

Anderhub, Vital; Güth, Werner; Engelmann, Dirk

Working Paper

An experimental study of the repeated trust game with incomplete information

SFB 373 Discussion Paper, No. 1999,97

Provided in Cooperation with:

Collaborative Research Center 373: Quantification and Simulation of Economic Processes, Humboldt University Berlin

Suggested Citation: Anderhub, Vital; Güth, Werner; Engelmann, Dirk (1999) : An experimental study of the repeated trust game with incomplete information, SFB 373 Discussion Paper, No. 1999,97, Humboldt University of Berlin, Interdisciplinary Research Project 373: Quantification and Simulation of Economic Processes, Berlin, <https://nbn-resolving.de/urn:nbn:de:kobv:11-10046848>

This Version is available at:

<https://hdl.handle.net/10419/61770>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

An Experimental Study of the Repeated Trust Game with Incomplete Information*

Vital Anderhub[†], Dirk Engelmann[‡], and Werner Güth[†]

Abstract

In the trust game first player 1 decides between non-cooperation or trust in reciprocity and then, in the latter case, player 2 between exploiting player 1 or rewarding him. In our experiment player 2 can be a notorically rewarding player (this type is implemented as a robot strategy) or a human participant who may decide opportunistically. To allow for reputation formation, this base game is played repeatedly. Learning can be analysed since participants play successively several repeated games with changing partners. In our computerized experiment participants can explicitly rely on mixed strategies which allows testing the qualitative and quantitative aspects of reputation equilibria also at an individual level.

*The financial support of the Deutsche Forschungsgemeinschaft (SFB373, C5) and the EU-TMR Research Network ENDEAR is gratefully acknowledged.

[†]Humboldt-Universität zu Berlin, Wirtschaftswissenschaftliche Fakultät, Institut für Wirtschaftstheorie III, Spandauer Str. 1, D – 10178 Berlin, Germany.

[‡]Humboldt-Universität zu Berlin, Wirtschaftswissenschaftliche Fakultät, Institut für öffentliche Wirtschaft und Wirtschaftspolitik, Spandauer Str. 1, D – 10178 Berlin, Germany, email: engelman@wiwi.hu-berlin.de.

1 Introduction

The trust game is a simple sequential game with just two players 1 and 2 who can engage in some mutually profitable cooperation if they can trust each other. As illustrated by Figure 1, first player 1 chooses between N (on-cooperation) or T (trust in reciprocity). In case of N the game is over and both players receive s . In the case of T now player 2 decides between E (exploiting) and R (rewarding). Whereas R yields r for both players, only player 2 receives t and player 1 nothing in case of E .




Figure 1: The trust game with $t > r > s > 0$

Although a pen and paper-experiment (Güth, Ockenfels, and Wendel, 1997) with the large exploitation incentive $t-r = DM 40.00$ indicated a low willingness to trust (T) and nearly no tendency to reward (R), one naturally would expect that at least some people feel obliged to reward a trusting player 1. If such rewarding behavior by player 2 cannot be excluded, reputation equilibria (see Kreps, Milgrom, Roberts, and Wilson, 1982) predict that even an opportunistic player will at least initially mimic such a rewarding player 2 when the trust game is played repeatedly by the same partners. By imitating an always rewarding player 2, such an opportunistic player 2 strategically builds up the reputation (the conditional probability) of being of the always rewarding type. This may induce player 1 to continue initial cooperation in the course of the finitely repeated trust game.

Earlier experimental studies of reputation formation have used other base games like finitely repeated Prisoner's Dilemma games (Andreoni and

Miller, 1993) or centipede games (McKelvey and Palfrey, 1992) or have used similar base games but a quite different experimental design (e.g. Camerer and Weigelt, 1988, Neral and Ochs, 1992).

Our main treatment explicitly introduces the possibility of a (second mover) type who always feels obliged to reward trust. Instead of changing systematically its positive a priori probability we have preferred to vary its relevance by changing systematically the (finite) number of repetitions.

Trust is automatically rewarded with probability p whereas with complimentary probability $1 - p$ player 2 is free to choose. Actually when trust was automatically rewarded, it meant that one did not encounter a human player 2 but a robot strategy. This is similar to Andreoni and Miller (1993), whereas Neral and Ochs (1992) let the always rewarding type be played by an experimenter and Camerer and Weigelt (1988) change the payoffs for a part of the players to make rewarding a dominant strategy. McKelvey and Palfrey (1992) have simply introduced the possibility of an altruistic type and estimated its a priori-probability needed to explain the experimentally observed behavior. In our instructions we did not explicitly mention that the automatically rewarding player 2 is a robot (without deceiving the subjects by indicating the presence of a real partner either). Although it does not influence the solution behavior, knowing that there is no human partner might affect behavior.

Andreoni and Miller (1993) have introduced robot strategies playing tit-for-tat. However, imitating a tit-for-tat strategy is more complicated than just imitating a simple always rewarding strategy in the repeated trust game. Hence we believe that our experiment constitutes an easier test of the reputation building hypothesis.

We conducted a control treatment with no automatically rewarding type of player 2, i.e. $p = 0$. Apart from this modification the design was the same. The motivation for this control treatment is twofold. First, we wanted to check whether the play without automatically rewarding players qualitatively resembles the play in the presence of such players. If this is the case, reputation formation could also be an explanation for the observed behavior in the absence of such players. Second, the existence of robot players somehow might influence behavior. This possible influence can be ignored if the results in the control treatment are qualitatively similar.

