



Learning (not) to yield: An experimental study of evolving ultimatum game behavior



Judith Avrahami^a, Werner Güth^b, Ralph Hertwig^c, Yaakov Kareev^a, Hironori Otsubo^{d,*}

^a The Hebrew University of Jerusalem, Center for the Study of Rationality, Givat Ram, Jerusalem, 91904, Israel

^b Max Planck Institute of Economics, Strategic Interaction Group, Kahlaische Strasse 10, 07745 Jena, Germany

^c Max Planck Institute for Human Development, Center for Adaptive Rationality, Lentzeallee 94, 14195 Berlin, Germany

^d Soka University, Faculty of Economics, 1-236 Tangi-machi, Hachioji, Tokyo 192-8577, Japan

ARTICLE INFO

Article history:

Received 24 August 2012

Received in revised form 23 August 2013

Accepted 30 August 2013

Keywords:

Ultimatum bargaining game

Reputation

Regret

Learning

Experiment

ABSTRACT

Whether behavior converges toward rational play or fair play in repeated ultimatum games, depends on which player yields first. If responders conceded first by accepting low offers, proposers, would not need to learn to offer more. Play would thus converge toward unequal sharing. If proposers, learnt fast that low offers are doomed to be rejected and adjusted their offers accordingly, pressure, would be lifted from responders to learn to accept such offers. Play would thus converge toward equal, sharing. Here, we tested the hypothesis that it is regret—both material and strategic—which determines, how players adapt their behavior. We conducted a repeated ultimatum game experiment with, randomly changing strangers. One treatment offers players only feedback about the outcome of their, play. Another treatment offers additional information about the median outcomes in the population. We find that regret is a good predictor of the dynamics of play, in particular of proposer behavior., Except for a very short endgame phase, in which more tolerance of less equitable sharing surfaced, behavior converges toward equal sharing. Population information hardly speeds up this convergence.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

Over more than thirty years, a 2-person bargaining game (Güth et al., 1982) has piqued the curiosity of experimenters in economics and beyond. Exasperated by the plethora of ultimatum game investigations (see the recent review of Güth and Kocher, 2013), Camerer (2003) proposed a ban on further studies. Although one can easily relate to his sentiment, one should not overlook that often the ultimatum game is only employed as an experimental tool to investigate new issues such as the effect of cheap talk messages or different degrees of anonymity. Recently, however, our interest in the game per se has been reawakened when noticing that only relatively little is known about the game's learning and population dynamics. Specifically, we will turn to the dynamics that unfold when players make repeated decisions in the ultimatum game with randomly changing opponents, and when they not only learn about their own outcome in the previous round but also find out how the population, on average, adapts to previous results (path dependence).

Section 2 briefly reviews the literature on repeated ultimatum game. In Section 3, we discuss how feedback information about the outcome of the previous round may influence future behavior, especially when assuming that participants are guided by post-decisional regret as suggested by the theory of direction learning and the quantitative yet static concept of impulse balancing. Section 4 describes our experimental method. Section 5 reports the main results, focusing on the test of how post-decisional regret shapes behavior before concluding.

2. Repeated ultimatum play

The paradigmatic finding in studies of one-shot implementations of the ultimatum game is fair play: That is, proposers tend to offer relatively equal divisions and responders are inclined to rebuff relatively unequal offers (see Camerer, 2003). Relatively few studies have examined proposer and responder behavior in recurrent ultimatum games in which the same two partners interact repeatedly (see, for example, Cooper and Dutcher, 2011). Under these circumstances, the game is assumed to morph into a reputation game (Kreps et al., 1982). Pursuing the goal of building up an advantageous reputation, players' early moves take on special importance because they signal a player's type. For instance, by being unwavering at the beginning—proposing the smallest possible share and rebuffing any proposal that deviates from equality—proposers and

* Corresponding author. Tel.: +81 42 691 9316.

E-mail addresses: judith.avrahami@mail.huji.ac.il (J. Avrahami), gueth@econ.mpg.de (W. Güth), hertwig@mpib-berlin.mpg.de (R. Hertwig), yaakov.kareev@mail.huji.ac.il (Y. Kareev), ootsubo@soka.ac.jp (H. Otsubo).

responders bet on later rounds in which they hope to reap the fruits of their unbending stance. Thwarting their calculations, such unbending postures can result in outrageous conflict rates (e.g., Slembeck, 1999). Suggestive evidence for reputational concerns can also be found in experimental studies of closely related games, played in a partner design, e.g., of repeatedly played social dilemma games (see Ledyard, 1995; Selten and Stoecker, 1986).

The question is what to expect in recurrent play in stranger design. Fischer et al. (2006) have already observed convergence to fair play rather than opportunistic behavior with random strangers matching. This may indicate that reputation concerns are weaker or non-existent in a random matching protocol (Cooper and Dutcher, 2011). Such expectations are in line with several theoretical analyses that have explored when and why convergence to fair play ought to be expected including Binmore and Samuelson (1995) and Huck and Oechssler's (1999) evolutionary approaches, and Roth and Erev's (1995) reinforcement learning approach.

3. The current study

In our view, however, more evidence is necessary to warrant the prediction that a stranger design will give rise to equal sharing. In addition, no previous studies have scrutinized in a step-by-step fashion the question of how feedback for a decision affects the following one in a repeated ultimatum game with strangers. Henceforth, we focus on effects of two types of feedback: (i) information about one's particular proposer-responder dyad, and (ii) additionally to (i), also information about the population of which one's dyad is a member. Specifically, we investigate two treatments differing only in whether players' receive (or fail to receive) information about how others, faced with the same situation, responded.

