

Axelrod, R. (2009). *The Evolution of Cooperation: Revised Edition.* Basic Books.

Bolton, G. E. and A. Ockenfels (2000). ERC: A theory of equity, reciprocity, and competition. *The American Economic Review 90*(1), 166–193.

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthey Weather Review 78*(1), 1–3.

Burton-Chellew, M. N., C. El Mouden, and S. A. West (2016). Conditional Cooperation and Confusion in Public-Goods Experiments. *Proceedings of the National Academy of Sciences of the United States of America 113*(5), 1291–1296.

Burton-Chellew, M. N., H. H. Nax, and S. A. West (2015, February). Payoff-based learning explains the decline in cooperation in public goods games. *Proceedings of the Royal Society B: Biological Sciences 282*(1801), 20142678.

Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics 117*(3), 817–869.

---

[1]`https://www.github.com/seyhunsaral/stabilitycondcoop`

Clark, K. and M. Sefton (2001). The sequential prisoner's dilemma: Evidence on reciprocation. *The Economic Journal 111*(468), 51–68.

Cooper, D. J. and C. K. Stockman (2002). Fairness and learning: an experimental examination. *Games and Economic Behavior 41*(1), 26 – 45.

Cooper, D. J. and C. K. Stockman (2011). History dependence and the formation of social preferences: an experimental study. *Economic Inquiry 49*(2), 540–563.

Croson, R. T. A. (1996). Partners and Strangers Revisited. *Economics Letters 53*, 25–32.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics 10*(2), 171–178.

Fischbacher, U. and S. Gächter (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review 100*(1), 541–56.

Fischbacher, U., S. Gächter, and E. Fehr (2001). Are People Conditionally Cooperative? Evidence From a Public Goods Experiment. *Economics Letters 71*(3), 397–404.

Fisman, R., S. Kariv, and D. Markovits (2007). Individual preferences for giving. *The American Economic Review 97*(5), 1858–1876.

Gächter, S. and B. Herrmann (2009). Reciprocity, culture and human cooperation: Previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society of London B: Biological Sciences 364*(1518), 791–806.

Gale, J., K. G. Binmore, and L. Samuelson (1995). Learning to be imperfect: The ultimatum game. *Games and Economic Behavior 8*(1), 56–90.

Gintis, H., S. Bowles, and R. Boyd (2005). *Moral sentiments and material interests: The foundations of cooperation in economic life.* The MIT Press.

Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, R. McElreath, M. Alvard, A. Barr, J. Ensminger, and et al. (2005). "economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences 28*(6), 795–815.

Herrmann, B. and C. Thöni (2009). Measuring Conditional Cooperation: a Replication Study in Russia. *Experimental Economics 12*(1), 87–92.

Houser, D. and R. Kurzban (2002). Revisiting Kindness and Confusion in Public Goods Experiments. *American Economic Review 92*(4), 1062–1069.

Kimbrough, E. O. and A. Vostroknutov (2016). Norms Make Preferences Social. *Journal of the European Economic Association 14*(3), 608–638.

Kocher, M. G., T. Cherry, S. Kroll, R. J. Netzer, and M. Sutter (2008). Conditional Cooperation on Three Continents. *Economics Letters 101*(3), 175–178.

Krupka, E. L. and R. A. Weber (2013). Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association 11*(3), 495–524.

Kurzban, R. and D. Houser (2005). Experiments investigating cooperative types in humans: A complement to evolutionary theory and simulations. *Proceedings of the National Academy of Sciences 102*(5), 1803–1807.

Neugebauer, T., J. Perote, U. Schmidt, and M. Loos (2009). Selfish-Biased Conditional Cooperation: On the Decline of Contributions in Repeated Public Goods Experiments. *Journal of Economic Psychology 30*(1), 52–60.

Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science 314*(5805), 1560–1563.

Palfrey, T. R. and J. E. Prisbrey (1997). Anomalous behavior in public goods experiments: How much and why? *The American Economic Review 87*(5), 829–846.

Plott, C. R. (1996). Rational individual behavior in markets and social choice processes: The discovered preference hypothesis. In E. C. Kenneth Arrow (Ed.), *The rational foundations of economic behavior. IEA Conference Vol. 114*, pp. 225–50. London: Macmillan.

Rand, D. G., J. D. Greene, and M. A. Nowak (2012). Spontaneous giving and calculated greed. *Nature 489*, 427–430.

Thöni, C. and S. Volk (2018). Conditional cooperation: Review and refinement. *Economics Letters 171*, 37–40.

Volk, S., C. Thöni, and W. Ruigrok (2012). Temporal stability and psychological foundations of cooperation preferences. *Journal of Economic Behavior & Organization 81*(2), 664–676.

Supplementary Information for The Stability of Conditional Cooperation: Egoism Trumps Reciprocity in Social Dilemmas

Andreozzi, Ploner, Saral

October 2019

Classification of Conditional Strategies

We denote each conditional strategy by three letters that represent the conditional responses to L, M, and H, respectively. Since there are only three responses, in total 27 different conditional strategies are possible. As the conditional strategies in our game are simple, the conditional type classification is rather self-evident; thus, we classify conditional types without the need of a calculation or a subjective evaluation.

We say that a player is *selfish* if he/she maximizes his/her own payoff (LLL); *conditional cooperator* if he/she increases its contribution monotonically. A special case of this type of strategy is the *perfect conditional cooperator* (LMH), who responds to the counterpart with the same action the counterpart played. We refer to conditional cooperators who are not perfect conditional cooperators as *imperfect conditional cooperators*[1]. We call a strategy *hump-shaped* if the most generous response of the player is a response to M (LML, LHL, LHM, MHL, MHM). We classified conditional strategies that do not fit any of the definitions above as *other* type. Figure 1 demonstrates the conditional strategies graphically and their classifications.