One further novelty of our design is that we explicitly allow players to use mixed strategies. Previous attempts to elicit individual mixed strategies (e.g. Albers and Vogt, 1997, and Ochs, 1995) have used normal form games like 2x2-bimatrix games. Observing individual mixing behavior allows to test the qualitative and quantitative predictions of reputation equilibria on

the individual level (and not just for the entire population).

Like earlier studies we find that reputation equilibria can quite well explain the aggregate data of experienced subjects. However, individual data, in particular individual mixed strategies, differ substantially from the prediction.

In the following we first analyze finitely repeated trust games with incomplete information (section 2). The procedural details of our computerized experiment are described in section 3. After a description of the main behavioral tendencies in section 4 we test statistically some of the qualitative aspects of reputation equilibria (section 5). Our conclusions relate our findings to the relevant literature.

2 Repeated trust games with incomplete information

Let p with $0 < p < 1$ denote the probability for a player 2-type who automatically rewards trust, i.e. who always reacts by R to player 1's move of T . When choosing between N and T , player 1 is, however, unaware of whether he confronts this type of player 2 or a player 2 who opportunistically decides between E and R . All what player 1 knows are the probabilities p of a notorically rewarding player 2 and $1 - p$ of a free to choose-type of player 2.

In our experiment the notorically rewarding type of player 2 is implemented by a corresponding robot strategy whereas the free to choose-type is a human participant whose monetary rewards are those given by Figure 1. Such a human participant may, of course, also feel morally obliged to reward trust in reciprocity. Our game theoretic benchmark solution relies on the usual assumption, namely that human decision makers are only interested in their own monetary reward and perfectly rational.

To illustrate reputation equilibria in the simplest way assume that the trust game is not only played once, but twice by the same partners. Generally let y_t denote the probability by which the free to choose-type of player 2 (we will refer to him as player 2) rewards trust in the t^{th} period. Similarly, let x_t denote player 1's probability of using T in period t . In period 2, after the decisions (T, R) in period 1, player 1's conditional probability $1 - p_2$ of confronting a real player 2 is given by

$$1 - p_2 = \frac{(1 - p) y_1}{p + (1 - p) y_1}.$$

Thus $y_1 = 0$ yields $1 - p_2 = 0$ which would imply the optimal decisions $x_2 = 1$ and $y_1 = 1$, contradicting the initial assumption $y_1 = 0$. So the solution must rely on $y_1 > 0$. The assumption $y_1 = 1$ implies $1 - p_2 = 1 - p$, i.e. player 1 will use T in period 2 if $p > \frac{s}{r}$ and N if $p < \frac{s}{r}$. The latter, however, implies $y_1 = 0$, i.e. a contradiction. Thus the solution for $y_1 = 1$ requires $p > \frac{s}{r}$ and $x_2 = 1$.

Finally assume that $0 < y_1 < 1$. If this implied $x_2 = 1$, then $y_1 = 1$ would be optimal. Thus the only consistent constellation with $0 < y_1 < 1$ is $x_2 < 1$. This means that player 1 is indifferent between T and N in period 2, therefore $p_2 = \frac{s}{r}$ and hence

$$y_1 = \frac{p}{1-p} \cdot \frac{r-s}{s}.$$

To render this choice y_1 optimal, player 2 must be indifferent between E and R in the first period which requires

$$x_2 = \frac{t-r}{t-s}.$$

Anticipating y_1 and x_2 the optimal choice for player 1 in period 1 is $x_1 = 1$ if $p > \left(\frac{s}{r}\right)^2$ and $x_1 = 0$ if $p < \left(\frac{s}{r}\right)^2$.

From these results one easily concludes that for all generic parameters p with $0 < p < 1$ there exists just one perfect Bayesian equilibrium (Fudenberg and Tirole, 1991, section 8.2, or perfect equilibrium, Selten, 1975, or sequential equilibrium, Kreps and Wilson, 1982, which are equivalent in the present game).

If for $\left(\frac{s}{r}\right)^2 > p > 0$ player 1 would mistakenly choose T_1 rather than N_1 , the play would continue like after the move T_1 in the parameter region $\frac{s}{r} > p > \left(\frac{s}{r}\right)^2$. If, conversely, in the region $\frac{s}{r} > p > \left(\frac{s}{r}\right)^2$ player 1 chooses mistakenly N_1 instead of T_1 , he would continue with N_2 as in the region $\left(\frac{s}{r}\right)^2 > p > 0$. A wrong choice N_1 for $1 > p > \frac{s}{r}$ would, however, be followed by T_2 , i.e. would not change the behavior in the second period $t = 2$.

More generally, let m denote the number of successive periods of playing the trust game by the same partners. When describing the implications of reputation equilibria, we distinguish what happens on the solution path and what would happen off the solution path. Experimentally one cannot expect rational behavior and, even less likely, common knowledge of rationality. It is therefore important to know the conditionally rational behavior off the solution path as it has been qualitatively described above.

On the solution path one will start with trust if there will be trust later on. Off the solution path it is possible that there is first no cooperation (due

to mistakes), but that trust is chosen and rewarded later. Here trust and rewarding trust simply mean that the respective behavior is realized with positive probability. The main feature of reputation equilibria is that after an initial phase of full cooperation, i.e. (T, R) in the initial periods $1, \dots, t-1$, or off the solution path, a phase of mutual mixing starts with $0 < y_\tau < 1$ for $t \leq \tau < m$ and $0 < x_\tau < 1$ for $t < \tau \leq m$.

