

Avrahami, Judith; Güth, Werner; Hertwig, Ralph; Kareev, Yaakov; Otsubo, Hironori

Working Paper

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

Jena Economic Research Papers, No. 2010,092

Provided in Cooperation with:

Max Planck Institute of Economics

Suggested Citation: Avrahami, Judith; Güth, Werner; Hertwig, Ralph; Kareev, Yaakov; Otsubo, Hironori (2010) : Learning (not) to yield: An experimental study of evolving ultimatum game behavior, Jena Economic Research Papers, No. 2010,092, Friedrich Schiller University Jena and Max Planck Institute of Economics, Jena

This Version is available at:

<https://hdl.handle.net/10419/56853>

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.



JENA ECONOMIC RESEARCH PAPERS



2010 – 092

Learning (Not) To Yield: An Experimental Study of Evolving Ultimatum Game Behavior

by

**Judith Avrahami
Werner Güth
Ralph Hertwig
Yaakov Kareev
Hironori Otsubo**

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a joint publication of the Friedrich Schiller University and the Max Planck Institute of Economics, Jena, Germany. For editorial correspondence please contact markus.pasche@uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

Max Planck Institute of Economics
Kahlaische Str. 10
D-07745 Jena
www.econ.mpg.de

© by the author.

**Learning (Not) To Yield:
An Experimental Study of Evolving Ultimatum Game Behavior***

by

Judith Avrahami,^a Werner Güth,^b Ralph Hertwig,^c Yaakov Kareev,^a Hironori Otsubo^b

December 8, 2010

Abstract

Whether behavior converges toward rational play or fair play in repeated ultimatum games depends on which player yields first. If responders concede first by accepting low offers, proposers would not need to learn to offer more, and play would converge toward unequal sharing. By the same token, if proposers learn fast that low offers are doomed to be rejected and adjust their offers accordingly, pressure would be lifted from responders to learn to accept such offers. Play would converge toward equal sharing. Here we tested the hypothesis that it is regret—both material and strategic—which determines how players modify their behavior. We conducted a repeated ultimatum game experiment with random strangers, in which one treatment does and another does not provide population feedback in addition to informing players about their own outcome. Our results show that regret is a good predictor of the dynamics of play. 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 learn not only about their own outcome in the previous round but also find out how the population on average has adapted to previous results (path dependence).

JEL classification: C78, C92

Keywords: Ultimatum bargaining game; Reputation; Regret; Learning; Experiment

* We gratefully acknowledge generous financial support from the Max Planck Society.

^a The Hebrew University of Jerusalem, Center for the Study of Rationality and School of Education.

^b Max Planck Institute of Economics, Strategic Interaction Group.

^c University of Basel, Department of Psychology.

1. Introduction

For nearly thirty years, a 2-person bargaining game (Güth, Schmittberger, & Schwarze, 1982) has piqued the curiosity of numerous experimenters. Exasperated by the plethora of investigations involving the ultimatum game, Camerer (2003) recently proposed a ban on further studies of it. Although one can easily relate to Camerer's sentiment, one should not overlook that quite often the ultimatum game is only employed as an experimental vehicle to investigate 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— notwithstanding the myriad investigations of the ultimatum game—only a handful of studies have examined 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 has adapted to previous results (path dependence).

In section 2, we will discuss how feedback information about the outcome of the previous round can influence future behavior, especially when assuming that participants are guided by post-decisional regret as suggested by the qualitative theory of direction learning and the quantitative but static concept of impulse balancing. Section 3 describes our experimental method. Section 4 contains the main results, focusing on the the test of how post-decisional regret shapes behavior. Section 5 concludes.

2. Recursive ultimatum play and post-decisional regret

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, should proposers make the offer (for a recent survey, see Camerer, 2003). Some studies have examined proposer and responder behavior in recurrent

ultimatum games in which the same two partners interact repeatedly. Under these circumstances, the game is assumed to morph into a reputation game (Kreps et al., 1982). Pursuing the goal of building up a reputation, players' early moves take on special importance because they signal a player's type. For instance, by being unwavering at the beginning—that is, proposing the smallest possible share and rebuffing any proposal that deviates from equality—proposers and responders bet on the 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).