---

[1] Please note that we classify some of the generous strategies such as MMH and MHH as "imperfect" since they do not perfectly copy the opponents' actions. In that sense, it differs from the "selfish-biased" meaning. In practice, however, this do not make any difference, since no subjects use either of these imperfectly generous strategies.
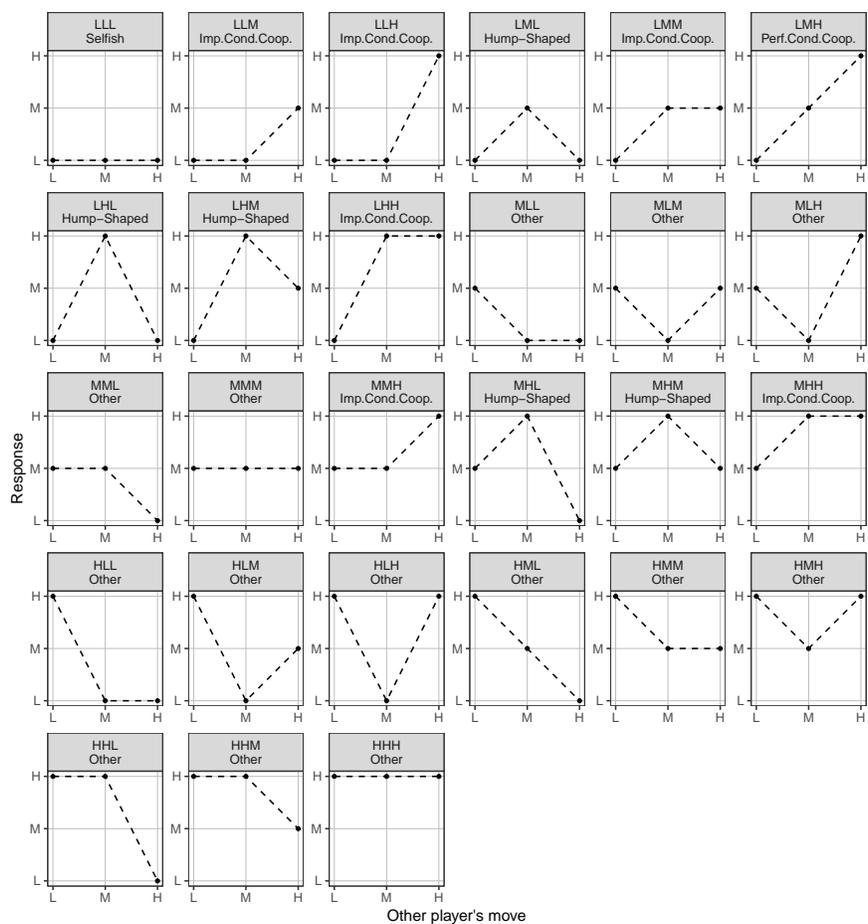
Figure 1: All possible conditional strategies and their classifications

Additional Figures
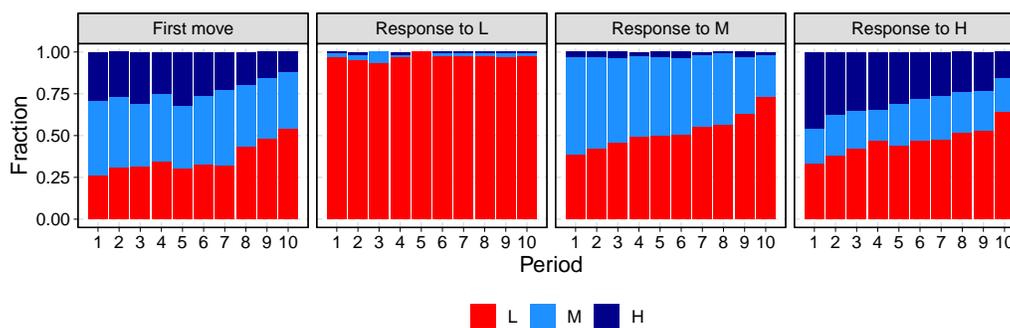
Figure 2: Average fraction of actions chosen as the first player and as the second player for as a conditional response
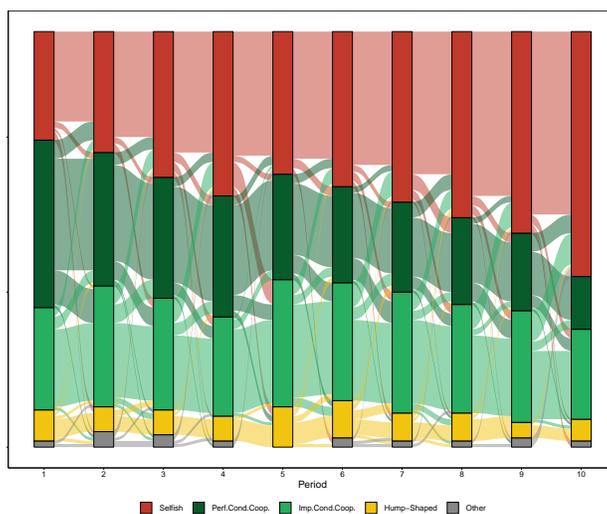

Figure 3: Fraction of conditional strategy types and transitions between types

Belief Elicitation Method

We used the Quadratic Scoring Rule(1) to elicit subjects' beliefs on the opponents' action in an incentive-compatible manner. To increase the

iii

comprehension, we used an alternative presentation of the Quadratic Scoring Rule (2). We designed an interface which contains a slider for each possible outcome to inform subjects about their possible earnings from the belief elicitation task. We also provided a table in the printed instructions which shows the relevant reward and penalty for the probability assigned to an action.

The reward of a subject from one choice in the belief elicitation task in which a subject revealed his/her guess as $p = (p_L, p_M, p_H)$ and action $j \in \{L, M, H\}$ is chosen by the opponent, denoted $Q_A(p)$, and calculated as:

$$Q_j(p) = \alpha + \beta - \beta(1 - p_j)^2 - \beta \sum_{i \neq j}(p_i)^2. \tag{1}$$

We used parameters $\alpha = 20$ and $\beta = 20$ in our experiment, which give each subject the possibility to earn between 0 Experimental Currency (ECU) and 40 ECU in the belief elicitation task stage. For instance, if a subject assigns a 100% probability to the action chosen by the counterpart, he/she receives 40 ECU as a reward and 0 tokens as a penalty, which adds up to 40 ECU. Similarly, if a subject assigns a 50% probability to the action chosen by the counterpart and a 25% probability for the other two actions, he receives 35 ECU as reward and 5 ECU as a penalty for each of the other alternatives that are not selected, which adds up to 25 tokens received from the belief task for this particular action.