In our *D-treatment*, players received information about the decision of their partner in the current round: responders received the offer proposed to them and proposers received the threshold at which their responder would have been willing to accept or decline the offer. Such a threshold provides proposers with extra information that may enable them to fine-tune their behavior. Moreover, it provides us with a dimension—over and above the binary “accept versus reject” response—on which basis we can observe how people's behavior evolves over time.

Imposing monotonic acceptance strategies on responders enables us to provide accessible information concerning a responder's decision for the current round. Let us admit, however, that in doing so, we accept two potential disadvantages. The first is that some participants may prefer non-monotonic strategies and reject not only meager but also over-generous offers (see Güth et al., 2003, 2007; Hennig-Schmidt et al., 2008). Second, the strategy method is assumed to engage a “cold” state of cognition, that is, people are required to imagine how upset they will be when facing an unfair offer rather than when actually facing it. In sum, participants obtained their dyad-specific information, that is, each of the two players learnt about (i) the amount offered by the proposer and the acceptance threshold chosen by the responder, (ii) whether the proposer's offer was accepted or rejected, and (iii) both players' payoffs for the current round.

In our *P-treatment*, players received information about the behavior in the *population*, that is, about all other dyads, in addition to the outcome in one's own dyad. Population behavior can be represented in different ways. For instance, players could be provided with distribution information about how frequently any possible division occurred, and the likelihood with which it was accepted and rejected, respectively. Such complex distributive information, however, may overtax participants' comprehension, and thus compromise its value. For this reason, we presented information about

population behavior in terms of simple aggregated properties of the distribution, namely, the median offer and the median acceptance threshold. In sum, in the *P-treatment*, in every period each dyad (a proposer and a responder) learnt not only their own dyad's information, as in the *D-treatment*, but also the median offer and the median acceptance threshold of the other dyads.

What kind of behavior will evolve in a repeated ultimatum game involving strangers who receive such detailed feedback following each period? Moreover, how will population feedback shape behavior, relative to a condition in which a player's feedback refers merely to one's dyad? We approach these questions using *learning direction theory* (Selten and Buchta, 1999; Selten and Stoecker, 1986; Ockenfels and Selten, 2005)—a theory embodying ex-post rationality particularly suited to repeated decision tasks. It applies to situations in which feedback after each period permits drawing retrospective inferences about what behavior would have yielded better outcomes. Egged on by these postdictions, so the theory predicts, the decision maker tends to move future decisions into the direction suggested by the comparison of actual and counterfactual outcomes. The motivational force behind such behavioral adaptations can be interpreted to represent (i) feelings of regret that are experienced in the wake of past decisions and that have an effect on future decisions (Selten and Buchta, 1999) and (ii) feelings of anticipated regret (Ritov, 1996) emerging when deliberating the possible outcomes of the next move. Whereas the first interpretation can be modeled in terms of a quantitative specification of directional learning, the second refers to the forward-looking deliberations as specified by impulse balance equilibria. It should be noted, however, that because players knew that they would find out whether they have cause for regret, any choice could be interpreted either as response to regret over the previous choice or as an act to preempt future, anticipated regret. It is therefore impossible to distinguish between the two in such sequences of repeated decisions. Henceforth, when we speak of the proposer and responder's reaction to regret, we refer to both interpretations.

In a repeated ultimatum game with the feedback we provide, regret for a decision can arise from two sources: experienced and hypothetical. Experienced regret can be triggered by material loss in the recent past or by strategic (reputational) considerations dealing with future material losses as a consequence of past behavior. We refer to the two types of regret as *material regret* and *reputational regret*, respectively. A proposer's material regret may manifest in the following cognition: “I could have earned more, had I just made a higher offer.” A responder, in contrast, may feel that a bird in the hand is worth two in the bush: “Had I just accepted the quite low offer, I would be better off.” Looking beyond the recent past, players may also regret that their past behavior may have revealed too much about their breakpoint. For instance, the proposer may feel remorse about having revealed the willingness to offer more equitable shares: “Now, with responders knowing that there is the willingness to offer fair amounts, they will not accept less.” A responder, in turn, may feel sorry for having made public the willingness to be content with less: “Now, with proposers knowing that there is the willingness to accept less, they will offer less.” Such manifestations of reputational regret could be stronger in the *P-treatment* than in the *D-treatment* where players may be concerned about what their future counterparts read off from the population feedback.

The notion of reputational regret suggests that, like in a repeated ultimatum game involving a partners design, our strangers design may give rise to the same strategic considerations in which players regret having revealed their breakpoint too early. If so, we may observe behavioral escalation (an unusually high conflict ratio), of the kind reported for the ultimatum game involving a partners design (Slembeck, 1999), with proposers who aim to demonstrate that they offer “little” and responder, in turn, who aim to show that

they are not willing to accept “little”. This would be particularly true in the *P-treatment*.

Of course, as a consequence of random rematching in each round, such reputational concerns hinge on the specific role that is being assigned. If people are cognizant of this dependency, a strong incentive could emerge to free ride on the reputation established by others in the same role. When others, in turn, experience such free riding, their role-based reputational concerns may be put in question. To put it differently: The role of reputational regret will depend on whether proposer and responder participants sense or develop some strong enough group or role solidarity that averts or sufficiently attenuates free-riding on the reputation that has been established by others sharing the same role.