Except for the study by Slembeck (1999) that used the ultimatum game as a base game, we are not aware of experiments exploring how partners play the ultimatum game recursively. Suggestive evidence may, however, 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). Because Fisher et al. (2006) have already observed convergence to fair play rather than to opportunistic behavior with random strangers matching, this tendency could be even stronger for a partners matching protocol. But, as shown by Slembeck (1999), reputation formation may question this. Such reputation formation has also inspired theoretical attempts as to when and why convergence to fair play ought to be expected, e.g., Binmore and Samuelson (1995), Huck and Oechssler (1999), and Roth and Erev (1995).

In our view, however, these results are insufficient to warrant the prediction that a stranger design will give rise to equal sharing. Moreover, previous studies have not scrutinized in much detail the question of how feedback in a repeated ultimatum game with strangers affects outcomes and people's propensity to learn. Henceforth, we will focus on two types of feedback: (i) information about one's particular proposer-responder dyad, and (ii) information about the population of which one's dyad is a part. Specifically, our study investigates two

treatments in which players' receive (or fail to receive) information about what others, faced with the same situation, have been doing.

In our *P-treatment*, players received information about the behavior in the *population*. There are different ways to represent information about population behavior. For instance, players could receive distribution information about the relative frequency of occurrence of any possible division, and the likelihood with which it was accepted and rejected, respectively. Such complex distributive information, however, may overtax participants' comprehension, thus compromising its value. For this reason, we presented information about population behavior in terms of simple aggregated properties of the distributions, namely, the median choices and the median acceptance thresholds. To gauge the latter, we asked each responder to provide us with the threshold at which he or she would have been willing to accept or decline a distribution. 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—through which we can observe how people's behavior evolves over time.

In sum, in the *P-treatment*, in every period each participant learnt the median choices of players in both roles, namely the median offer of all proposers and the median acceptance threshold of all responders. In addition to population information, participants also obtained dyad-specific information in every trial, 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. Unlike in the *P-treatment*, the *D-treatment*, our second treatment, only offered dyad-specific information to players; no population information was given.

What kind of behavior will evolve in a repeated ultimatum game involving strangers? 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 *impulse*

balance 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. The theory applies to situations in which feedback after each period permits drawing retrospective inferences about what behavior would have yielded better outcomes. The theory predicts that, egged on by these postdictions, 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 about past decisions that have an effect on future decisions (Selten and Buchta, 1999) and (ii) feelings of anticipated regret (Ritov, 1996) emerging when deliberating 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 has also been suggested that impulse balance equilibria capture in a static way the cyclic behavior of directional learning (Selten and Buchta, 1999). Henceforth, when we speak of the proposer and responder's feelings of regret, we refer to both interpretations, namely, feelings of regret and anticipated regret.

In the repeated ultimatum game, feelings of regret can be triggered by material loss in the recent past or by strategic 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. A proposer's material regret may manifest in the following cognition: "I could have earned more, had I had just made a higher offer." A responder, in contrast, may ponder 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 proposers may feel remorse about having revealed their willingness to offer more equitable shares: "Now, with responders

knowing that we are willing to offer fair amounts, they will not accept less.” Responders, in turn, may feel sorry for having made public their willingness to be content with less: “Now, with proposers knowing that we are willing to accept less, they will offer less.” These manifestations of reputational regret capture the notion that players may be concerned about what their future counterparts read off from the population feedback in the P-treatment.

The notion of reputational regret suggests that, as 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 in the P-treatment the same kind of behavioral escalation, i.e., an unusually high conflict ratio, reported for the ultimatum game involving a partners design (Slembeck, 1999), with proposers aiming to demonstrate that they offer “little” and responder participants, in turn, showing that they not willing to accept “little”.

Our D-treatment does not provide population feedback but merely feedback about one’s outcome in the last period. Relative to the P-treatment, behavior may thus more likely converge toward equal division. However, by requiring responders to also provide their acceptance threshold, responders’ reputational regret may creep into their considerations. But if it does, its impact is likely to be less forceful than in the P-treatment. Of course, one may also speculate in other ways, e.g. that the P-treatment fosters more equal sharing, for instance, when the medians rely on equal shares. This, in our view, demonstrates that, contrary to Camerer (2003), the ultimatum game still can generate new and interesting questions.