Table 1 shows the rewards and penalties table given to subjects.

Table 1: Rewards and Costs of Correct Beliefs according to Quadratic Scoring Rule

Assigned Probability(%)	Reward if chosen (EC)	Cost if not chosen (EC)
100	40	20
95	39.95	18.05
90	39.8	16.2
85	39.55	14.45
80	39.2	12.8
75	38.75	11.25
70	38.2	9.8
65	37.55	8.45
60	36.8	7.2
55	35.95	6.05
50	35	5
45	33.95	4.05
40	32.8	3.2
35	31.55	2.45
30	30.2	1.8
25	28.75	1.25
20	27.2	0.8
15	25.55	0.45
10	23.8	0.2
5	21.95	0.05
0	20	0

Conditional Cooperation and Beliefs

Our incentivized belief elicitation allows us to obtain relatively rich data for beliefs. Although proper scoring rules, such as the Quadratic Scoring Rule, assume risk neutrality, and in reality subjects that are risk averse tend to distribute the risk (3), the change in beliefs over the course of ten periods gives us a hint as to how subjects' beliefs change their conditional strategies.

To simplify our analysis, we aggregate each subject's belief on what the counterpart would do in a specific decision node, as the expected payoff from the action of the opponent in this node. Table 2 shows that if subjects expect a higher response to a transfer of H, they are more likely to be conditional

cooperators and less likely to be selfish. Nevertheless, subjects become more likely to be selfish and less likely to be conditional cooperators in each period. A higher expectation in terms of payoff, as a response to the medium transfer M, is associated with a higher likelihood of choosing a hump-shaped strategy.

Figures 4 to 7 demonstrate the aggregated beliefs and the actual choices made by subjects. We concentrate on the two most interesting cases, in which beliefs concern the first player's unconditional choice (Figure 4) and the second player's choice conditional to the first player choosing the high transfer H (Figure 7). Figure 4 reveals that, on average, subjects' beliefs about the first player's choices are remarkably accurate in the initial rounds and show no tendency of becoming more or less accurate in later rounds. Figure 7 reveals that in the first stages of the game subjects are also clever at guessing the second player's transfer in response to an high transfer. Interestingly, however, their guesses become *less* accurate as the game unfolds, as they become more selfish than they think the other subjects are.

	Dependent Variable		
	isCondCoop	*isSelfish*	*isHumpShaped*
(Intercept)	$-1.64\ (0.75)^{*}$	$1.67\ (0.75)^{*}$	$-7.90\ (1.44)^{***}$
Period	$-0.17\ (0.03)^{***}$	$0.20\ (0.04)^{***}$	$-0.08\ (0.06)$
Treatment *CondInfo*	$0.00\ (0.57)$	$-0.10\ (0.57)$	$0.40\ (0.94)$
Belief - First Player (Uncond.)	$1.93\ (0.71)^{**}$	$-1.25\ (0.71)$	$-1.69\ (1.35)$
Belief - Response to L	$-1.23\ (1.25)$	$-0.69\ (1.27)$	$-0.65\ (1.95)$
Belief - Response to M	$1.08\ (1.14)$	$-2.28\ (1.18)$	$3.50\ (1.78)^{*}$
Belief - Response to H	$4.01\ (0.76)^{***}$	$-3.73\ (0.79)^{***}$	$-1.14\ (1.19)$
AIC	1118.56	1064.38	425.74
BIC	1160.16	1105.98	467.35
Log Likelihood	-551.28	-524.19	-204.87
Num. obs.	1340	1340	1340
Num. groups: subject_unq	134	134	134
Var: subject_unq (Intercept)	8.67	8.45	34.76

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 2: Mixed-Effects Logistic Regression on Binary Variable of Conditional Classification. Beliefs are measured in terms of expected transfers that are scaled to [0,1] interval.
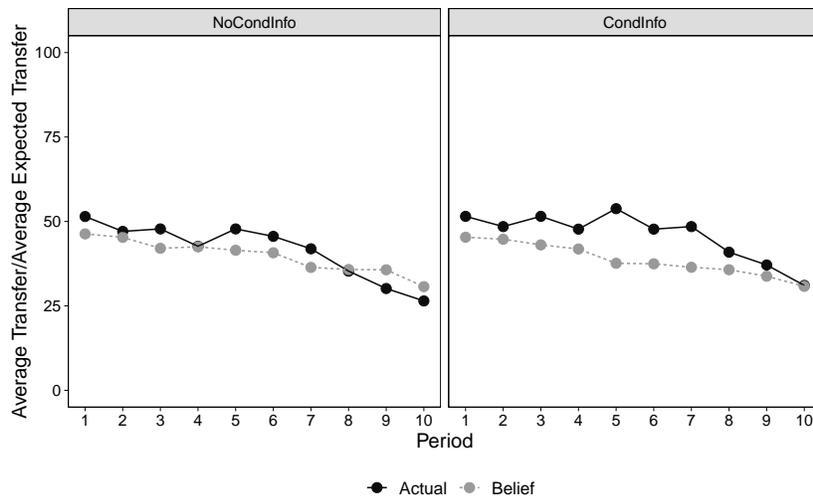
Figure 4: Mean Unconditional Transfer and Beliefs on the Expected Transfers in Each Treatment.
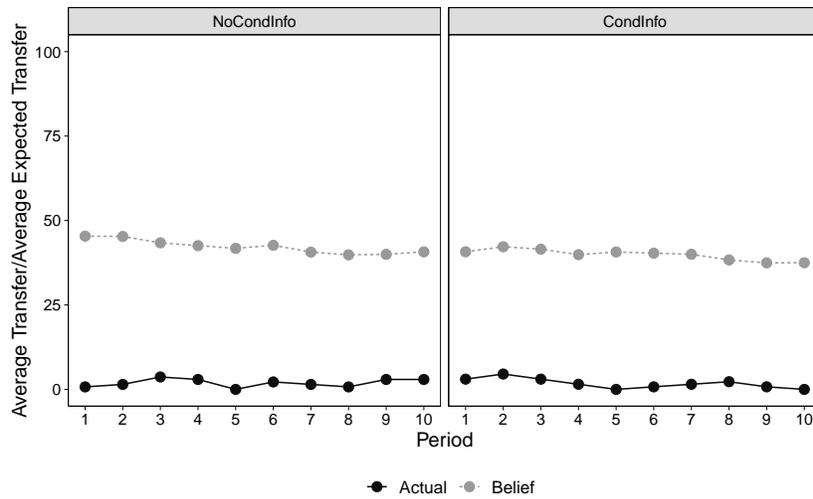