Because reputational regret is expected to be weaker in our *D-treatment*, which does not offer population feedback but merely feedback about the play of one’s own dyad in the last period, behavior may more easily converge toward equal division than in the *P-treatment*. Of course, one may also speculate that the opposite will occur: equal sharing will evolve more swiftly in the *P-treatment* than in the *D-treatment* if, for instance, the feedback about median offers and median acceptance threshold reveals a strong proclivity toward equitable sharing.

In the *P-treatment*, offering feedback about potential partners, yet another type of regret may arise, namely hypothetical regret. It captures reactions like: “Had I played against a prototypical other my offer (or threshold) would have turned out to be too low (or too high)”. Knowing that through random matching one is likely to face any of the others as future partners one may change one’s behavior in response to such hypothetical regret as well.

To summarize, we are concerned with the questions of if and how the repeated play of the ultimatum game is shaped by regret, and how regret’s impact depends on the presence versus absence of population information. In particular, which treatment would result in faster convergence toward equal sharing. Relatedly, we examine whether population information decreases or increases the conflict ratio. As suggested earlier, we expect regret—material, reputational and hypothetical—to matter and to be influenced by the different feedback, as varied by our experimental treatments. Do we expect strong effects? It depends. If fairness is due to concerns for strategic reputation, there should be a drastic deviation from equal sharing toward the end—but only, of course, if equitable sharing was the rule before. With regard to treatment, which involved dyad feedback versus both dyad and population feedback, we are less certain that the effect will be both statistically and economically significant. In sum, even after three decades of investigations of the ultimatum game, it can still generate novel and non-trivial questions.

4. Experimental method

One hundred and twenty-eight students from various fields and enrolled at the Friedrich-Schiller University of Jena were recruited via the ORSEE software (Greiner, 2004). They were divided into four groups of 32 participants, two groups in the *D-* and *P-treatment*, respectively. Each session involved only one group, and none of the participants was admitted to more than one session, i.e., a between-subject design was employed. They were all *experienced* insofar as they had participated in ultimatum bargaining game experiments at least once in the past.¹

The experiment was conducted in the laboratory of the Max Planck Institute of Economics, which contains a network of 32 PCs.

Table 1
Mean proposals and requests, by treatment.

		Role	
		Proposer	Responder
Treatment	D-Treatment	.52 (<i>n</i> = 19)	.28 (<i>n</i> = 22)
	P-Treatment	.53 (<i>n</i> = 18)	.27 (<i>n</i> = 21)

It was programmed and conducted with the software *z-Tree* (Fischbacher, 2007). Upon arrival at the laboratory, participants were seated in individual cubicles separated from one another by partitions. Any form of communication between participants was prohibited. Once all participants were seated, they began to read instructions silently at their own pace. Next, an experimenter read the instructions (see Appendix A) aloud so that all information became common knowledge. The experimenter answered questions individually.

At the beginning of each session, four groups of eight participants each were randomly formed, and each person was randomly assigned to one of two different roles, proposer or responder. Each group thus consisted of four proposers and four responders, respectively. Participants retained their role throughout. The instructions emphasized that no interaction between groups was possible and that roles remained the same.

A session lasted about 90 min, including 20 min of instruction and payment. Participants played 100 identical rounds of the simultaneous-move ultimatum bargaining game (as described in Section 2). Each of the 100 rounds was structured as follows. At the beginning of a round, a proposer and a responder in the same group randomly formed a dyad. Once a round began, a decision screen appeared on the computer display on which they were asked to submit their decisions: The proposer stated how much of the “pie” of 10 points s/he is willing to hand to the responder; the responder stated the lowest amount s/he would agree to receive. After all 32 participants per session submitted their decisions, a result screen presented each player with the following information about her dyad: (a) the amount offered by the proposer and the responder’s acceptance threshold, (b) whether the proposer’s offer was accepted or rejected, and (c) her payoff and the payoff of the other player in her dyad for the current round. In the *P-treatment*, the result screen presented the following additional information about the other three dyads per group: (d) the median offer of the other three dyads, (e) the median acceptance threshold of the other three dyads, and (f) the number of other dyads that achieved acceptance because offers had passed the responders’ acceptance thresholds.

At the end of the session, a summary screen displayed the total points players had accumulated and the corresponding cash earnings in Euros. The total points were converted at the rate of €1 per 25 points in both treatments. The average individual payoff was €21.89 in the *P-treatment* and €21.84 in the *D-treatment*, respectively, including a €2.5 show-up bonus.

5. Results

Before turning to the dynamics of play, namely, how the decisions in one period depend on the regret in the previous one, let us report some summary statistics. A considerable number of players either offered or requested a fair share throughout the game (‘Equity Insistent’ players). Their numbers, in the different roles (27 and 21 for proposers and responders, respectively) and in the different treatments (23 and 25 for the *D-* and *P-treatments*, respectively) did not differ significantly. Mean offers and requests, separately for the two treatments, are presented in Table 1.

Averaged across rounds, groups and treatments, proposers offered 4.97 points of the 10-point pie, and responders demanded,

¹ We initially planned to recruit participants who were new to the ultimatum bargaining game experiments. However, we learnt it was impossible to recruit a sufficiently large number of such novices from our subject pool.