In the following section 3, we introduce the experimental scenario and its two treatments (P and D) and discuss our main hypotheses. Section 3 describes the experimental implementation. Our major findings are reported in section 4. Section 5 concludes.

3. Experimental method

As discussed above, we are concerned with how the repeated play of ultimatum games is influenced by post-decisional regret and how this depends on whether additional population feedback information is provided that should speed up convergence – according to our hypothesis – to equal sharing. We are also interested in whether population feedback lowers or increases the conflict ratio and affects the variance in behavior.

3.1. Participants

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 subjects, two groups in the D- and P-treatment, respectively. Each session involved only one group, and none of the subjects 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

3.2. Procedure

The experiment was conducted in the experimental laboratory of the Max Planck Institute of Economics, which contains a network of 32 PCs. The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). Upon arrival at the laboratory, the subjects were seated in individual cubicles separated from one another by partitions. Any form of communication between subjects was prohibited. Once all subjects were seated, they began to read instructions silently at their own pace. Next, an experimenter read the instructions (see Appendix) aloud so that all information became common knowledge. Questions were answered individually by the experimenter.

At the beginning of each session, four equal-size groups of eight subjects were randomly formed, and each subject was randomly assigned to one of two different roles, proposer or

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

responder. Each group thus consisted of four proposers and four responders, respectively. Subjects 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 minutes, including 20 minutes of instruction and payment. During the session, subjects played 100 identical rounds of the simultaneous-move ultimatum bargaining game (as described in Section 2.1). 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 pair. Once a round began, a decision screen appeared on the computer display on which they were asked to submit their decisions. After all 32 subjects 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 pairs per group: (d) the median offer of the other three pairs, (e) the median acceptance threshold of the other three pairs, and (f) the number of other pairs 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 points were converted at the rate of €1 per 25 points in both treatments. The average individual payoff was €1.89 in the P-treatment and €1.84 in the D-treatment, respectively, including a €2.5 show-up bonus.

4. Results

We begin with summary statistics. Averaged across rounds, proposers offered 4.97 points and responders demanded, on average, 4.80 points. The difference between these values is significant ($F(1,127) = 7.35$, $p = .008$, partial $\eta^2 = .056$). Relative to the D-treatment, the P-

treatment proposers 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 ($F(1,127) = 2.47, p = .119$, partial $\eta^2 = .020$ for the overall difference $F < 1$ for the interaction).