Figure 5: Mean Conditional Transfers and Beliefs on the Expected Transfers in Response to L in Each Treatment.
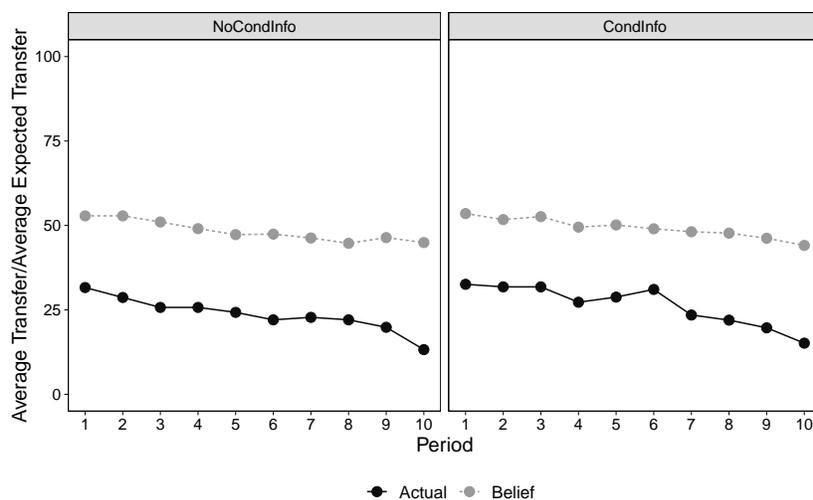
Figure 6: Mean Conditional Transfers and Beliefs on the Expected Transfers in Response to M in Each Treatment.
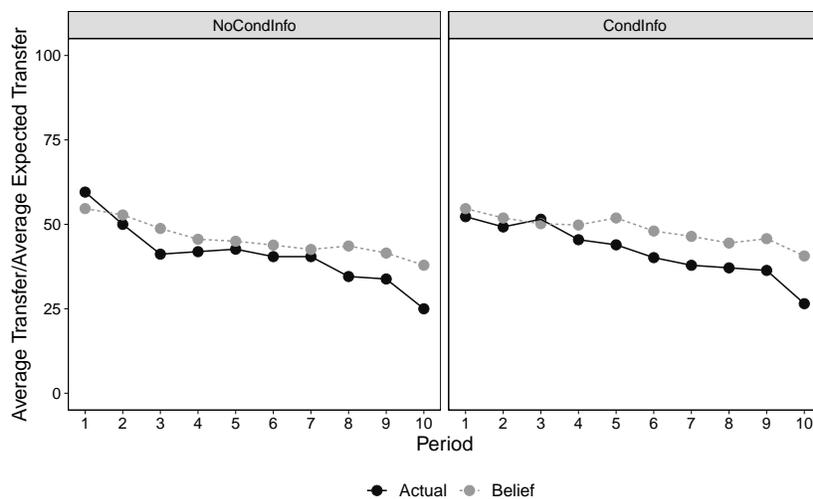

Figure 7: Mean Conditional Transfers and Beliefs on the Expected Transfers in Response to H in Each Treatment.

Comparison of the Proportion of Selfish and Conditional Types in Two Treatments

Period	*CondInfo*	*NoCondInfo*	p-value
1	0.62	0.68	0.62
2	0.61	0.62	1.00
3	0.62	0.50	0.22
4	0.53	0.53	1.00
5	0.56	0.56	1.00
6	0.47	0.56	0.39
7	0.48	0.53	0.73
8	0.47	0.47	1.00
9	0.50	0.41	0.39
10	0.36	0.32	0.76

Table 3: Proportions of Conditional Cooperators in Two Treatments and p-value of Chi-squared Test of Equal Proportions. We see no difference between treatments in any period.

Period	*CondInfo*	*NoCondInfo*	p-value
1	0.29	0.24	0.62
2	0.29	0.29	1.00
3	0.30	0.40	0.34
4	0.39	0.40	1.00
5	0.32	0.37	0.67
6	0.35	0.40	0.69
7	0.41	0.41	1.00
8	0.44	0.46	0.99
9	0.45	0.51	0.60
10	0.56	0.62	0.62

Table 4: Proportion of Selfish Players in Two Treatments and p-value of Chi-squared Test of Equal Proportions. We see no difference between treatments in any period.

Statistical Methods

The statistical analysis was run in the *R environment* (4) and our raw data were imported with the *zTree* package (5). Figures in the Results section are produced via the *ggplot2* package (6). In Table 2 we report p-values obtained of Pearson's Chi-squared tests with Yates' continuity correction on two-way frequency tables. Due to small sample size, we obtained p-values with Monte-Carlo simulation with 2000 replicates. None of the significance levels have changed according to simulated p-values. Thus we report non-simulated values. Table 3 summarizes the outcomes of a Generalized Linear Mixed Model (GLMM) fitted by maximum likelihood (Laplace Approximation) via the *lme4* package (7). Given the dichotomous nature of the dependent variable, we adopted a binomial family specification with a logit link function. To account for repeated choices over rounds, we added a random effect at the individual level. The Akaike information criterion (AIC) is provided for the sake of model comparison.

Subject Screens (English Translation)

Figure 8-15 shows the subject screens' translations to English. The original experiment is conducted in English and z-Tree (8) files are available with zBrac(9) compatible treatment and language files.