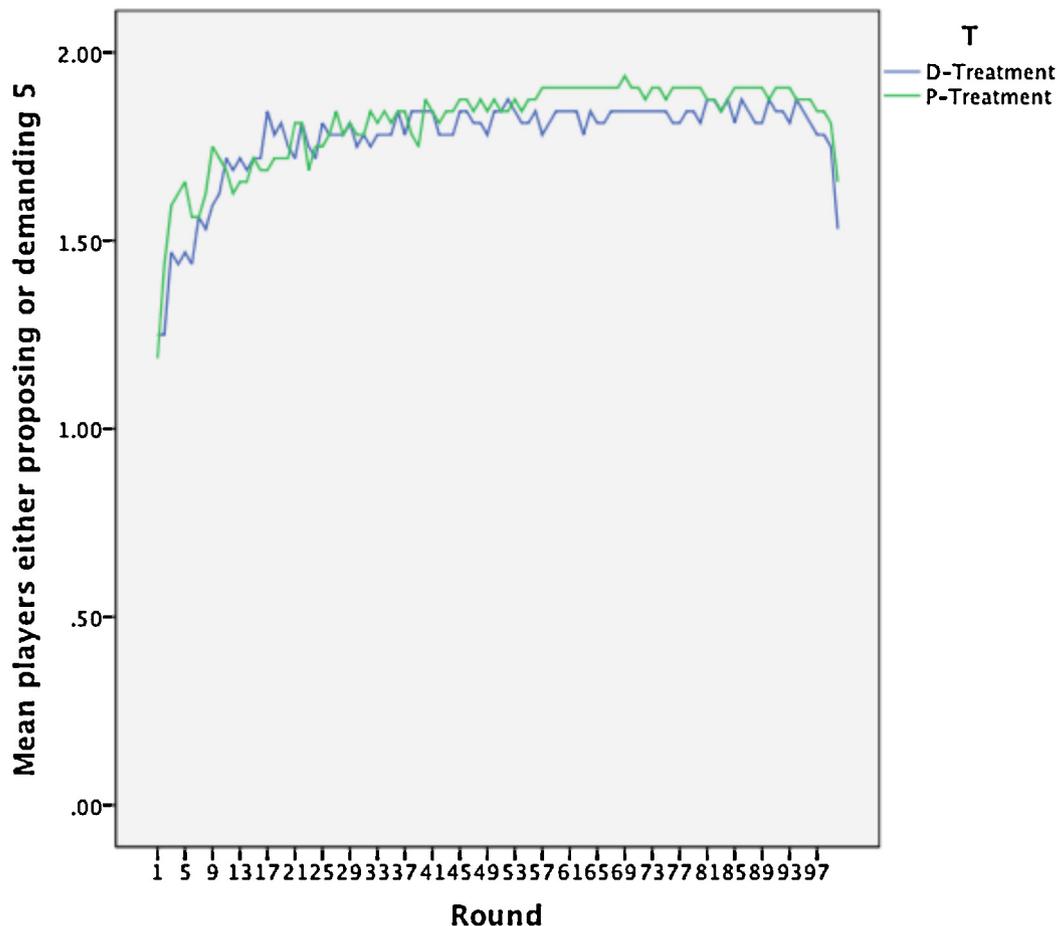


Fig. 1. Mean number of players in a dyad choosing the equitable division (i.e., proposers offering five points or responders demanding at least five). Note that, on average, the number of players in a dyad choosing the equitable division is never less than one.

on average, 4.79 points. The difference between these values is significant ($F_{(1,28)} = 6.75, p = .015$, partial $\eta^2 = .194$). Relative to the D-treatment, the proposers in the P-treatment offered slightly more and the responders demanded slightly more. However, neither the overall difference between treatments nor the interaction between treatment and role were significant ($F_{(1,28)} = 2.18, p = .151$, partial $\eta^2 = .072$ for the overall difference, $F < 1$ for the interaction). Thus, the additional information conveyed in the P-treatment did not escalate conflict. Relatedly, hypothetical regret (i.e., regret that could only be felt in the P-treatment by comparing one's decision to the median value of other members in one's matching group) did not trigger significantly different aggregate behavior.²

How does players' behavior change with experience? Fig. 1 describes the evolution of equitable offers and demands, namely, the mean number of players in a dyad who offered or demanded five points, over rounds, separately for the two treatments. The literature on finitely repeated games (see, for instance, Selten and Stoecker, 1986) would suggest "cooperation" until an endgame effect takes over. In the ultimatum game, cooperation in the sense of efficient behavior, however, only excludes conflict. Actually, the multiplicity of equilibria of the ultimatum game permits for

finite-horizon Folk Theorems (Benoit and Krishna, 1985). Although the argument is thus not a very stringent one, we interpret "cooperation" as equal sharing up to the last round or two (at which proposers may be tempted to exploit their ultimatum power, and responders, possibly realizing this, change their minimal requests accordingly).

Fig. 1 shows that "cooperation" in the sense of equal sharing until the endgame effect takes over is initially not a ubiquitous behavior. Equal sharing (in the sense of proposers offering 5 and responders accepting only offers of at least 5) is at first at about 60%, but then rather quickly evolves, with the learning, however, lasting during the first half of the total of 100 rounds. As Fig. 1 shows, the proportion of players who opted for the equitable division is somewhat higher in the P-treatment than in the D-treatment.

Another obvious observation is that the proportion of equitable divisions rises with experience, but drastically declines toward the end of the game. Did this decline occur because of proposers' increasing greed toward the end or because responders lessened their demands? Fig. 2 plots the number of points offered and points demanded, separately for the two roles and the two treatments. It was the responders who relaxed their demands (most pronounced in the D-treatment): Apparently, once their reputation ceased to be of consequence, they were ready to make do with whatever was offered.