Two important aspects of this mixing phase are the constancy of player 1's mixing, i.e.

$$x_\tau = \frac{t-r}{t-s} \text{ for all } t < \tau \leq m$$

and the constancy of the product of the y_τ -values, namely

$$y_t \cdot y_{t+1} \cdot \dots \cdot y_{m-1} = \frac{p}{1-p} \cdot \frac{r-s}{s}.$$

Of course, the latter equality does not suffice to actually derive the unique reputation equilibrium for all generic probability parameters p with $0 < p < 1$. This can be done by repeatedly using the formula

$$1 - p_{\tau+1} = \frac{(1 - p_\tau) y_\tau}{p_\tau + (1 - p_\tau) y_\tau}$$

defining player 1's posterior probability $1 - p_{\tau+1}$ in period $\tau+1$ of confronting a real player 2 after observing (T, R) in period τ (Engelmann, 1999, solves such games). The specific solutions will be described after introducing our experimental design, especially the numerical specification of the parameters s, r, t, p , and m defining a repeated trust game with incomplete information.

3 Experimental Design

The computerized experiment always relied on

$$s = \frac{1}{2}, r = \frac{3}{4}, t = 1, \text{ and } p = \begin{cases} \frac{1}{3} & \text{in the main treatment} \\ 0 & \text{in the control treatment} \end{cases}$$

so that $\frac{s}{r} = \frac{2}{3}$ and $(\frac{s}{r})^3 = \frac{8}{27} < p = \frac{1}{3} < \frac{4}{9} = (\frac{s}{r})^2$. We only vary the number m of successive periods for playing the trust game by the same two partners. Whether or not player 2 is actually morally unrestricted is experimentally

uncontrolled. If, for instance, trust is rewarded in the last period m , i.e. if $y_m > 0$, this is inconsistent with opportunism.

The m periods of the basic trust game define one (repeated) game. Each participant played several repeated trust games with changing partners where the numbers of repetitions were

$$m = 3, 6, 2, 10, 3 \text{ and } 6.$$

To inspire learning a participant constantly assumed the role of player 1 or player 2. All participants received the same instructions (see Appendix A for an English translation) informing them that only $\frac{2}{3}$ of all players 1 confronted a free to choose-type partner as player 2 whereas $\frac{1}{3}$ of all players 1 could be sure that T is rewarded by R (implemented as a robot strategy). In each new repeated game it was again randomly decided whether or not a particular player 1 was matched with a real player or not. In the control treatment with $p = 0$ such information was, of course, missing.

A participant was not restricted to pure strategies, but could explicitly randomize. More specifically, a participant had to specify the probabilities in percent for his two decision alternatives. Since reputation equilibria at least partly rely on mixed strategies, one can thus check their qualitative and quantitative aspects not only populationwise, but also at an individual level.

For our numerical specification the unique reputation equilibria depend on m as follows:

$$\begin{aligned} m = 2 : & \quad ((N_1, N_2), (y_1 = \frac{1}{4}, E_2)) \\ m = 3 : & \quad ((T_1, x_2 = \frac{1}{2}, x_3 = \frac{1}{2}), (y_1 = \frac{5}{8}, y_2 = \frac{2}{5}, E_3)) \\ m = 6 : & \quad ((T_1, T_2, T_3, T_4, x_5 = \frac{1}{2}, x_6 = \frac{1}{2}), (R_1, R_2, R_3, y_4 = \frac{5}{8}, y_5 = \frac{2}{5}, E_6)) \\ m = 10 : & \quad ((T_1, T_2, T_3, T_4, T_5, T_6, T_7, T_8, x_9 = \frac{1}{2}, x_{10} = \frac{1}{2}), \\ & \quad (R_1, R_2, R_3, R_4, R_5, R_6, R_7, y_8 = \frac{5}{8}, y_9 = \frac{2}{5}, E_{10})) \end{aligned}$$

Independent of m in the range $m \geq 3$ the mixing phase of both players extends over two periods, namely the last two periods for player 1 and the third and second to last period for player 2.

We did not expect such clear results. Players who engage in randomization are likely to do so over more periods. An important qualitative aspect of reputation equilibria is, however, that after an initial phase of nearly certain cooperation, its likelihood should drop considerably when the end is near.

Another interesting qualitative aspect is that behavior in the end phase does not depend on m when $m \geq 3$. Path dependence, e.g. in the sense that a longer initial phase of cooperation induces more cooperation during the end phase, would suggest a different result (Güth, Ockenfels, and Wendel, 1993, report such behavior).

A human player 2 may not actually feel free to choose E , but rather committed to reward 1's trust. A purely intrinsic obligation to reciprocate should not depend on m suggesting that a participant, who rewards trust always or at least mostly with high probability, will do so regardless of the repetition number $m = 2, 3, 6$, or 10.

To exclude unwarranted repeated game effects resulting from playing repeatedly (repeated) games a participant confronted a new partner in each (repeated) game. Also indirect reputation effects were prevented by the matching procedure. More specifically, for each session nine players 1 and six players 2 were invited and were matched accordingly. The five sessions of the main treatment thus provide data of altogether 45 players 1 and 30 players 2. Each of the $p = 0$ sessions of the control-treatment involved six players 1 and six players 2. A disadvantage of the matching scheme is that the behavior of two participants cannot be regarded as independent. In a rigorous statistical sense only the five sessions in the main treatment and the three sessions in the control treatment qualify as independent observations. Our statistics that assume individual decisions by different participants as independent should therefore be considered with care. However, they can only be wrong in the direction of finding significance where a stricter test would not. When we do not find significant results a test that takes into account possible dependences would not either.

