

Rationality and Emotions in Ultimatum Bargaining

Shmuel Zamir¹
The Hebrew University and LEI/CREST

Abstract

The Ultimatum Bargaining paradigm is often thought of as a demonstration of extreme disagreement between experimental evidence and ‘game theoretical predictions’ and the basic assumption of rationality from which they are derived. Using the data of four experiments on Ultimatum Bargaining which I am involved in, I argue that, quite differently from this general impression, rationality in the sense of self-interested motives, is very much present in the observed behavior of both proposers and responders in the Ultimatum Bargaining game. Part of the argument calls for a *broader interpretation* of the notion of rationality than just immediate money maximization and the backward induction argument.

Introduction

Ultimatum Bargaining (UB) is one of the most extensively studied and most intensely discussed paradigms in *game theory*, probably second only to the Prisoners’ Dilemma. Even as a ‘part time’ experimental game theorist, I find myself involved in four works on Ultimatum Bargaining which will be the basis for this talk.

What is the paradigm?

This is an extremely simple-looking, two-player bargaining environment : One player, called the *proposer*, makes a proposal of how to divide a certain sum of money with the other player, called the *responder*, who has the opportunity to accept or reject the proposed division. If the responder accepts, each player earns the amount proposed for him by the proposer, and if the responder rejects, then each player earns zero.

What is so interesting and so intriguing about it?

- Just like the Prisoners’ Dilemma, in spite of its simplicity (even over-simplicity), the Ultimatum Bargaining environment captures a rather common situation, in which the fruits of cooperation are to be shared by two agents who are, for some external reasons, in asymmetric positions.
- As is the case with the Prisoners’ Dilemma, it is claimed that the ‘theoretical predictions’ for Ultimatum Bargaining are ‘not in line’ with the empirical and

¹ I am deeply indebted to R.J. Aumann for valuable discussions and suggestions. I would like to thank my two discussants, Rosemarie Nagel and Jean-Louis Rulliere, for their stimulating remarks. Any further remarks, comments or suggestions are welcome: zamir@ensae.fr

experimental data. In other words, this paradigm seems to present a discrepancy between the *normative* and *descriptive* aspects of game theory.

What are the ‘theoretical predictions’?

A basic rationality assumption asserts that a person will prefer to receive any amount of money to receiving nothing. Under this assumption, the proposer expects the responder to accept any offer that yields a positive share of the stack, and therefore should make the smallest possible offer, leaving virtually the whole stack for himself. This argument reflects the basic rationale underlying the concepts of *backward induction* and *subgame-perfect* equilibrium, which are central notions in economic modeling.

What are the empirical and experimental findings?

Experimental results in a variety of designs and setups have shown that human subjects’ behavior differs considerably from the argument presented above (see Thaler 1988, Roth 1995, Camerer and Thaler 1995, for surveys of previous experiments on UB and some of its variants). Typical offers average about 40%, the 50-50 equal split is quite frequent (often the mode), and offers that deviate substantially from an equal division are typically rejected. These results are rather robust and were confirmed in many variants of the experiments, including some which addressed the importance of the stack’s size² and one addressing the possibility of cultural artifacts (see Roth, Prasnikar, Okuno-Fujiwara and Zamir (1991)).

These results have often been interpreted as an intriguing discrepancy between experimental results and game-theoretic predictions, which leads me to the subject of my presentation: How should these facts be interpreted? Can they be interpreted so as to be in line with the main ideas of game theory and, thus, be incorporated into its descriptive dimension?

Let me make it clear that I do not attempt to provide an exhaustive review of the literature on this topic, which is rather extensive. (A search on *EconLit* found 81 items to which one might add unpublished papers and works in progress, including three that I myself am involved in and will talk about here.) My rather modest objective here is to address what I consider to be the main issues: the compatibility of the observed data with game theory, and in particular with the notion of rationality. I will provide my personal view on this issue in relation to my works on the subject.

Rationality: Is there any in observed UB behavior?