Next, we turn to the question of how regret in one period has an effect on the decision in the succeeding period. To that end, we calculated players' regret and used it to predict the change in the offer or the demand. Regret was defined as follows: If a transaction failed, proposers would feel regret for having offered too little (i.e.,

² The means for only the players who were not 'Equity Insistent' are quite similar: D-treatment: 4.91 and 4.59 for proposer and responder, respectively; P-Treatment: 4.97 and 4.81 for proposer and responder, respectively. The difference between proposer and responder was significant ($F_{(1,27)} = 8.69, p = .007$, partial $\eta^2 = .244$) but the difference between treatments was not ($F_{(1,27)} = 2.85, p = .103$, partial $\eta^2 = .095$) and there was no interaction between treatment and role ($F < 1$).

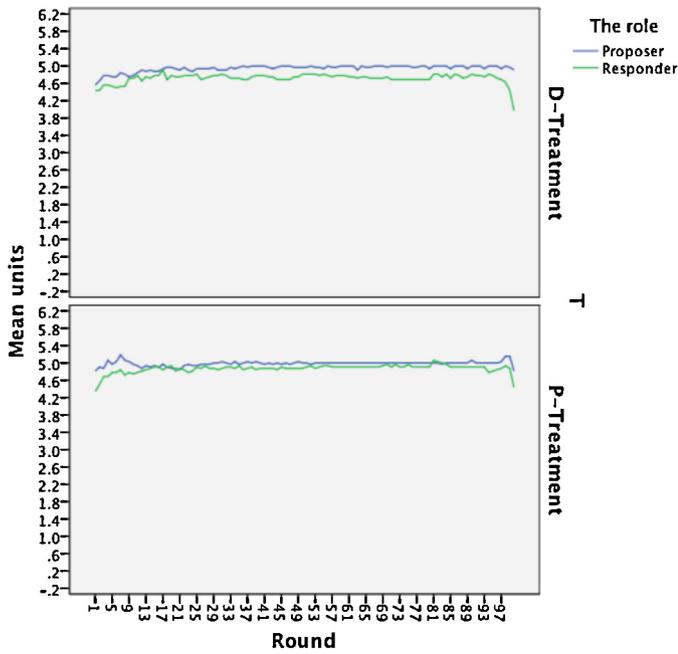


Fig. 2. The average points offered (proposers) and demanded (responders), separately for the P- and D-treatments.

their offer was below the responder’s threshold). Such regret would create an impulse to propose more in the future. The strength of this impulse equals what the proposers could have earned, relative to going out empty handed. Had they proposed exactly as much as the responder’s threshold demands, the proposer could have earned the “pie” minus the threshold.

The responders, too, would have cause for regret if the transaction failed: Had they demanded less, the transaction would have come through and the responder would have secured the proposer’s offer. The size of the responder’s regret equals the size of the offer. The regret of both proposer and responder may be alleviated by reputational satisfaction for having sent a signal to the other: The proposers telegraphed to the responders they should better not expect that much and the responders let the proposers know they should be more generous next time. Obviously, if the proposal equals the responder’s threshold no regret is expected.

Finally, let us turn to successful transactions and the possible regret they may evoke. If the proposers’ offer exceeded the responders’ threshold, they would regret having been too generous. Such regret would create an impulse to propose less in the future. The size of such regret equals the difference between the actual offer and the responders’ threshold. The proposers’ regret is material and reflects the sorrow for forgone payoffs. This material regret may be further aggravated by reputational regret for having disclosed the willingness to be generous. The responders’ regret in the case of a successful transaction in which proposal and acceptance threshold differed is purely reputational. Although having demanded more would not have boosted the gain, it would have avoided signaling willingness to make do with less than was offered.

Taken together, we see then, that in the case of failed transactions reputational regret alleviates the material regret of both players. In contrast, in the case of successful transactions in which proposal and responder’s threshold differ, reputational regret aggravates the material regret of the proposer; furthermore, reputational regret is the only type of regret that the responder should have. These definitions of regret imply that the regret on the part of proposers is more varied, relative to that by responders. To appreciate this, consider the case of a failed transaction. Because thresholds are practically never higher than 5, the material loss for offering too

Table 2

The average correlation between regret in a round and the change in the number of points proposed (by the proposer) or demanded (by the responder), by treatment.

		Role	
		Proposer	Responder
Treatment	D-Treatment	4.94	4.72
	P-Treatment	4.99	4.87

little would be at least 10–5 = 5, whereas the responder’s material loss for demanding too much is practically always less than five. In the case of a successful transaction with a discrepancy between offer and threshold such that the offer exceeds the threshold, the proposer experiences a combination of material and reputational regret whereas the responder experiences reputational regret only. Therefore, we expect regret to better predict the behavior of the proposer than that of the responder.