4 Behavioral tendencies

Let $x_t(i)$ and $y_t(i)$ denote participant i 's probability of trust (T), respectively of reward (R). The mean probability of trust and reward are

$$x_t = \sum_i x_t(i)/N_1 \text{ and } y_t = \sum_i y_t(i)/N_2,$$

where $N_1 = 45$ and $N_2 = 30$ are the numbers of participants playing the role of player 1 or player 2, respectively. How x_t and y_t develop over time is illustrated by Figure 2. Here only the x_t -choices of the players 1 confronted with real players 2 are considered. Since participants play successively repeated games with $m = 3, 6, 2, 10, 3$, and 6 periods, they play altogether

30 periods. A vertical dashed line in Figure 2 indicates the last period of a (repeated) game. The end or termination effect (x_t drops when the last period of the game is near) is minor only for the first two games. Afterwards there is always a sharp decline of x_t . A reason for this is suggested by the development of y_t . From the end of the second game on (period 9) players 2 on average reward only rarely in the last period of a (repeated) game.




Figure 2: Average trust- and reward-probabilities per period

Did players 1 confronting a real player 2 behave differently than players 1 facing a robot strategy? In Figure 3 the (non-)dashed line illustrates the x_t -development of the players 1 confronting a (real partner) robot strategy. On average the dashed line lies above the non-dashed one what can be attributed to the path dependence (a robot never exploits whereas real players 2 sometimes do, see the y_t -development in Figure 2). However, these differences are significant only in periods 3, 19, and 21 (Mann-Whitney-U test, 5%-level).




Figure 3: Average trust per period (confronting a real partner or a robot)

Figure 3 gives the number of trust moves T (dashed bars) and the number of those which were not rewarded (solid bars). For $p = 1/3$ trust is most often exploited in period 9, i.e. the last period of the second game. Afterwards such a degree of exploitation is impossible since players 1 were warned that this might happen. The relative rates of exploited trust in last period are 50% (50%) in game 1, 72% (57%) in game 2, 50% (80%) in game 3, 88% (33%) in game 4, 82% (50%) in game 5, and 78% (100%) in game 6 for the main treatment (control treatment).




Figure 4: Total number of actual trust choices and exploited ones per period $p = 1/3$ (left, pairs with real players 2 only) and $p = 0$ (right)

How do participants react to the possibility of explicit randomization? The possibility that this invites extensive randomization (demand effect) is not confirmed by Figure 5 reporting the numbers of pure strategies by the 45 players 1 and the 30 players 2. Except for two minor exceptions (periods 1 and 6 for players 1) at least $2/3$ of all participants use pure strategies. The overall probability of pure strategies is 81.0% for players 1 and 88.7% for players 2. According to the normative solution both players randomize in two periods when $m > 2$ (in case of $m = 2$ only player 2 randomizes once). Thus the expected probabilities of pure strategies are $2/3 \approx 66.7\%$ for player 1 and $19/30 \approx 63.3\%$ for player 2. Figure 6 is the $p = 0$ analogue of Figure 5. The share of pure strategies is with 84.6% for players 1 and 89.6% for players 2 very similar to the main treatment.




Figure 5: Number of pure strategy choices by players 1 (left) and players 2 (right), $p = 1/3$




Figure 6: Number of pure strategy choices by players 1 (left) and players 2 (right), $p = 0$

How many players 2 rewarded player 1 even in the last period of a (repeated) game is illustrated by Table 1, which distinguishes between reward ($0 < y_m \leq 1$) and full reward ($y_m = 1$). Since each participant played 6 games with 6 different partners the frequency is 0 at least and 6 at most. Only 2 of the 30 players 2 did never try to exploit a trusting player 1 in the last period of a game. One of these participants chose full reward in all 30 periods, the other did so except for period 1. All other players 2 yielded to opportunistic incentives by deviating from full reward $y_m = 1$ in the last period of at least three games. The two players 2 who have chosen $y_m = 1$ in (almost) all 30 periods could be viewed as an endogenous increase of p from $1/3$ to $p = \frac{1}{3} + \frac{2}{3} \cdot \frac{2}{30} = \frac{17}{45}$. Players 2 who sometimes reward and sometimes exploit in the last period of a repeated game are either still experimenting or variety seeking.

degree of reward	frequency of reward in the last period						
	0	1	2	3	4	5	6
reward ($0 < y_m \leq 1$)	7	6	6	3	3	0	5
full reward ($y_m = 1$)	11	7	7	3	0	0	2

Table 1: Number of players 2 choosing reward in the last period with the given frequency

How did players 1 react to being exploited? 10 of the 45 players 1 were never exploited before the last period of a game. 10 of the remaining 35 players 1 never trusted again in the same game after being exploited,

whereas 25 did. These 25 clearly violate a basic requirement of reputation equilibria.

player i	earnings				
	min_i	μ_i	μ_i^*	max_i	σ_i
1	156	215.4	222.75	255	24.35
2	222	249.3	249	273	19.52

Table 2: Actual and equilibrium payoffs

According to Table 2 players 2 on average earned $\mu_2 = 249.3$ and thus roughly the amount $\mu_2^* = 249$ predicted by game theory. The variance of earnings is smaller than for players 1, both according to $\sigma_1 > \sigma_2$ as well as according to $max_1 - min_1 = 99 > max_2 - min_2 = 51$. The deviation $|\mu_1 - \mu_1^*| = 7.35$ is also minor. It seems that the payoff losses, due to non-optimal choices, have been nearly compensated by the additional chances of cooperation offered by more trustworthy players 2 than those exogenously induced by $p = 1/3$.