Figure 8: Decision Screen as the first player

Figure 9: Decision Screen as the second player
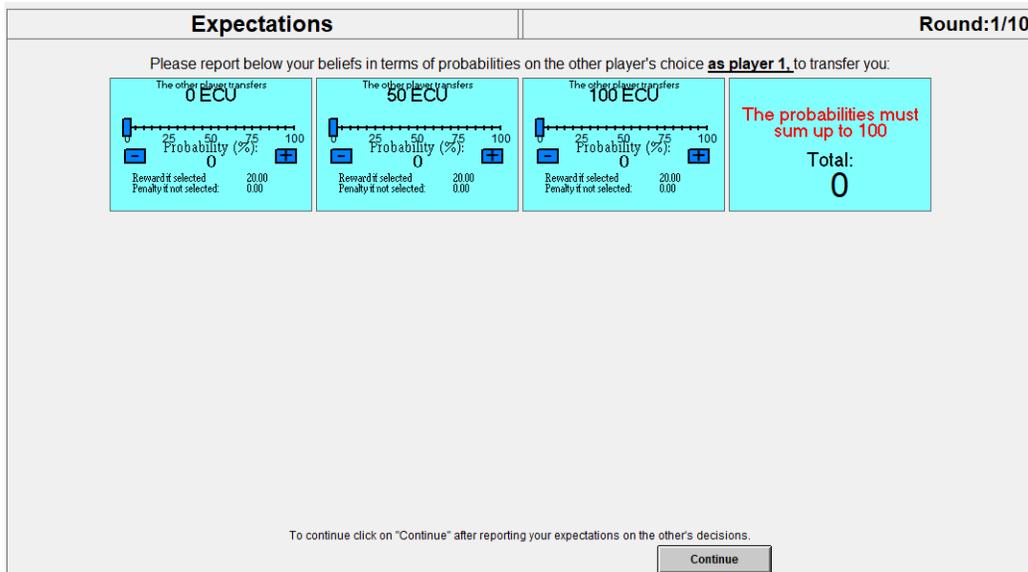
Figure 10: Belief elicitation: Beliefs on the opponent's actions as the first player
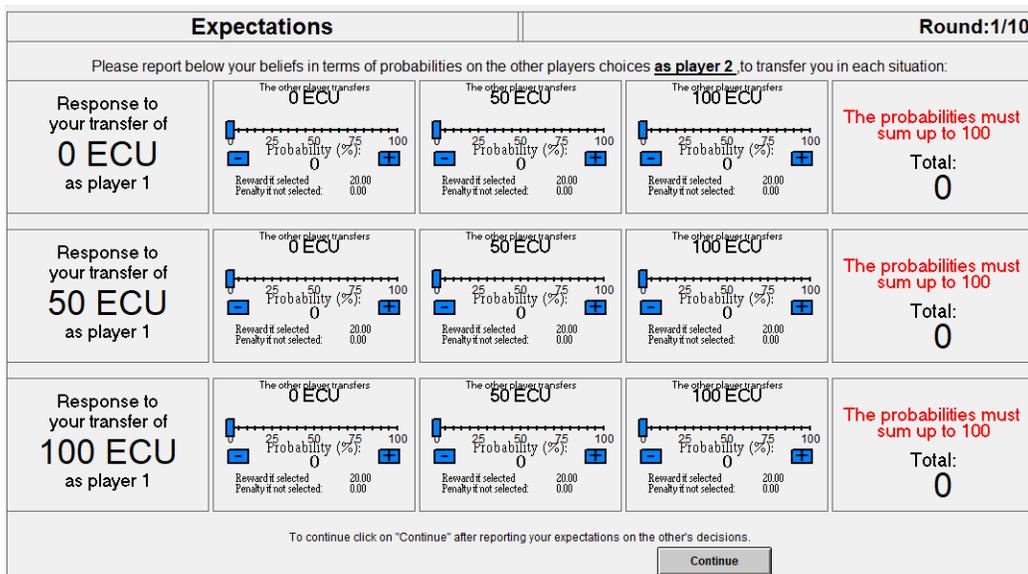
Figure 11: Belief elicitation: Beliefs on the opponent's actions as the second player