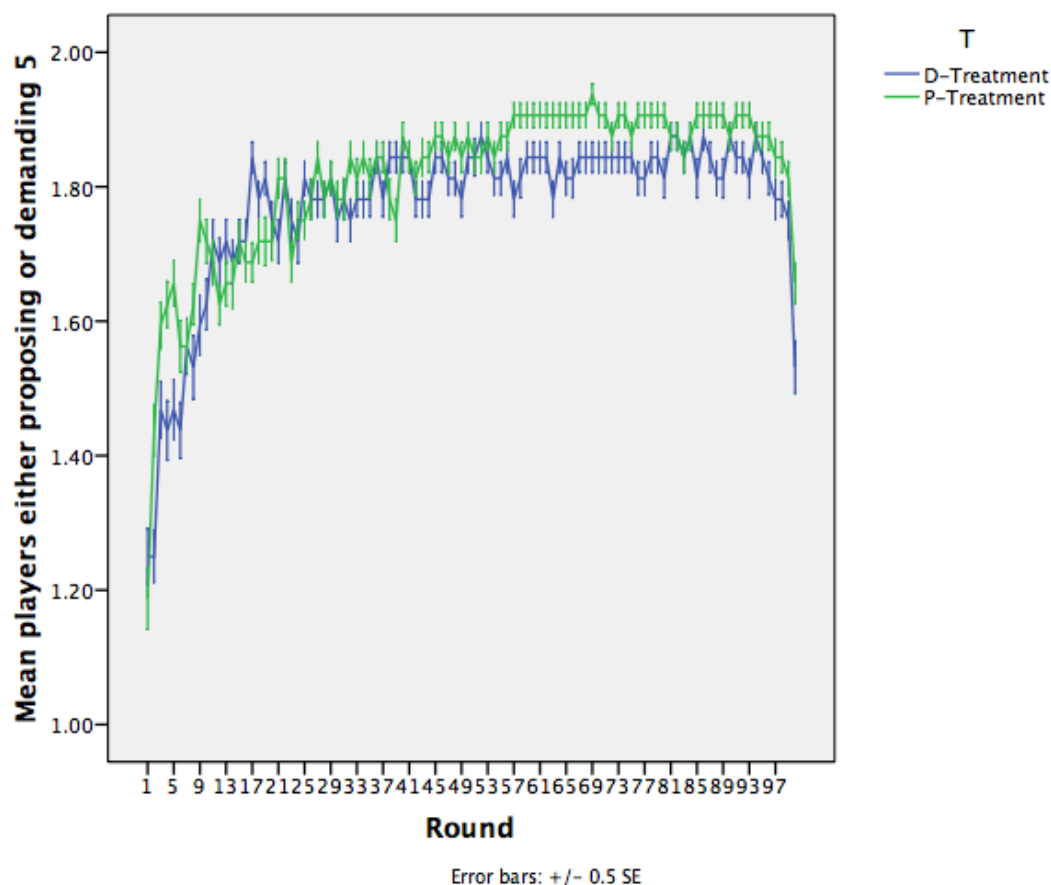


Figure 1. Mean number of players in a pair choosing the equitable division (proposers offering five points or responders demanding five).

How does players' behavior change over time? Figure 1 describes the evolution of equitable offers and demands, namely, the mean number of players in a pair who offered or demanded five points, over rounds, separately for the two treatments. The literature on finitely repeated games (see, for instance, Selten and Stöcker, 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 allows for finite-horizon Folk Theorems (Benoit and Krishna, 1985). Although the argument is thus not a very stringent one, we interpret "cooperation" as peaceful equal sharing up to the end when proposers may be tempted to exploit their ultimatum power.

Figure 1 clearly demonstrates that "cooperation" in the sense of peaceful equal sharing until the endgame effect takes over is not a readily available disposition but rather has to be learned. Maintaining equal sharing (in the sense of proposers offering 5 and responders

accepting only offers of at least 5) is at first rather low but then rather quickly learned, with the learning, however, lasting during the first half of the total of 100 rounds. As Figure 1 shows, the proportion of players who opted for the equitable division is somewhat higher in the P-treatment than in the D-treatment.