A strength of our experimental design is that we can test the qualitative structure of reputation equilibria also on the individual level. We focus on the main treatment ($p = 1/3$) and experienced play (games 4,5, and 6). The high proportion of pure strategies (see Figures 5 and 6) suggest that most participants are not compromising by mixing. Precisely eight of 45 players 1 and three of 30 players 2 choose a mixed strategy in more than 1/3 of the periods, whereas the others rarely mix.

Eight of the players 1 always choose full trust ($x_t = 1$) as long as they are not exploited. Two of them are never exploited and three of them switch after exploitation to no trust ($x_t = 0$) for the remainder of the game. Another eight players 1 choose full trust unless they have been exploited or the last period of the game is reached. Two players 1 always choose no trust ($x_t = 0$).

A relatively large proportion of players 2, namely 18 of 30 choose full reward ($y_t = 1$) until at least the third to last period in all the three games considered. Seven of these always choose full reward until the second to last period and no reward ($y_t = 0$) in the last period. Another two always choose full reward. One player 2 always chooses no reward ($y_t = 0$). Hence on an individual level, the share of players 2 whose behavior is close to the equilibrium prediction is larger than that of players 1.

5 On reputation formation

Game theoretically, except for $m = 2$, players should start with full trust ($x_1^* = 1$) which should be partially rewarded ($m \leq 3$) or fully rewarded ($m > 3$). Only towards the end, i.e. in the last three periods, trust and reward of trust should sharply decline.

A rudimentary aspect of such behavior is the monotonicity in the sense that neither x_t nor y_t should increase with t . In the main treatment ($p = 1/3$) for player 1 the proportion of monotone choices is initially below 50%, but increases afterwards (see Table 4). For players 2 it is relatively high from the beginning. The share of always monotonically choosing players 1 and 2 is, however, very low (22.2% and 30.0%, respectively). The results for the control treatment ($p = 0$) are similar.

monotonicity	p=1/3		p=0	
	player 1	player 2	player 1	player 2
game 1 (3 periods)	42.2	63.3	61.1	66.7
game 2 (6 periods)	33.3	60.0	38.9	61.1
game 3 (2 periods)	88.9	96.7	83.3	94.4
game 4 (10 periods)	55.6	60.0	61.1	61.1
game 5 (3 periods)	71.1	76.6	88.9	83.3
game 6 (6 periods)	62.2	80.0	72.2	72.2
all games (30 periods)	22.2	30.0	27.8	33.3

Table 3: Percentage of monotone (decreasing) strategies by player types and treatment

To check for experience effects one can compare the 1st and the 5th as well as the 2nd and the 6th repeated game extending over 3 and 6 periods, respectively. In Figure 7 (left) we have confronted the x_t -development of the 1st game (solid squares) with the one of the 5th game (hollow squares) in the upper part and the y_t -development in the lower part. The analogous comparison of the 2nd and 6th game can be found in Figure 7 (right). While average trust- and reward rates fall (almost) monotonically in the 5th and 6th game, this is not the case for the inexperienced participants in the 1st and 2nd game. Learning shifts average behavior towards the normative benchmark. For $m = 6$ the learning effects (players 1 learn to trust less and players 2 to exploit when the end is near) shift behavior into the direction of rational behavior. Average trust is initially nearly constant ($x_t \sim .8$ for $t = 1, 2, 3, 4$)

before it drops to $x_6 \sim .4$. Similarly there is a sharp decline from $y_t \sim .9$ for $t = 1, 2, 3, 4$ to $y_6 < .2$.




Figure 7: Average trust (top) and reward (bottom) rates for games with 3 (left) and 6 (right) periods

Table 3 gives the results of a Wilcoxon-test comparing for each individual subject the decisions in corresponding periods in different games of equal length. The learning from earlier to later games towards the normative benchmark is significant (5%-level) for the first period and both types of players, i.e. trust and reward increase significantly from game 1 to game 5 and from game 2 to game 6. For the last period of the 6 period games the opposite effect is significant for both types of players.

		games 1 and 5			games 2 and 6			
		1	2	3	4	5	6	
player 1	Z	-2.750	-0.679	-1.391				
	sign. p	.006*	.497	.164				
player 2	Z	-3.158	-2.432	-1.862				
	sign. p	.002*	.015*	.063				

		games 1 and 5			games 2 and 6			
		1	2	3	4	5	6	
player 1	Z	-2.301	-2.159	-2.837	-1.426	-1.314	-2.626	
	sign. p	.021*	.031*	.005*	.154	.189	.009*	
player 2	Z	-2.812	-1.869	-1.051	-0.406	-1.647	-2.014	
	sign. p	.005*	.062	.293	.684	.100	.044*	