To begin with, let me make a point regarding the statement of the problem we are dealing with: I find it somewhat misleading to say that there is a discrepancy between experimental results and game-theoretic predictions. I think that neither Nash, who provided us with his

² See Hoffman, McCabe and Smith (1996) who ran \$100 versions of the Ultimatum Game and Cameron (1995) who ran a UB game in Indonesia with a stack of size 200,000 Rupiahs, the equivalent of 3 months’ income.

concept of equilibrium, nor Selten, who introduced subgame perfectness, would say that subgame Nash equilibrium, even if it is unique, is *the prediction of game theory*. Even if we are testing game theory as a descriptive theory, there are other solution concepts in addition to subgame perfect equilibrium, of which some, I will argue, may be relevant to the Ultimatum Bargaining data. This may appear to be, and to some extent is in fact, a semantic point. By raising it, I have no intention to avoid or to play down the real problem, but rather to sharpen it. The apparent discrepancy is between the observed behavior in UB and the ‘basic assumption of rationality’ as we stated it – *a rational person prefers receiving any positive amount of money to receiving nothing* – along with its consequence in the form of subgame perfect equilibrium. Most of the research and debate is related, therefore, to the following questions:

- *Is Game Theory, in general, and the notion of rationality in particular, relevant to the Ultimatum Bargaining environment?*
- *To what extent is the observed behavior of proposers and responders in Ultimatum Bargaining compatible (or incompatible) with rational behavior?*
- *What is the appropriate notion of rationality, suitable for analyzing this game?*

Notice that I do not pretend to fully explain the observed behavior in Ultimatum Bargaining. I tend to agree that there are many components of this behavior which lie in disciplines far out of my competence e.g. psychology, sociology, biology and possibly others. My goal is restricted to the issue of the relevance or irrelevance of rationality to this behavior. I will argue that (paraphrasing from Camerer and Thaler (1995)):

Rationality and self-interested behavior are alive and well, even in Ultimatum Bargaining.

In their seminal paper, Güth, Schmittberger and Schwarze (1983) conclude their experimental findings by stating that:

Our experimental results show that in actual life the ultimatum aspects of easy games will not have such extreme consequences: Independent of the game form, subjects often rely on what they consider a fair or justified result.

Although there is wide agreement today that the ‘notion of fairness’ is indeed relevant and plays an important role in explaining behavior in Ultimatum Bargaining, it does not entirely replace strategic and rational motivation but rather affects, supplements and expresses it. Furthermore, the fact that the theoretical rationality predictions are extreme is apparently not the main reason for adopting ‘fairness motivated’ behavior, “independent of the game form.” In the first part of the paper with Roth, Prasnikar and Okuno-Fujiwara, we report on a variant of the Ultimatum Game with nine proposers and one responder. (See also Prasnikar and Roth (1992).) The very extreme subgame perfect outcome, according to which the responder collects the whole stack (10), leaving zero to all other nine players, was perfectly observed in all sessions in all four countries.

In this presentation, I will try to convince you that there is rather clear evidence of self-interest and rationality considerations in the observed behavior of both proposers and responders. As I said before, a full account of all the determinants of observed behavior in

Ultimatum Bargaining is an overly ambitious task, but it seems safe (and quite agreed upon) to admit the existence of two categories of motives: *Rationality and self-interest* (in some larger sense) and all other motives which do not seem to be directly related (if at all) to rationality. For lack of a better name, we may refer to this second category as *emotions* (which could be *sociological, cultural, etc.*) The definitions and the boundaries of these two categories are not clear and may depend on one's personal point of view. I know at least one prominent game theorist, R.J. Aumann, who argues that practically *all* Ultimatum Bargaining behavior is rational if imbedded in a larger evolutionary game theoretic framework, provided we adopt a population interpretation of the Nash equilibrium concept. Although I do not disagree with such a viewpoint, I do not adopt it 'operationally' at this point. *Firstly, I have not yet seen it proved formally.* Secondly, and more importantly, it seems to me interesting to identify considerations of 'short term rationality' (the first category of motives) and contrast them with other motives in the second category, which may (or may not) be given to explanation by rationality in a much larger evolutionary game.

Rationality in proposers' behavior

Although the behavior patterns of proposers and responders are strongly inter-related and constitute two facets of the phenomenon we are looking at, it will be convenient to focus our attention on each of them separately. As we will see, they are driven by rather different motivations. Starting with proposers' behavior, the basic fact is that:

Proposers in Ultimatum Bargaining offer significantly more than they would, were they to act in accordance with the subgame perfect outcome.

The first question to ask is:

- Is this purely a result of 'good nature', or is it the product of a deeply rooted *norm* of behavior?