Having thus defined the impulses evoked by material and/or reputational regret, we tested how well a decision in a given period can be predicted by the regret that was likely experienced in the previous period. We calculated, for every player, the correlation between the change in the number of points (in the offers and demands) chosen from a period to its next and the size (and direction) of the regret in the previous period. Before reporting the results of the correlation a caveat is in order (Table 2). As already discussed above, about a third of our participants chose the same equitable value of five throughout the 100 periods. These participants were obviously not reacting to the previous outcome or, more specifically, to regret. Therefore, the mean values of the correlation analyses to be reported below apply only to those participants whose decisions varied, 80 out of the total of 128 points (offered or demanded) between that round and the next. Regret can indeed better explain changes in proposer than in responder behavior in both treatments. First, an analysis of variance on the Fisher’s Z transformed correlations revealed a strong effect of regret ($F_{\text{intercept}(1,75)} = 135.81, p < .001, \text{partial } \eta^2 = .644$). Second, consistent with our expectation, the effect of regret was stronger for the proposer than for the responder ($F_{(1,75)} = 13.69, p < .001, \text{partial } \eta^2 = .154$). Third, the effect of treatment and the interaction between role and treatment were not significant (both F ’s < 1).

In the P-treatment players received information not only about their own dyad but also about the median offers and demands in their group. They may thus have also sensed regret fueled by the behavior in the population, which we term ‘hypothetical’ regret (see above). The correlations between population-associated regret and changes in behavior were .51 for the proposer and .41 for the responder. In order to compare these effects to those fueled by the outcome of one’s dyad, we conducted a repeated measure analysis of variance on the Fisher’s Z transformed correlations with type of regret (experienced or counterfactual) as a within-subject factor and role as a between-subjects factor. The analysis revealed that here too the effect of regret was stronger for the proposer than for the responder ($F_{(1,34)} = 5.31, p = .027, \text{partial } \eta^2 = .135$). Neither the type of regret nor the interaction between role and type of regret proved significant (both F ’s < 1). It deserves to be emphasized, however, that counterfactual material regret—possibly suggesting how to best respond to the matching group as a whole—was at least as good a predictor of change in behavior as was the regret experienced.

Finally, we analyzed how the conflict ratio evolved in the two treatments. Because the only inefficiency in the ultimatum game is conflict, this analysis will also speak to the efficiency dynamics of both treatments. In order to avoid this problem we did not use the actual frequencies of conflict. Instead, we simulated for each responder how many of the four offers by the four proposers in the

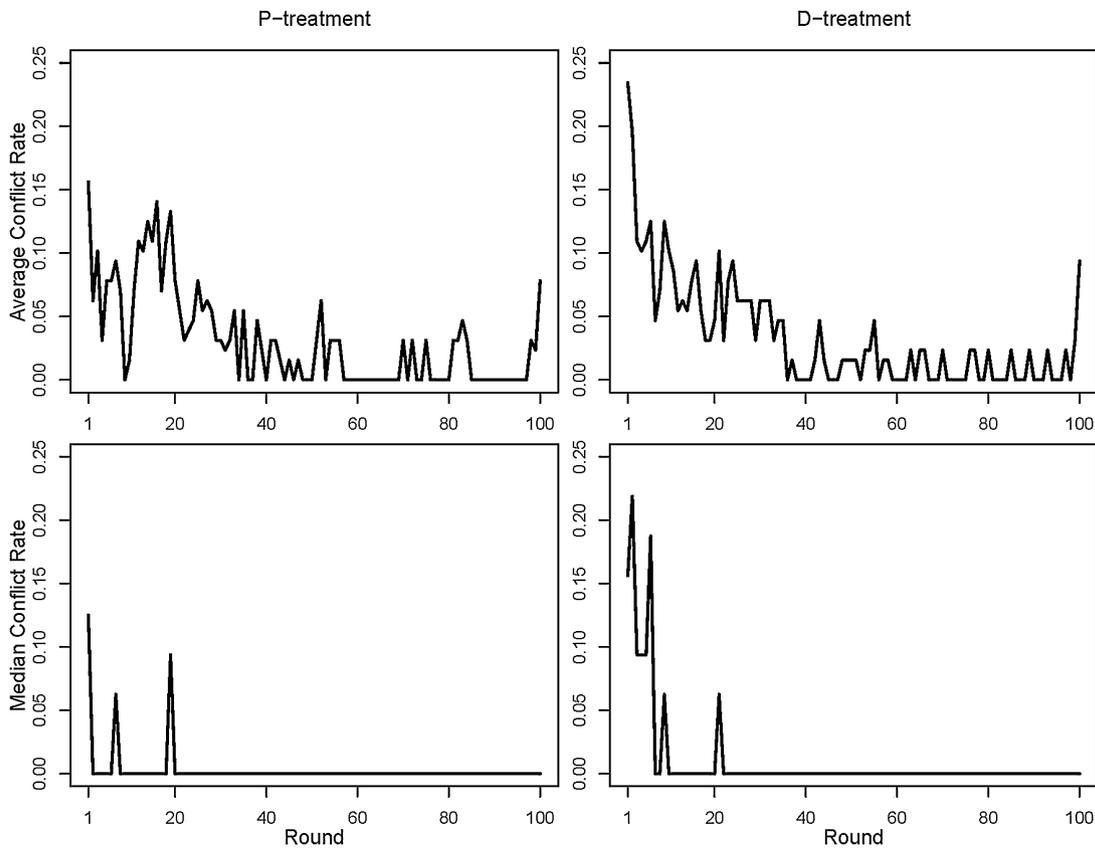


Fig. 3. Average and median conflict rates, separately for the P- and D-treatments.

same matching group (s)he would have rejected. The conflict rate is computed as follows:

$$CR_g(t) = \frac{1}{16} \sum_{i=1}^4 \sum_{j=1}^4 I\{p_{ig}(t) < r_{jg}(t)\}$$

where $p_{ig}(t)$ and $r_{jg}(t)$ denote the proposed points by proposer i and the demanded points by responder j of matching group g in round t , respectively, and where

$$I\{p_{ig}(t) < r_{jg}(t)\} = \begin{cases} 1 & \text{if } p_{ig}(t) < r_{jg}(t), \\ 0 & \text{otherwise.} \end{cases}$$