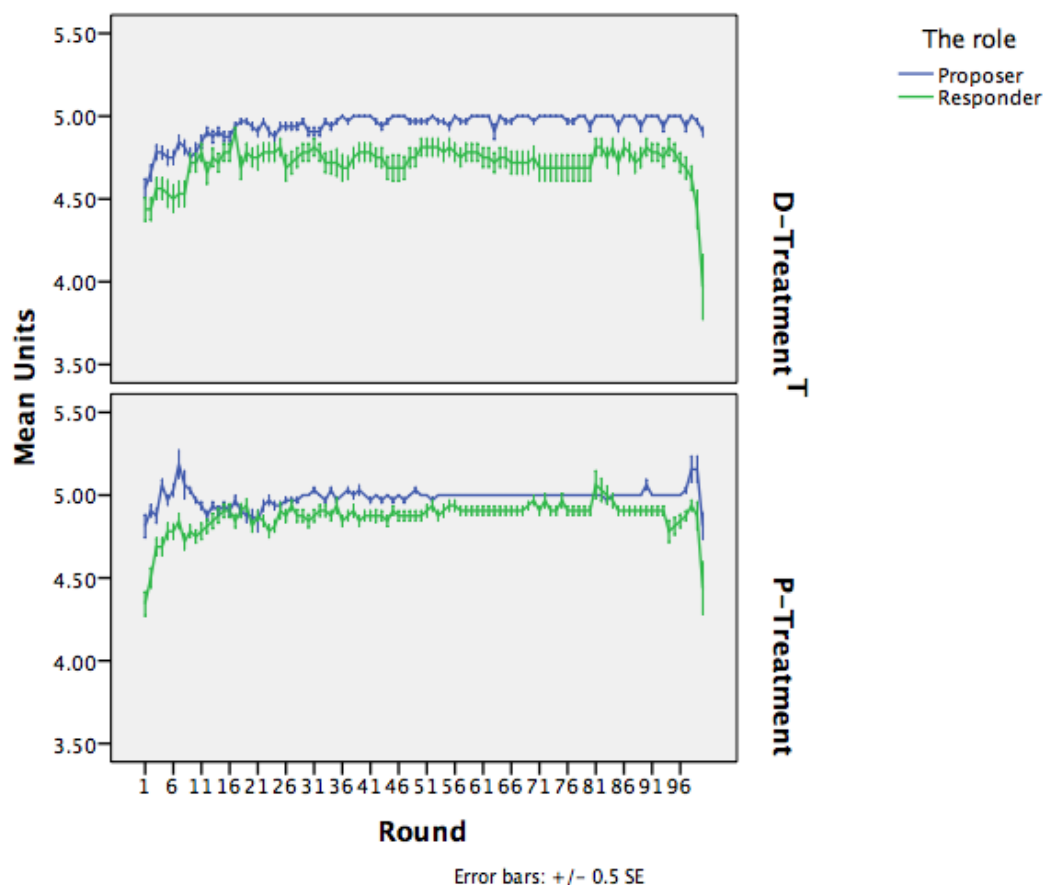


Figure 2. The average points offered (proposer) and demanded (responder), separately for the P- and D-treatments, respectively.

Another obvious property 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 towards the end or because responders lessened their demands? Figure 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 affected the decision in the subsequent 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., 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 the nil amount they earned. 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 what the proposer had offered. 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.

Now let us turn to a successful transaction and the possible regret it may evoke. If the proposers' offers 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.

The above definitions of regret imply that the regret on the part of the proposers ranges more widely. To appreciate this, consider the case of a failed transaction. Because thresholds are practically never higher than 5, the material loss for offering too 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 demand such that the offer exceeds the demand, the proposer experiences a combination of

material and reputational regret whereas the responder experiences only reputational regret. Therefore, we expect regret to better predict the behavior of the proposer than the behavior of the responder.

Having thus defined the impulses evoked by regret, we tested how well a decision in a given period can be predicted by the regret that likely was experienced in the previous period. To find out we regressed, for every player, the change in the number of points chosen between a period and the next (in the offer of proposers and in the demand of responders) on the size (and direction) of the regret in the previous period. Before reporting the results of the regression a caveat is in place. Some participants chose the same value throughout the 100 periods—almost exclusively the equitable value of five. Since a regression cannot be applied to a constant variable, the mean values of the regression analyses to be reported below apply only to participants who varied their decisions, of which there were 80 (out of the total of 128).²

Across roles, the average goodness of prediction in both treatments was $R^2 = .237$. As expected, the proposers' behavior was better predicted by the impulse of regret than that of responders: Specifically, the goodness of prediction was $R^2 = .325$, and $R^2 = .150$, respectively). The average goodness of prediction was very similar for the two treatments. Subjecting these R^2 values to an ANOVA indicated that the difference between roles was significant ($F(1,79) = 13.25$, $p < .001$, partial $\eta^2 = .148$), but neither treatment nor the interaction between treatment and role were significant (both F values < 1)

Recall that in the P-treatment players received information not only about their own pair but also about the median of both offers and demands in their group. They may have thus experienced not only the actual regret but also a hypothetical regret: “Had I played against a prototypical other my offer (or threshold) would have turned out to be too low (or too high)”.