Table 4: Wilcoxon-test for single periods in games of equal length

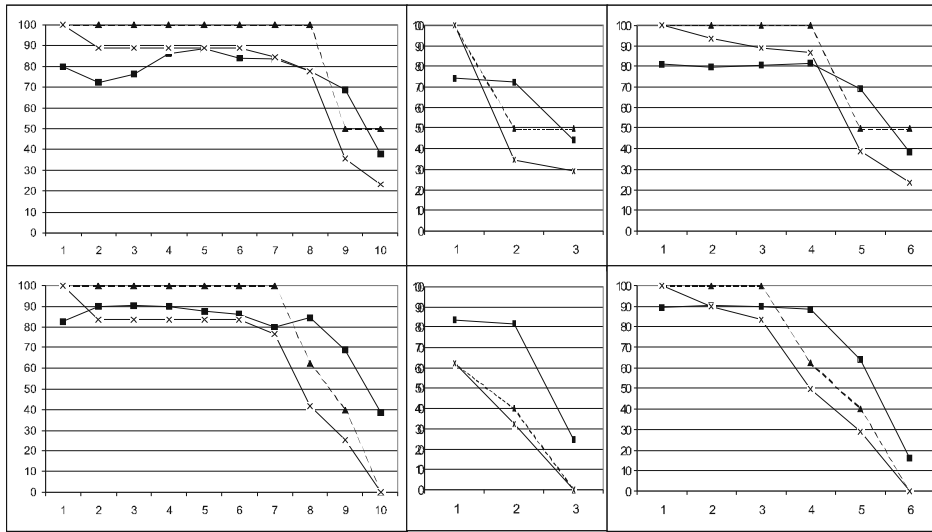


Figure 8: Observed (boxes), equilibrium (triangles), and conditionally optimal behavior (crosses) for games 4, 5, and 6 – player 1 (top) and player 2 (bottom)

When testing reputation formation it seems reasonable to focus on experienced behavior. We therefore concentrate in the following on the 4th, 5th,

and 6th game with $m = 10, 3$, respectively 6 periods. For all these games Figure 8 confronts the observed average with the equilibrium and the conditionally optimal behavior for the given previous path.

In view of the extreme predictions of (conditional) optimality, e.g. in the sense of $x_t = 1 = y_t$ for the early periods t of a game, it is clear that the empirical distribution differs significantly in the early periods as well as in the last period $t = m$ where $y_m = 0$. We therefore restrict the analysis in Table 5 to the periods with non-trivial conditional best replies i.e. last and second to last period for player 1 and second and third to last period for player 2. We conduct the analysis both including (incl.) all players and excluding (excl.) those players whose conditional best reply is $x_t = 0$ or $y_t = 0$, since they can deviate only in one direction.

Game	Period	x_t		y_t	
		incl.	excl.	incl.	excl.
4	8			-4.778 $p = .000$	-4.472 $p = .000$
	9	-4.757 $p = .000$	-4.505 $p = .000$	-3.982 $p = .000$	-3.455 $p = .001$
	10	-2.527 $p = .012$	-1.322 $p = .186$		
5	1			-2.535 $p = .011$	-2.535 $p = .011$
	2	-4.725 $p = .000$	-4.219 $p = .000$	-4.575 $p = .000$	-4.463 $p = .000$
	3	-1.987 $p = .047$	-.366 $p = .714$		
6	4			-4.410 $p = .000$	-4.041 $p = .000$
	5	-3.752 $p = .000$	-3.288 $p = .001$	-3.681 $p = .000$	-3.194 $p = .001$
	6	-1.900 $p = .057$	-.138 $p = .891$		

Table 5: Wilcoxon-test of actual decisions against conditionally best replies

According to Table 5 players 1 are significantly more trusting than predicted in the second to last period. Also, players 2 are significantly more rewarding than predicted in both the second and third to last period. In the

last period players 1 are significantly more trusting than predicted in games 4 and 5. If those players 1 for whom the conditional best reply is $x_m = 0$ are excluded from the analysis the difference between actual decisions and conditional best replies almost completely disappears.

Comparing individual trust and reward rates in the control and the main treatment, i.e. for $p = 0$ and $p = 1/3$, for the single periods, we find that reward rates never differ significantly. The trust rates are significantly higher (Mann–Whitney–U test, 5%–level) for $p = 1/3$ than for $p = 0$ only in periods 2, 3, 8, and 24, if we include all players 1. If we include only those players 1 with a real partner the same holds for periods 2, 3, 10, and 24. Appendix B compares the average behavioral trends for $p = 1/3$ (black squares) with those for $p = 0$ (grey triangles) in all six successive games. Except for game 1 where $p = 0$ induces less trust and more (voluntary) rewarding the differences appear as minor. Thus compared to the theoretical prediction the effect of robot strategies is surprisingly small.

To test the qualitative structure of the basic feature of rational play, namely initial constant trust and reward followed by monotone decreasing, we assume a simple linear structure, i.e.

$$x_t = \begin{cases} \bar{x} & \text{for } t \leq t^* < m \\ \bar{x} - \alpha(t - t^*) & \text{for } t^* < t \leq m \end{cases}$$

where $\bar{x} \geq 0, 0 \leq \alpha < \frac{\bar{x}}{m - t^*}$ and

$$y_t = \begin{cases} \bar{y} & \text{for } t \leq \tilde{t} < m \\ \bar{y} - \beta(t - \tilde{t}) & \text{for } \tilde{t} < t \leq m \end{cases}$$

where $\bar{y} \geq 0, 0 \leq \beta < \frac{\bar{y}}{m - \tilde{t}}$

We first estimate for given t^* and \tilde{t} the coefficients (\bar{x}, α) and (\bar{y}, β) , respectively and then search for the values t^* and \tilde{t} for which the explained variance of the individual decisions is maximal. The results (\bar{x}, α, t^*) and $(\bar{y}, \beta, \tilde{t})$ are displayed in Table 6.