Fig. 3 displays the average and median values of conflict rates of both treatments. The initial conflict rates, based on group averages, are very similar in both treatments. Thus, population feedback does not appear to discourage conflict more than personal experience alone. The clearest evidence for strategic reputation concerns (or lack thereof) is the drastic decline of equal sharing toward the end. When there is “no future”, there is also little hope for equality concerns and peaceful voluntary cooperation evaporates. The treatment effect is qualitatively as predicted but relatively minor. At least in part, this may be due to the fact that after many rounds even limited information feedback about own past plays provides richer experiences, comparable to population feedback. Another reason is, of course, that even if additional information feedback on others’ past plays matters, one would not expect a large effect size.

6. Conclusions

In a repeated ultimatum bargaining experiment we studied whether adaptation to the outcome of the previous round can be explained in terms of directional learning theory (Selten

and Stoecker, 1986) and the static concept of impulse balancing (Ockenfels and Selten, 2005). To this end, we distinguished two types of regret—material and reputational regret for proposers and responders, respectively—and discussed how they may shape future behavior. Reputational regret would in all likelihood have been stronger in a partners design. But because of the small population size—there were only four possible counterparts—reputational regret could even be expected in this stranger design. Indeed, our findings showed it to be influential even when employing random rematching.

Based on reinforcement learning (e.g. Roth and Erev, 1995) or evolution (Binmore and Samuelson, 1995; Huck and Oechssler, 1999), previous studies theoretically demonstrated that equal sharing could evolve. Our results show sharing equally does evolve. The fact that it declines toward the end of the 100 rounds speaks in favor of the importance of reputational concerns. Players first adopt equality seeking and only reveal naked self-interest when future dealings are no longer imperiled by loss of reputation. It is interesting that this end-game effect is primarily due to the responders’ willingness to accept meager offers. Only a few proposers seemed to be aware of the responders’ change of mind or seem to be willing to exploit it.

The two treatments only differ in the information they offer after each round. Whereas the D-treatment only informs about own play and outcome the P-treatment offers additional population information by informing about the median behavior of the players in the other dyads. We expected and found faster convergence to equal sharing for the P-treatment. The likely reason is that population feedback permits for swifter recognition of what is generally acceptable and what not.

Concerning the effects of regret, we found that experienced regret is significantly correlated with behavior in the hypothesized ways and that the impact is significantly stronger for proposers

than for responders. We did not find a significant treatment effect concerning regret. This suggests that that regret is primarily based on a post-decisional analysis of one's own behavior. Population information does not add much to this analysis.

Finally, in light of our random strangers design, strategic reputational concerns should have been quite weak and, as in repeated public goods experiments, should have become weaker with learning (Ledyard, 1995). But there was no convergence toward rational play in terms of lower offers and lower acceptance thresholds. Except for a very short endgame phase, in which more tolerance of less equitable sharing surfaced, we observed convergence toward equal sharing.

Appendix A. Interactive Decision-Making Experiment

Subject Instructions

Instructions (Original instructions in German. Text in square brackets appeared only in the instructions for the P-treatment.)

A.1. Introduction

Welcome to our experiment! During this experiment you will be asked to make several decisions and so will the other participants.

Please read the instructions carefully. Your decisions, as well as the decisions of the other participants will determine your payoff according to rules which will be explained shortly. The points that you earn during the experiment will be converted to Euros at the rate of 25 points = €1.00. In addition to your earnings from your decisions over the course of the experiment, you will receive a show-up fee of €2.50 for having shown up on time.

Please note that hereafter any form of communication between the participants is strictly prohibited. If you violate this rule, you will be excluded from the experiment with no payment. If you have any questions, please raise your hand. The experimenter will come to you and answer your questions individually.

A.2. Description of the experiment

This experiment is fully computerized. You will be making your decisions by clicking on appropriate buttons on the screen. All the participants are reading the same instructions and taking part in this experiment for the first time, as you are. During the experiment, you will participate in a series of **100** identical rounds.

A total of 32 persons are participating in this experiment. Before the experiment starts, four groups of 8 participants will be randomly formed, and you will be interacting with the same 7 other participants of your group throughout the experiment (how to interact with them will be explained shortly). In other words, you will never be interacting with the participants of other groups.

A.3. Description of the task

During the experiment, you will be interacting with another participant in your group in each round. At the beginning of the experiment, you and the 7 other participants of your group will be assigned to one of two different roles, Proposer or Responder, so that there are 4 Proposer participants and 4 Responder participants in your group. The computer will once again randomly determine your role (Proposer or Responder), and your role will remain the same throughout the experiment. At the beginning of each round, you will be randomly matched with another participant assigned to the opposite role in your group.

What is the task? Proposer decides how much of a given sum to offer Responder, and Responder chooses the minimum she or he would accept. If this minimum is lower than or equal to the offer—the sum is shared according to the offer; if the minimum is

higher—no one gets anything. Here is in more detail on how Proposer and Responder in a pair would proceed. On each of the 100 rounds:

- Proposer has to choose how many point(s) out of **10 points** she/he wants to give to the Responder; in other words, Proposer offers to Responder one of the following integer amounts: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10.
- Without knowing the Proposer's offer, Responder has to choose how many points she/he demands **at least** from the Proposer; in other words, Responder chooses an acceptance threshold from one of the following integer amounts: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10.
- If Proposer's offer is larger than or equal to Responder's acceptance threshold, Proposer will earn 10 points minus the offered amount, and Responder will earn the offered amount. Otherwise, i.e. if the offered amount is smaller than the acceptance threshold, both earn nothing.