This was tested by the design of 'dictator games' and 'impunity games' (Forsythe, Horowitz, Savin, and Sefton (1994), Bolton and Zwick (1995)). In these games, the proposer is asked to share the cake with the other person who, unlike in the Ultimatum paradigm, is passive and has no power to affect the proposer's payoff. Independently of the responder's action, the proposer gets the part of the cake he asks for. The results are that offers are much lower than in the Ultimatum Game but (in most of the experiments) they are still significantly positive. This suggests that part of the proposer's generosity in UB is due to *fear of rejection* of low offers by the responder (whether explainable or not). As Ochs and Roth (1989) put it:

We do not conclude that players 'try to be fair'. It is enough to suppose that they try to estimate the utilities of the players they are bargaining with...

Addressing the same issue, Kagel, Kim and Moser (1996) asked the question: Do proposers want to *be* fair or to *look* fair? They designed an Ultimatum Game of incomplete information in which the stake is 100 points but the monetary value of the points may be different for the two bargainers; in particular, the responder may not know the monetary

value of the offer for the proposer and may thus accept a 50-50 split (of the points) even though in monetary terms, it is far from being equitable. Roughly speaking, their result is that *the appearance of fairness is enough*; the proposers offered less whenever they thought they could get away with it.

These results, and many others of similar nature, indicate that there is a strong component of rationality in the proposers' observed behavior. Whether we call it fear of rejection (even if the rejection, as such, is 'irrational') or a desire to look fair just to please the responder, or adaptation to the environment they are in, these are demonstrations of self-motivated behavior which is an important aspect of rationality, fully in line with game theory. This point of view, which by now is quite commonly accepted, is in contrast with the following statement of Thaler (1988):

We have seen that game theory is unsatisfactory as a positive model of behavior. It is also lacking as a prescriptive tool. While none of the subjects in Ochs and Roth's experiment came very close to using game-theoretic strategies, those who most closely approximated this strategy did not make the most money.

While I agree of course with the facts, I totally disagree with the interpretation, which I find to be self-contradictory or, at best, a misinterpretation of the notion of rationality. It basically says that players are not rational since they do not play close to the subgame perfect equilibrium, and those who play it are not making good profits...

Then maybe IT WAS rational not to play the subgame perfect equilibrium?!

Incidentally, I happened to encounter this type of reasoning also among subjects: At the end of one of the Jerusalem sessions of the four-country experiment, as I was paying the subjects the money they earned, one subject came up to me, visibly upset, and said:

'I did not earn any money because all the other players are *stupid!* How can you reject a positive amount of money and prefer to get zero? They just did not understand the game! You should have stopped the experiment and explained it to them...'

In other words, this student clearly understood backward induction and subgame perfectness and had a very clear (but rather narrow) notion of rationality, according to which he was the only rational player there. Yet he exhibited the poorest performance.

In our four-country paper, to express this observed 'rational behavior' more precisely and quantitatively, we provided the following analysis of the data. For each country, we pooled all the response data to compute the 'empirical rejection probability', which is to say, the rejection rate for each value of offer (where there were enough observations). From this we computed the 'expected payoff' for each offer.³ The resulting functions are given in the following figure (taken from the original paper):

³ Thus, for example, if a proposer offers 300 points, he will earn 700 if it is accepted and zero if it is rejected. In the US data, 300 was offered 15 times and accepted 4 times (26.7%). Thus, the expected payoff of offering 300 is $700 \times 0.267 = 186.9$.