² Some players always chose the same value; they were therefore excluded from the regression analyses

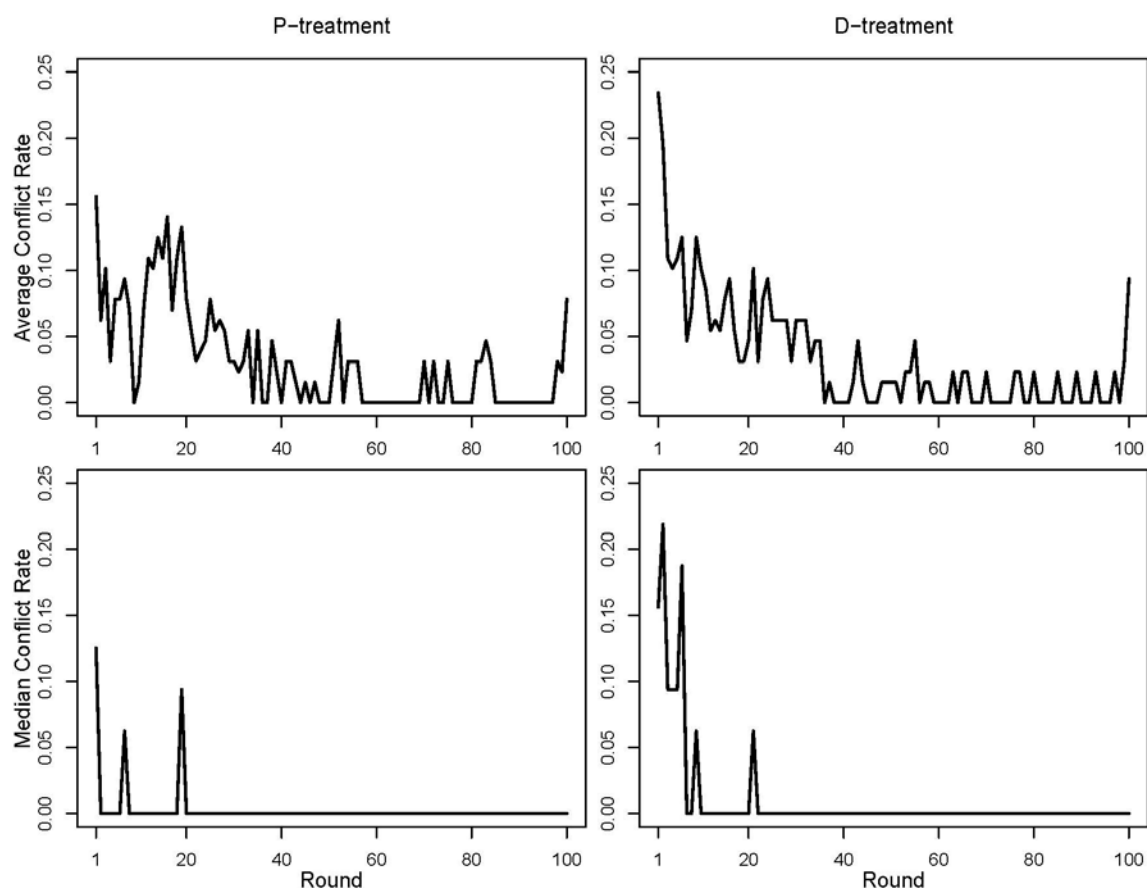


Figure 3. Average and median conflict rates, separately for the P- and D-treatments.

Does such hypothetical regret also have an effect on the subsequent decision? To find out we calculated what we call hypothetical regret and used this as an additional variable to actual regret in predicting future play. We then compared the goodness of prediction based on actual regret only with prediction based on both actual and hypothetical regret. However, since the second predictor includes two variables and the first only one, we compared not the R^2 of the regressions but the BIC values (see Schwarz, 1978) instead.

The BIC values were lower (hence the prediction better) when the hypothetical regret was included to predict current play (-28 without versus -53 with the hypothetical regret; $F(1,36)=10.2$, $p=.003$; $\eta^2 = .22$). The BIC values were also lower (but not significantly) for the proposer than for the responder, in line with earlier analyses (-58 and -23, for proposer and responder, respectively; $F(1,36)=3.4$, $p=.075$, $\eta^2 = .09$).

Let us finally analyze 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. Because conflict may be rare due to random rematching it may be influenced by random pairing. 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 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

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

Figure 3 displays the average and median values of conflict rates of both treatments. The initial conflict rates, based on group averages, seem smaller in the P-treatment than in the D-treatment. Thus, population feedback appears to discourage conflict much earlier apparently because players learn about its detrimental effects not only by their own but also by others' experiences.