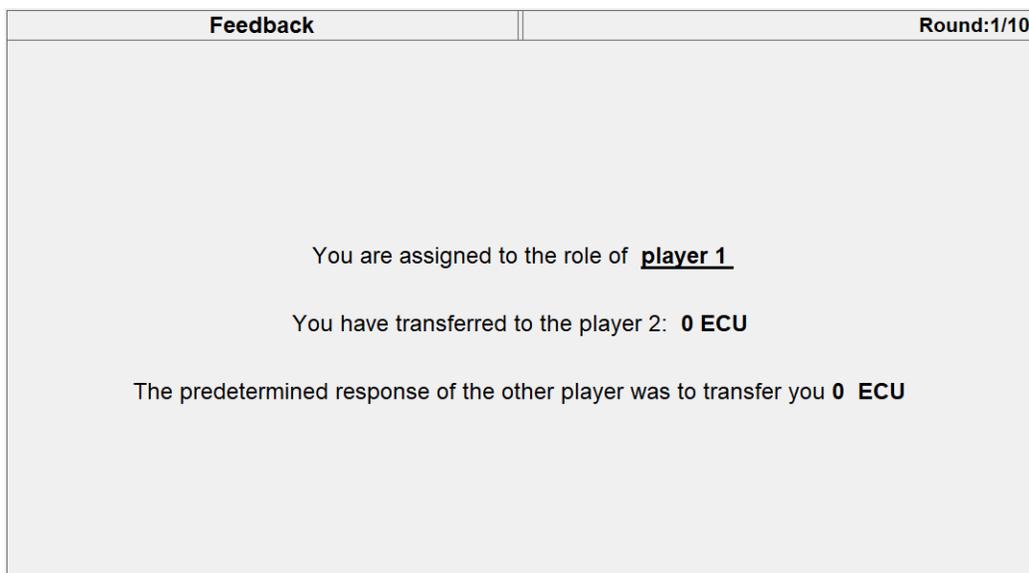
Figure 12: Feedback for the first player in *NoCondInfo* treatment
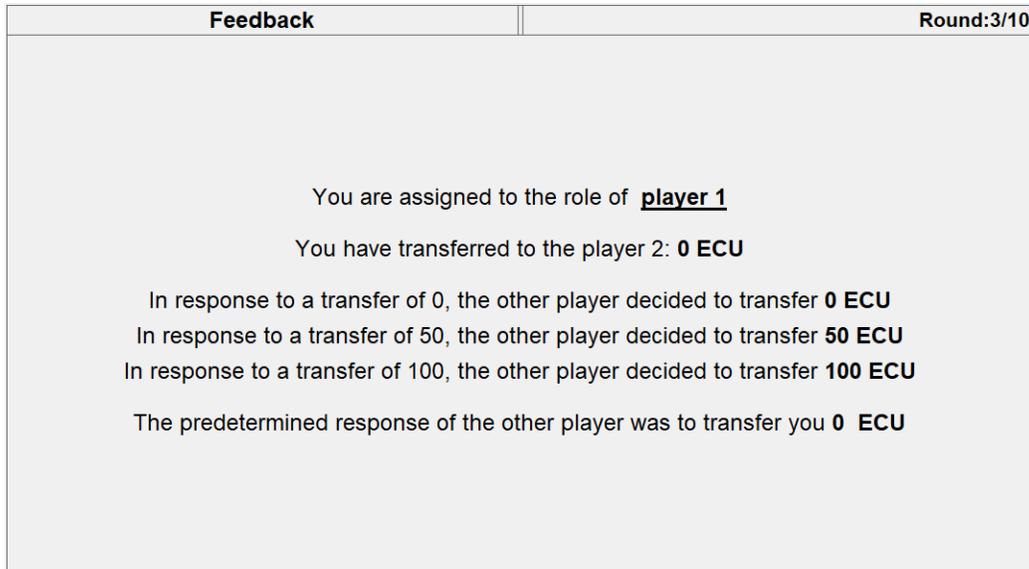
Figure 13: Feedback for the first player in *CondInfo* treatment
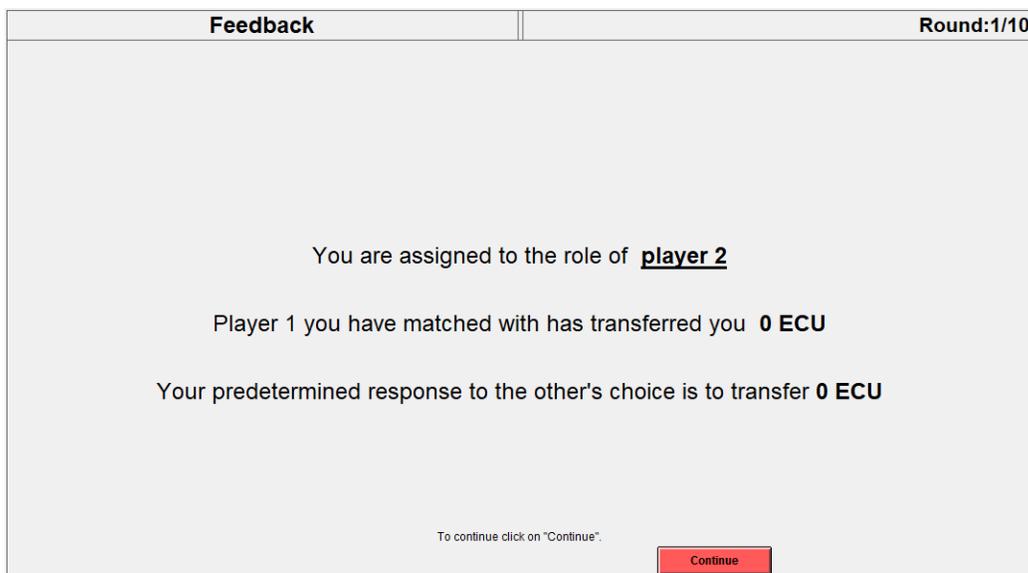
Figure 14: Feedback for the second player
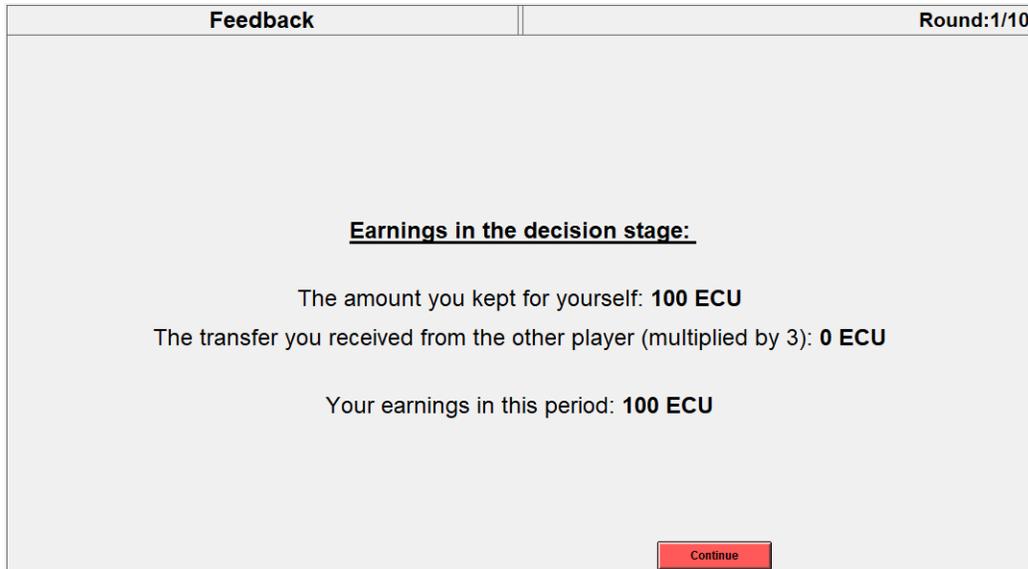
Figure 15: Payoff feedback