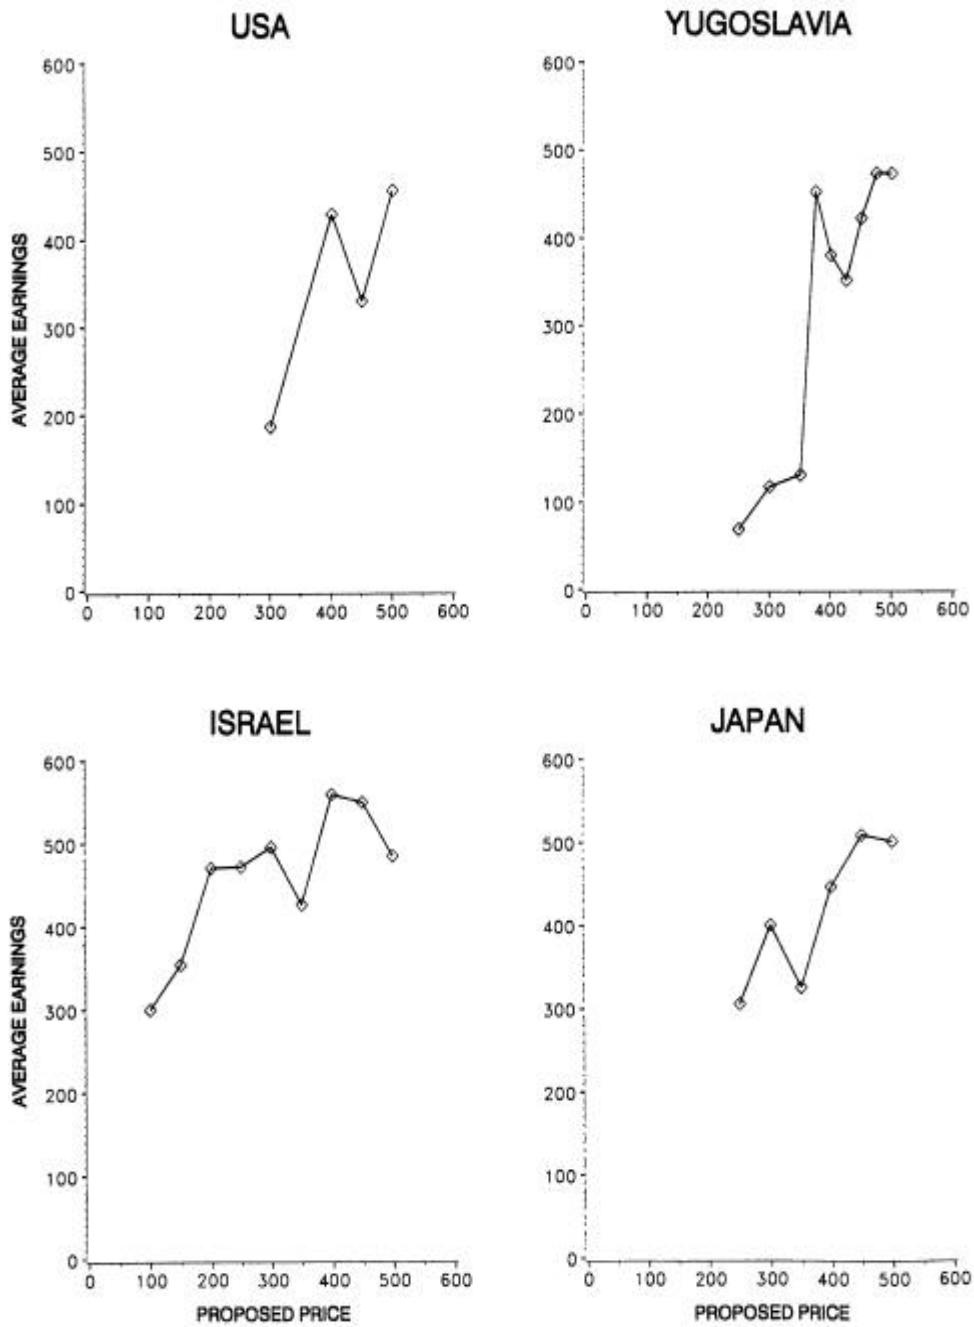


FIGURE 6. BUYERS' EARNINGS IN BARGAINING SESSIONS, BY PROPOSED PRICES
(POOLED DATA OF 10 ROUNDS)

Source: Roth, Prasnikar, Okuno-Fujiwara, and Zamir, 1991: 1090.

We then compared these graphs with the modal offers observed in round 10 (the last round) in each country. The modal offer in the final round in both the US and Yugoslavia was 500, which was also the offer that maximized the proposer's average earnings in these countries. The modal offer in the final round in Israel was 400; here too, this was the offer which maximized the average earning. Finally, in Japan there were two modal offers in the last round, 400 and 450, the latter of which, maximized the average earning of the proposer in this country. We concluded:

Thus, by round 10, the buyers (proposers) seem to be adapting to the experience of the prior rounds in a manner roughly consistent with simple income-maximization.