5. Concluding Remarks

We conducted a recursive ultimatum bargaining experiment to study the adaptation to the outcome of the previous round, based on the quantitative extension of directional learning (Selten and Stoecker, 1986) as suggested by the static concept of impulse balancing (Ockenfels and Selten, 2005). In section 2, we have distinguished two types of post-decisional regret—material and reputational regret for both roles—and have discussed how they should impact future behavior. Reputational regret would in all likelihood have been stronger and

possibly even dominating in a partners design. But due to the small population size, there were only four possible others, and thus reputation regret could be expected. Indeed, our findings showed it be influential even when employing random rematching.

Previous theoretical studies, based on some form of (reinforcement) learning (e.g. Roth and Erev, 1995) or evolution (Binmore and Samuelson, 1995; Huck and Oechssler, 1999) have mainly demonstrated that equal sharing can evolve. We have shown here that sharing equally does actually evolve. The fact that it declines toward the end of the 100 rounds clearly speaks in favor of the reputation hypothesis: one first mimics equality seeking and only reveals one's greed when future dealings are no longer endangered by loss of reputation. It is interesting that the termination effect is mainly due to an endgame willingness of responders to accept meager offers whereas only a few proposers seemed to be aware of this or, even if they were aware of it, were unwilling to exploit it.

The two treatments only differ in the feedback information that they provide after each round. Whereas the D-treatment only informs about own play and outcome, the P-treatment provides additional population feedback information by revealing the median behavior of the other pairs. We expected and confirmed faster convergence to equal sharing for the P-treatment because population feedback allows for better and earlier recognition of what is generally acceptable and what not.

Concerning the effects of regret, we could confirm that experienced regret is significantly correlated with behavior in the hypothesized ways but more strongly so for proposers than for responders; the difference between roles is significant. We could not confirm a significant treatment effect concerning regret, suggesting that regret is mainly based on a post-decisional analysis of one's own behavior, independently of population information. Only when using both, a person's own actual regret as well as hypothetical regret (if one would have encountered "the median other"), was the impact of regret boosted by population feedback.

Finally, our experiment questions the as-if justification of common(ly known)

opportunistic rationality in economics. Due to our random strangers design, rational reputational concerns should have been quite weak and, similar to recursive public goods experiments, should have become weaker with learning (Ledyard, 1995). But rather than convergence to rational play in terms of lower offers and lower acceptance thresholds, we observed convergence to equal sharing, except for a very short endgame phase, in which more tolerance of less equitable sharing emerged.

References

- Benoit, Jean-Pierre, and Vijay Krishna (1985). Finitely repeated games. *Econometrica*, 53, 905-922.
- Binmore, Kenneth G., and Larry Samuelson (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*. Princeton University Press, New York.
- Fischbacher, Urs (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10, 171-178.
- Fischer, Sven, Werner Güth, Wieland Müller, and Andreas Stiehler (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. Working paper, University of Cologne.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization*, 3, 367-388.
- Huck, Steffen, and Jörg Oechssler (1999). The indirect evolutionary approach to explaining fair allocations. *Games and Economic Behavior*, 28, 13-24.
- Kreps, David M., Paul Milgrom, and John Roberts Robert Wilson (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: John H. Kagal and Alvin E. Roth (eds), *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, pp.111-194.
- Ockenfels, Axel, and Reinhard Selten (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, and Ido Erev (1995). Learning in extensive-form games: Experimental data and simple dynamic models. *Games and Economic Behavior*, 8, 164-212.
- Schwarz, Gideon (1978). Estimating the dimension of a model. *Annals of Statistics*, 6, 461-464.
- Selten, Reinhard, and Joachim Buchta (1999). Experimental sealed bid first price auctions with directly observed bid functions. In: David V. Budescu, Ido Erev, and Rami Zwick (eds), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Mahwah, NJ: Lawrenz Erlbaum Associates, pp.79-102.
- Selten, Reinhard, and Rolf Stoecker (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. Discussion paper No.9904, Department of Economics, University of St. Gallen.

Appendix

Appendix 1: Instructions (Original instructions in German. Text in square brackets appeared only in the instructions for the P-treatment.)

INTERACTIVE DECISION MAKING EXPERIMENT SUBJECT INSTRUCTIONS

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.

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.

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.

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.]

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.