English translation of the written instructions

Instructions

General Information

Welcome to the experiment and thank you for your participation. This experiment is funded by **University of Trento** and all the data acquired will be used for scientific purposes.

From now on please remain silent, do not communicate with other participants and raise your hand if you have any question. The use of cell phones and of any other means of communication is forbidden. Please note that those violating these rules will be excluded from the experiment itself and all ensuing payments.

All data collected are held anonymous.

In this experiment, you can earn money depending on your decisions and those of other participants.

You will receive your payment in cash at the end of the experiment. You will be given **3 EUR** for taking part in the experiment. In addition, you may earn more according to the procedure explained below. During the experiment all the payments will be in **"Experimental Currency (ECU)"**. After the experiment, your earnings in ECU will be converted in Euros at the following conversion rate:

1 ECU = 0.05 EUROS

In some parts of the experiments, you will be interacting with other participants. During and after the experiment, your decisions and those of other participants will remain anonymous. This means that your identity of and that of the participant(s) you will interact with will be kept secret. The payment you will earn from the experiment will be given personally and will not be revealed to other participants. Moreover all the data acquired from the experiment will not be associated to the participants' identity.

Structure

The experiment is made of 10 repeated rounds identical to each other and each period will consist of two stages in which you will make decisions: **Decision Stage** and **Expectations Stage**. Prior to 10 identical periods of these two stages, there will be 4 rounds of a **Training Stage,** where you will become familiar with the experiment. At the end of the 10[th] round, you will go through a **Questionnaire Stage** where you will be asked to give feedback about the experiment and to answer a few questions.

The general structure of the experiment can be summarized by the table:

Stage Name	Repetition
Training Stage	**x4**
Expectations Stage	**x10**
Decision Stage	
Questionnaire	x1
Payments	x1

Procedure

Decision Stage

In this stage you will be matched anonymously with a random participant in the lab. In each round, you will have an initial endowment of **100 ECU** and you will be asked to make transfer decisions to the matched participant, as described below.

In each pair, one of the players will be the **Subject 1** and the other will be the **Subject 2.** *However, you will make your choices both as Subject 1 and as Subject 2 in each round*. After you and your partner chose as Subject 1 and as Subject 2, the roles will be assigned and decisions will be implemented accordingly.

Subject 1 will have the option to transfer **0**, **50** or **100 ECU** to the other and to keep the rest for himself/herself. The amount that Subject 1 transfers will be multiply by **3.** This means that if Subject 1 selects to transfer **0**, he/she will keep all 100 ECU to himself/herself and the other will get nothing; if he/she selects to transfer **50 ECU** to the other, he/she will keep the **50 ECU** to himself/herself and the other will receive **150 ECU;** finally, if he/she decides to transfer **100 ECU,** he/she will keep nothing and the other will receive **300 ECU.** The decision screen of Subject 1 will be the following:

[First Player Decision Screenshot Here]

Also **Subject 2** will have the option to transfer **0, 50** or **100 ECU,** as explained above. Similarly to Subject 1, also the amount transferred by Subject 2 to Subject 1 will be multiplied by **3.** However, Subject 2 will have the opportunity to choose his/her transfer to the other, conditional upon what Subject 1 chose transfer.

So Subject 2 will make three decisions:

To transfer **0, 50** or **100** ECU, if the other chooses **0 ECU;**
To transfer **0, 50** or **100** ECU, if the other chooses 5**0 ECU;**
To transfer **0, 50** or **100** ECU, if the other chooses **100 ECU;**

The actual transfer of Subject 2 will be the one made in correspondence to the actual choice of Subject 1.

The decision problem of Subject 2 will be shown in a decision screen like the following:

[Second Player Decision Screenshot Here]

As both subjects have 100 ECU and the amount transferred will be multiplied by 3, the earnings for both players will be as in the following table:

		Choice of Subject 2 conditional upon choice of Subject 1		
		0 ECU	**50 ECU**	**100 ECU**
Choice of Subject 1	**0 ECU**	Subject 1 earns 100 ECU, (100 kept, 0 received) Subject 2 earns 100 ECU. (100 kept, 0 received)	Subject 1 earns 250 ECU, (100 kept,150 received) Subject 2 earns 50 ECU. (50 kept, 0 received)	Subject 1 earns 400 ECU, (100 kept,300 received) Subject 2 earns 0 ECU. (0 kept, 0 received)
	50 ECU	Subject 1 earns 50 ECU, (50 kept, 0 received) Subject 2 earns 250 ECU. (100 kept, 150 received)	Subject 1 earns 200 ECU, (50 kept, 150 received) Subject 2 earns 200 ECU. (50 kept, 150 received)	Subject 1 earns 350 ECU, (50 kept, 300 received) Subject 2 earns 150 ECU. (0 kept, 150 received)
	100 ECU	Subject 1 earns 0 ECU, (0 kept, 0 received) Subject 2 earns 400 ECU. (100 kept, 300 received)	Subject 1 earns 150 ECU, (0 kept, 150 received) Subject 2 earns 350 ECU. (50 kept, 300 received)	Subject 1 earns 300 ECU, (0 kept, 300 received) Subject 2 earns 300 ECU. (0 kept, 300 received)

After both subjects had chosen both as Subject 1 and as Subject 2, the computer will randomly assign roles and the choice made by Subject and the choice made by Subject 2 conditional upon that of Subject 1 will be implemented. The matching of choices will define the earnings in the round.

Note that in each period you have the same chance to be the Subject 1 or the Subject 2. Therefore it is likely that in some rounds you will be selected as Subject 1 and in others as Subject 2.

The interaction will be repeated 10 times and in each round you will be randomly matched with another participant in the session.

At the end of the experiment, one of the rounds will be selected and the earnings from that round will define your final payment.

Feedback about the Decisions

[TREATMENT NoCondInfo]: After each round, you will receive a feedback about the other's decisions: if you are selected to be the Subject 1, you will see the other's response to your choice, and Subject 2 will see your choice. If you are selected to be Subject 2, you will see the other's choice and Subject 1 will see your corresponding response to his/her choice. Feedback will be delivered in a screenshot like the following:

[Treatment NoCondInfo Feedback Screenshot Here]

[TREATMENT CondInfo]: After each round, you will receive a feedback about the other's decisions: if you are selected to be the Subject 1, you will see the other's conditional choice to each possible choice of yours, and Subject 2 will see your choice. If you are selected to be Subject 2, you will see the other's choice and Subject 1 will see your conditional choices to each possible choice of him/her. Feedback will be delivered in a screenshot like the following:

xvii

[Treatment CondInfo Feedback Screenshot Here]

Summary of the Decision Stage

- Both subjects select the amount to transfer as Subject 1 (0 , 50 or 100 ECU).
- Both subjects select their choices as Subject 2 (0, 50, or 100 ECU conditional upon each 3 possible transfers of the Subject 1).
- One subject will be selected as Subject 1, one subject will be selected as Subject 2, randomly by the computer.
- The choices of the two subjects are implemented.
- Subjects receive a feedback about the choices of the other.