In fact, what we are seeing here is a kind of *best reply*, which is not only evidence for rationality and income-maximization but, as we all know, is the definitive concept behind Nash equilibrium. In other words, one is tempted to speculate and say that *we may be observing Nash equilibrium!!* Not subgame perfect Nash equilibrium, but Nash equilibrium nevertheless.⁴

At this point, this is no more than speculation since in a one-shot game, any rejection of a positive payoff is not rational and cannot be part of an equilibrium. If there is any equilibrium, it must be in some larger game and the equilibrium condition must be interpreted in some 'population sense'. As long as this is not done formally, it is still safe to say that proposers' observed behavior in the Ultimatum Game has a significant strategic and self-interest component. Roughly speaking, proposers adjust to their environment and try to extract high profits.

The observation that the proposer is basically 'best replying' to his environment led Winter and Zamir (1997) to further test this hypothesis by conducting an experiment with Ultimatum Bargaining in a changing environment. The idea is straightforward:

- If the players' behavior is, to a large extent, a reaction to their environment, can we then induce changes in their behavior by changing the environment?
- If such effects exist, what is the 'time scale' of these changes and how difficult is it to induce them?

To study the effect of the environment, we departed from the conventional setup of UB experiments, extending it by an additional experimental tool. Our population of players included both real subjects and virtual players; the latter were computer programs that played the roles of both proposers and responders by using fixed strategies specified at the beginning of the experiment. In each experiment, the UB game was played over and over again for a large number of rounds. At the beginning of each round, subjects were matched

⁴ As we know, the Ultimatum Bargaining game has many Nash equilibria. In fact, if the size of the stack is S then for any x between 0 and S , the split $(S-x, x)$ is an equilibrium outcome. The corresponding equilibrium strategies are: The proposer offers x to the responder and the responder's strategy is to reject any offer smaller than x and to accept any offer of x or more.

randomly either to another real player or to a virtual player. None of the real subjects knew about the presence of virtual players. From their point of view, they were playing a regular UB game with a conventional design.⁵

As I said, the objective of this design was to explore the way real subjects' behavior changes as a function of the type of virtual player in the experiment. The main question this paper addressed was again: what elements determine individual behavior in UB? Should one ascribe differences in behavior to differences in some deep cultural or educational attributes of individuals, or can they be explained as outcomes of responses to different environments?

We constructed two types of virtual proposers and responders. The first type, which we call "tough" (in two levels of 'toughness'), consists of proposers and responders who form an equilibrium that is closer to the subgame perfect outcome; i.e. the proposer makes low offers and the responder accepts low offers. The second type, which we call "fair", involves proposers and responders who form an equilibrium outcome which is close to the 50:50 division; i.e. proposers make offers around the equal share and responders reject offers yielding considerably less than 50%.

A virtual proposer is a computer program designed to submit offers at random from a fixed specified range. We designed two types of "tough" proposers, one (extremely tough) whose offers are sampled (randomly and uniformly) from the interval between 13 and 16 points and the other (moderately tough) whose offers are between 23 and 26 points. The "fair" virtual proposers all draw offers between 46 and 49 points.

For each type of virtual proposer we constructed a compatible virtual responder. For example, a virtual responder compatible with the 13-16 tough proposer is a computer program designed to draw an acceptance threshold value from the same set of offers, 13-16. If, for example, the threshold drawn was 14, then this virtual responder will accept any offer of more than 14 points and will reject all other offers. We denote by $P_{\{13,16\}}$, $P_{\{23,26\}}$ and $P_{\{46,49\}}$ the three types of virtual proposers and by $R_{\{13,16\}}$, $R_{\{23,26\}}$ and $R_{\{46,49\}}$ the corresponding three virtual responders. One may think of these virtual players as biological 'mutants' or members of a different society with different norms of behavior.

In each typical session a different group of subjects was received in the laboratory. Before commencing, a lottery determined the role of each subject (proposer or responder), which was fixed throughout the session. The subjects played the UB game for either 50 or 70 rounds, depending on the session. In each round, the set of proposers and virtual proposers was matched randomly with the set of responders and virtual responders.

⁵ In a second set of experiments (6 sessions), subjects were told of the presence of virtual players. Specifically, they were told that during the course of the game they may be matched to a computer program instead of to a real player. However, they were not told anything about the probability of this event or about the nature of these computer programs. No significant differences were observed between the data of the two parts of the experiment.