	game	1 (3 periods)			2 (6 periods)			3 (2 periods)		
$p = 0$	\bar{x}, α, t^*	45.19	12.75	1	73.92	13.37	3	22.22	-6.94	1
$p = 1/3$	\bar{x}, α, t^*	55.20	-2.37	1	64.48	-1.42	2	43.16	8.78	1
$p = 0$	$\bar{y}, \beta, \tilde{t}$	72.78	14.44	2	80.99	11.46	3	79.44	47.78	1
$p = 1/3$	$\bar{y}, \beta, \tilde{t}$	51.30	6.30	2	78.45	46.61	5	73.67	45.60	1
	game	4 (10 periods)			5 (3 periods)			6 (6 periods)		
$p = 0$	\bar{x}, α, t^*	76.56	13.76	7	84.91	39.44	1	79.64	30.95	4
$p = 1/3$	\bar{x}, α, t^*	81.46	19.93	8	73.13	28.96	2	81.74	19.89	4
$p = 0$	$\bar{y}, \beta, \tilde{t}$	84.07	45.19	9	86.48	25.00	1	91.16	29.92	4
$p = 1/3$	$\bar{y}, \beta, \tilde{t}$	86.72	22.91	8	82.67	58.00	2	90.56	35.14	4

Table 6: Optimal regression parameters t^* and \tilde{t}

Experienced players (games 4–6) conform qualitatively to the equilibrium prediction in the sense that their decisions are best approximated by a constant trust or reward rate until the second or third to last period followed by a substantial decrease. For players 1 the decrease of trust rates starts earlier (game 4 and 5) or is larger (game 6) for $p = 0$ than for $p = 1/3$.

6 Conclusions

Our experiment shows that, although individual decisions are quite far from the equilibrium prediction, on an aggregate level the data are quite well explained by reputation formation. In particular individual mixed strategies are far from equilibrium mixing. Players on average mix less often than predicted and rather in periods when they should not. This is no surprise given the complicated nature of the equilibria. Also, often choices are not rational in view of the given information. On the other hand, important features of the equilibrium prediction are reflected by the aggregate data. Given some experience, the aggregate trust and reward rates are monotonically decreasing, in particular there is a clear endgame effect. Furthermore, the length of the game does not have a significant influence on the play in the last three periods, thus there is no history dependence, except for that players 1 react to being exploited. Also for the control treatment aggregate decisions conform qualitatively to the equilibrium prediction (for a positive endogenous probability of rewarding, i.e. of intrinsically fair players 2). Hence by introducing the fair robot strategies we have not altered the nature of the game

but have just increased the subjective probability of meeting a fair player. The combination of the assumption of some intrinsic fairness and reputation formation thus seems to explain also the decisions in games even without fair robot strategies. The small difference between the main treatment and the control treatment (confirming the results of Andreoni and Miller, 1993) indicates that introducing fair robot players either leads to some crowding out of intrinsic fairness or to decreasing trust in it.

Overall our results are well in line with those of Andreoni and Miller (1993), Neral and Ochs (1992), Camerer and Weigelt (1988), and McKelvey and Palfrey (1992). They also find support for reputation formation, although Neral and Ochs observe that changes in the parameters induce changes in behavior not consistent with the equilibrium prediction. As we do they all find a clear endgame effect and clear evidence of learning in the direction of the equilibrium prediction, where learning appears to be relatively fast in our study. Also, subjects seem to have some homemade prior of altruism or fairness which is empirically justified. One of the aspects where our experiment differs is that we allow for mixed strategies on an individual level and find that the latter does not conform to the equilibrium prediction. Thus the conclusion that the reputation equilibrium yields a bad prediction for individual behavior is strengthened.

References

- [1] Albers, W. and Vogt, B. (1997): The Selection of Mixed Strategies in 2x2 Bimatrix Games, University of Bielefeld, Institute of Mathematical Economics, Working Paper No. 286.
- [2] Andreoni, J. and Miller, J.H. (1993): Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence, *The Economic Journal*, **103**, 570–585.
- [3] Camerer, C. and Weigelt, K. (1988): Experimental Tests of the Sequential Equilibrium Reputation Model, *Econometrica*, **56**, 1–36.
- [4] Engelmann, D. (1999): Perfect Bayesian Equilibria in the Finitely Repeated Trust Game with Imperfect Information, mimeo, Humboldt-Universität zu Berlin.
- [5] Fudenberg, D. and Tirole, J. (1991): *Game Theory*, MIT Press.

- [6] Güth, W., Ockenfels, P., and Wendel, M. (1993): Efficiency by Trust in Fairness? Multiperiod Ultimatum Bargaining Experiments with an Increasing Cake, *International Journal of Game Theory*, **22**, 51–73.
- [7] Güth, W., Ockenfels, P., and Wendel, M. (1997): Cooperation Based on Trust: An Experimental Investigation, *Journal of Economic Psychology*, **18**, 15–43.
- [8] Kreps, D., Milgrom, P., Roberts, J., and Wilson, R. (1982): Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma, *Journal of Economic Theory*, **27**, 245–252.
- [9] Kreps, D. and Wilson, R. (1982): Sequential Equilibria, *Econometrica*, **50**, 863–894
- [10] McKelvey, R.D. and Palfrey, T.R. (1992): An Experimental Study of the Centipede Game, *Econometrica*, **60**, 803–836.
- [11] Neral, J. and Ochs, J. (1992): The Sequential Equilibrium Theory of Reputation Building: A Further Test, *Econometrica*, **60**, 1151–1169
- [12] Ochs (1995): Games with Unique, Mixed Strategy Equilibria: An Experimental Study, *Games and Economic Behavior*, **10**, 202–217.
- [13] Selten (1975): Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, *International Journal of Game Theory*, **4**, 25–55.