A.4. Information feedback

What do you learn about interaction from the result of a round before proceeding to the next round? The computer will inform you and the other participant in your pair of:

- the amount offered by Proposer and the acceptance threshold chosen by Responder
- whether Proposer's offer is accepted or rejected, and
- your payoff and the payoff of the other participant in your pair for the current round.

[The computer will also inform you and the other participant in your pair about how the other three pairs in your group have made decisions. More specifically, you and the other participant in your pair will be informed of:

- the median offer of the other three pairs in your group
- the median acceptance threshold of the other three pairs in your group, and
- the number of other pairs in your group achieving acceptance because the offer had been larger than or equal to the Responder's acceptance threshold.

What is the “median offer/acceptance threshold” of the other three pairs in your group? If three offers take three different values, then the median offer/acceptance threshold is the second highest (or, equivalently, second lowest) value of them. If at least two of three offers/acceptance thresholds are equal, the median is given by the equal numbers.]

A.5. End of the experiment

After completing the experiment, a summary screen will display the total points that you have accumulated and the corresponding earnings in Euros. Please remain at your cubicle until asked to come forward and receive payment for the experiment.

Once you are ready to begin the experiment, please click on the “OK” button on the screen. When everyone is ready, the experimenter will read the instructions aloud, and then the experiment will start. Please remember that no communication is allowed during the experiment. If you encounter any difficulties, please raise your hand. The experimenter will come to assist you.

References

- Benoit, Jean-Pierre, Krishna, Vijay, 1985. *Finitely repeated games*. *Econometrica* 53, 905–922.

- Binmore, Kenneth G., Samuelson, Larry, 1995. An economist's perspective on the evolution of norms. *Journal of Institutional and Theoretical Economics* 150, 45–63.
- Camerer, Colin F., 2003. *Behavioral Game Theory*. New York, Princeton University Press.
- Cooper, David J., Dutcher, E. Glenn, 2011. The dynamics of responder behavior in ultimatum games: a meta-study. *Experimental Economics* 14, 519–546.
- Fischbacher, Urs, 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171–178.
- Fischer, Sven, Güth, Werner, Müller, Wieland, Stiehler, Andreas, 2006. From ultimatum to Nash bargaining: theory and experimental evidence. *Experimental Economics* 9, 17–33.
- Greiner, Ben, 2004. *The Online Recruitment System ORSEE 2.0 – A Guide for the Organization of Experiments in Economics*. University of Cologne (Working paper).
- Güth, Werner, Kocher, Martin, 2013. In: Altman, Morris (Ed.), *Real World Decision Making: An Encyclopedia of Behavioral Economics*. ABC-CLIO, Santa Barbara, CA.
- Güth, Werner, Schmittberger, Rolf, Schwarze, Berned, 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3, 367–388.
- Güth, Werner, Schmidt, Carsten, Sutter, Matthias, 2003. Fairness in the mail and opportunism in the Internet – a newspaper experiment on ultimatum bargaining. *German Economic Review* 4, 243–265.
- Güth, Werner, Schmidt, Carsten, Sutter, Matthias, 2007. Bargaining outside the lab – a newspaper experiment of a three-person ultimatum game. *Economic Journal* 117, 449–469.
- Hennig-Schmidt, Heike, Li, Zhu-Yu, Yang, Chaoliang, 2008. Why people reject advantageous offers – non-monotonic strategies in ultimatum bargaining: evaluating a video experiment run in PR China. *Journal of Economic Behavior and Organization* 65, 373–384.
- Huck, Steffen, Oechssler, Jörg, 1999. The indirect evolutionary approach to explaining fair allocations. *Games and Economic Behavior* 28, 13–24.
- Kreps, David M., Milgrom, Paul, Robert Wilson, John Roberts, 1982. Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory* 27, 245–252.
- Ledyard, John O., 1995. Public goods: a survey of experimental research. In: Kagal, John H., Roth, Alvin E. (Eds.), *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, pp. 111–194.
- Ockenfels, Axel, Selten, Reinhard, 2005. Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior* 51, 155–170.
- Ritov, I., 1996. Probability of regret: anticipation of uncertainty resolution in choice. *Organizational Behavior and Human Decision Processes* 66, 228–236.
- Roth, Alvin, Erev, Ido, 1995. Learning in extensive-form games: experimental data and simple dynamic models. *Games and Economic Behavior* 8, 164–212.
- Selten, Reinhard, Buchta, Joachim, 1999. Experimental sealed bid first price auctions with directly observed bid functions. In: Budescu, David V., Erev, Ido, Zwick, Rami (Eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 79–102.
- Selten, Reinhard, Stoecker, Rolf, 1986. End behavior in sequences of finite Prisoner's Dilemma supergames: a learning theory approach. *Journal of Economic Behavior and Organization* 7, 47–70.
- Slembeck, Tilman, 1999. *Reputations and Fairness in Bargaining: Experimental Evidence from a Repeated Ultimatum Game with Fixed Opponents*. Department of Economics, University of St. Gallen, Discussion paper No. 9904.