Expectations Stage

In each round, before the Decision Stage, you will take part in the Expectations Stage. You will be asked to assess for each of the options available how likely it is that it is going to be chosen by the other. Since in each round the other has three options (to transfer 0 ECU, 50 ECU or 100 ECU), you will be asked to report a probability for each option that the other has. The sum of probabilities must be equal to 100. The earnings in this stage are proportional to the accuracy of your estimate. You can report the probabilities either by using the sliders (dragging the blue handle to the left and right) or the plus and minus signs in each field.

Since in the Decision Stage, each player will make four decisions, you will be asked to state your beliefs for each of the four different choices. One for the other's choice as Subject 1 (first screen) and three for the other's choice as Subject 2 (second screen.

You will see a screen as the one below to state your beliefs about the other's decision as Subject 1:

[First Player Expectations Screenshot Here]

After completing this screen, you will be shown a screen like the following to predict the other's choices as Subject 2:

[Second Player Expectations Screenshot Here]

For each decision taken by the other, you will earn an amount of money for the probability you assigned to the choice taken. The larger the probability you give to the chosen option, the more ECU you will earn as a reward (values are shown on the second column of the table below). However for the probabilities you assign to a decision that is not chosen, you will incur a cost. The larger the probability you give to the unchosen option, the more ECU will be subtracted from your earnings (values are shown on the third column of the table below)

Since there are three options for each choice (to transfer 0, 50 or 100 ECU) you will get a reward for the right choice, while you will pay a cost for the other two choices, that will be subtracted from your earnings. procedure to compute rewards and costs are the same for all choices and for all rounds.
The table below shows rewards and costs associated to each assigned probability:

Assigned Probability to an option (%)	Reward if this option is chosen (ECU)	Cost if this option is not chosen (ECU)
100	40	20
95	39.95	18.05
90	39.8	16.2
85	39.55	14.45
80	39.2	12.8
75	38.75	11.25
70	38.2	9.8
65	37.55	8.45
60	36.8	7.2
55	35.95	6.05
50	35	5
45	33.95	4.05
40	32.8	3.2
35	31.55	2.45
30	30.2	1.8
25	28.75	1.25
20	27.2	0.8
15	25.55	0.45
10	23.8	0.2
5	21.95	0.05
0	20	0

On the screens where you make your choice, you will be able to see the rewards and the costs for the probability you assign to each option and the total earnings as well. However if you prefer, you can refer to the table during the experiment.

EXAMPLE:
Assume that you stated that the probability that the other chooses to transfer 50 ECU is 80%, the probability to transfer 0 ECU is 10% and the probability to transfer 100 ECU is 10%. If the other actually chooses to transfer 50 for this decision:

You will receive the reward for the 80% assigned to the right choice, that is 39.2 ECU
You will pay the costs for the 10% assigned to the two wrong choices, that is 0.2 ECU for each.
Thus, your earning will be 39.2 – 0.2 – 0.2 = 38.8 ECU.

EXAMPLE:
Assume that you stated that the probability that the other chooses to transfer 0 ECU is 50%, the probability to transfer 50 ECU is 30% and the probability to transfer 100 ECU is 20%. If the other actually chooses to transfer 50 ECU for this decision:

You will receive the reward for the 30% assigned to the right choice, that is 30.2 ECU.
You will pay the costs for the 50% assigned to the first wrong choice, that is 5 ECU.
You will pay the costs for the 20% assigned to the second wrong choice, that is 0.8 ECU.
Thus, your earning will be 30.2 – 5 – 0.8 = 24.4 ECU.

It can be seen that the maximum amount you can earn for your choice is 40 ECU when you assign 100% probability to an option chosen (40 – 0 – 0) and minimum amount that you can get is 0 ECU assigning 100% to an option that is not chosen (20 – 20 – 0).

Please note that with this mechanism it is in your interest to report your true beliefs.

At the end of the experiment, from all the guesses you make, one of the guesses will be selected and you will be paid according to accuracy that guess. There will be no feedback after Expectations Stages and your earnings will be notified at the end of the experiment and will be added to your payment.

Training Stage

In order to get familiar with the Decision Stage and the Expectations Stage, there will be a training stage of 4 rounds. In this stage instead of being matched with participants in the lab, you will interact with a computer choosing randomly.

In the first two rounds you will only go through the Decision Stage and in the other two rounds you will see the Expectations Stage first and the Decision Stage afterwards.

The choices made in the Training Stage do not contribute to your final earnings.

After the 4 rounds of the Training Stage, you will start the actual experiment interacting with other players for 10 rounds.

We reached the end of the instructions. Please raise your hand if you have any question.

References

[1] G. W. Brier, "Verification of forecasts expressed in terms of probability," *Monthey Weather Review*, vol. 78, no. 1, pp. 1–3, 1950.

[2] F. Artinger, F. Exadaktylos, H. Koppel, and L. Sääksvuori, "Applying quadratic scoring rule transparently in multiple choice settings: A note," tech. rep., Jena economic research papers, 2010.

[3] T. Offerman, J. Sonnemans, G. Van de Kuilen, and P. P. Wakker, "A truth serum for non-bayesians: Correcting proper scoring rules for risk attitudes," *The Review of Economic Studies*, vol. 76, no. 4, pp. 1461–1489, 2009.

[4] R Core Team, *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria, 2019.

[5] O. Kirchkamp, "Importing z-Tree data into r," *Journal of Behavioral and Experimental Finance*, vol. 22, pp. 1–2, 2019.

[6] H. Wickham, *ggplot2: elegant graphics for data analysis.* Springer, 2016.

[7] D. Bates, M. Mächler, B. Bolker, and S. Walker, "Fitting linear Mixed-effects models using lme4," *Journal of Statistical Software*, vol. 67, no. 1, pp. 1–48, 2015.

[8] U. Fischbacher, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental economics*, vol. 10, no. 2, pp. 171–178, 2007.

[9] A. S. Saral and A. M. Schröter, "zBrac—A multilanguage tool for z-tree," *Journal of Behavioral and Experimental Finance*, vol. 23, pp. 59–63, 2019.