A Appendix (Translated instructions)

[*The part in parentheses is missing in the instructions for the control treatment.*]

Welcome to our experiment. Please read the following instructions carefully. If there is anything that you do not understand either now or during the experiment, please do not ask publicly but raise your hand. We will then answer your question in private. Please stop communicating with other participants now for the remainder of the experiment. Otherwise you will not receive a payment.

The experiment consists of six rounds. Each round consists of several periods. In each period the same decisions are to be made. At the beginning of each round you will be informed about the number of periods in this round. In all periods of one round you will meet the same partner. However,

in different rounds you will always meet different partners. Hence you will never interact with someone with whom you have interacted in an earlier round. No participant will get to know any decision made by his current partner in previous rounds. Furthermore, you will never meet someone, who has met someone who previously had interacted with one of your previous partners etc. Thus whatever you do in one round can not influence later interactions, not even indirectly.

There are two different roles, participant X and participant Y. The roles will be assigned randomly at the beginning and will be constant for the whole experiment, i.e. you will either be a participant X in all rounds or a participant Y. You will be informed about your role before the start of the first round.

Participant X has to decide in each period between actions T and N. Participant Y chooses simultaneously between alternatives R and E. The payoffs for both participants are shown in the following table. Each cell represents a constellation of decisions of X and Y. The entries are the payments for X and Y (in ECU- Experimental Currency Unit). At the end of the experiment the sum of your earnings in ECU will be converted and paid in DM. 1 ECU corresponds to 0.10 DM.

		Participant X chooses			
		N		T	
		Payment X	Payment Y	Payment X	Payment Y
Participant Y chooses	R	6	6	9	9
	E	6	6	0	12

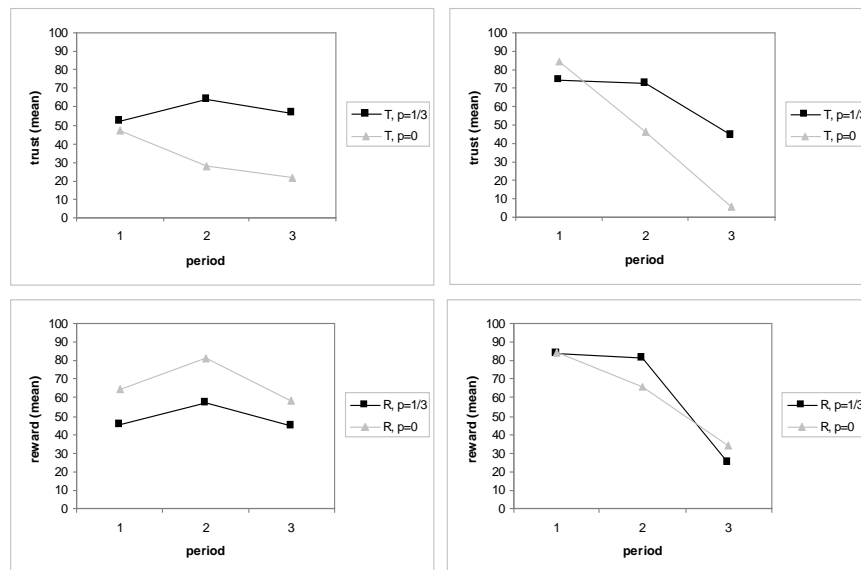
Verbally speaking this means: If participant X chooses alternative N, then the decision of the other participant is irrelevant and both receive a payment of 6 ECU. If participant X chooses alternative T and participant Y chooses alternative R then both receive a payment of 9 ECU. If participant X chooses alternative T and participant Y chooses alternative E then participant Y receives 12 ECU, but participant X nothing. Since decisions are made simultaneously, a participant will only learn the decision of the other participant after making the own decision. Participant X will only get to know the decision of participant Y if he has chosen T.

As noted before, each participant will be assigned a new partner in each round but keeps his role. [The following is to consider. For each participant X, at the beginning of each round a random draw determines with probability 1/3 that independent of the decisions of other participants a choice

of T will be followed by R. With the remaining probability $2/3$ he will be assigned a Y-partner who is completely free to choose between E and R. He will not be informed about the result of this random draw. However, if a participant X receives the reply E in a period after a choice of T, he is sure that he is confronted by a freely deciding Y-partner. The random draw is independent for each round. There is no such mechanism for participant Y.]





You can delegate the decisions that you have to make to a random device. For example as participant X, instead of deciding definitely for T or N, you can choose the probabilities for selecting T or N. In order to do this you fill in the probabilities in the corresponding fields on the screen. These have to add to 100%, of course. The computer will then select one of the possible decisions according to the probabilities you have chosen. The other participant will only be informed about the decision selected by the computer but not about the chosen probabilities. If you definitely want to decide for T, choose 100% for T and 0% for N and vice versa 0% for T and 100% for N in case of a definite decision for N. Participant Y can delegate his decision between R and E to the random device in an analogous way.

B Appendix (All six successive games)







Game 1

Game 5



Game 2

Game 6



Game 3

Game 4