For example, in one experiment the "society" consisted of a group of 12 real players (6 proposers and 6 responders) and a group of 8 virtual players (4 of $P_{\{23,26\}}$ and 4 of $R_{\{23,26\}}$). The random matching was designed to guarantee that all virtual players would be matched to real players. Usually, the number of virtual players was fixed throughout the session, but in two sessions (with virtuals $P_{\{13,16\}}$ and $R_{\{13,16\}}$) we gradually increased the population of virtual players.

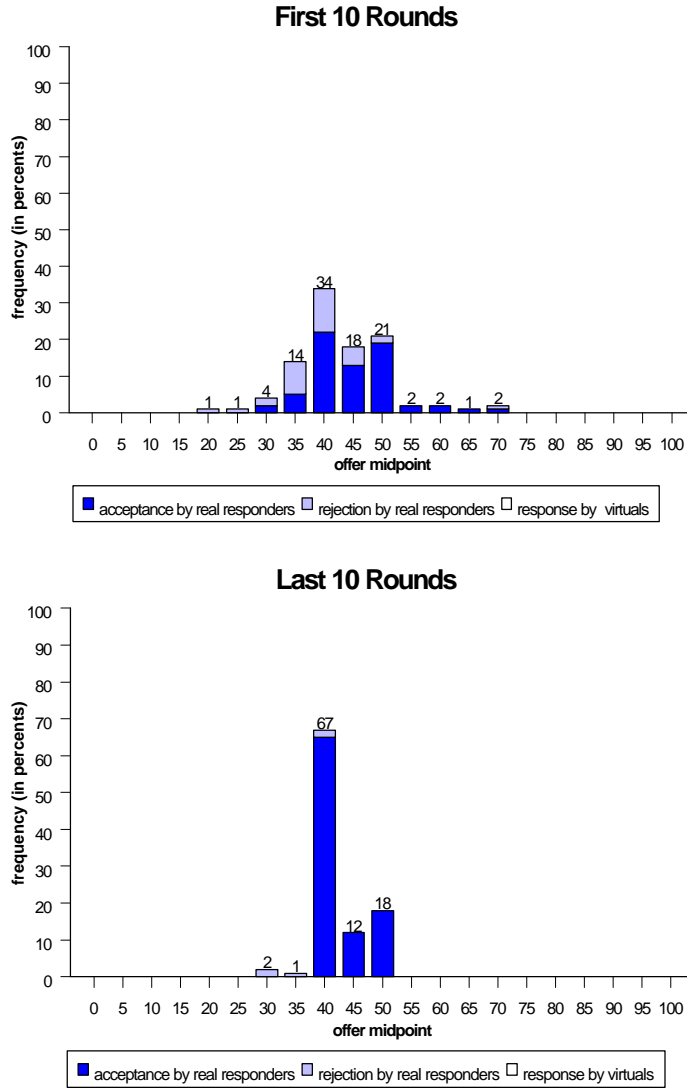
The main conclusion of these experiments was that the presence and identity of virtual players dramatically affects real subject behavior in UB. The most unambiguous result of this experiment is perhaps the effect of the presence of virtual players on the offers made by real players. This is summarized in the following table and figures:

Table 2: Mode, mean and standard deviation of offers by real players

Session		Total	First 10 Rounds	Last 10 Rounds
# Real Players	# and Type of Virtual Players			
12	-	40, 39.47, 7.06	50, 38.80, 13.37	40, 40.45, 2.81
20	-	40, 40.95, 6.81	40, 43.13, 8.26	40, 42.13, 4.21
12	4 – $P_{23,26}$; 4 – $R_{23,26}$	30, 36.89, 11.92	50, 47.35, 13.38	30, 32.17, 8.78
20	7 – $P_{23,26}$; 7 – $R_{23,26}$	30, 35.48, 7.76	40, 38.11, 9.48	30, 32.71, 5.71
12	½ are $P_{13,16}$; ½ are $R_{13,16}$	40, 39.44, 8.45	50, 45.50, 8.42	30, 35.07, 8.00
20	½ are $P_{13,16}$; ½ are $R_{13,16}$	20, 34.84, 12.26	50, 40.20, 17.33	20, 28.94, 10.43
12	4 – $P_{23,26}$; 4 – $R_{23,26}$	50, 45.20, 9.67	50, 42.85, 12.71	50, 46.53, 7.2
18	6 – $P_{46,49}$; 6 – $R_{46,49}$	50, 48.88, 3.17	50, 47.27, 4.84	50, 49.00, 3.47

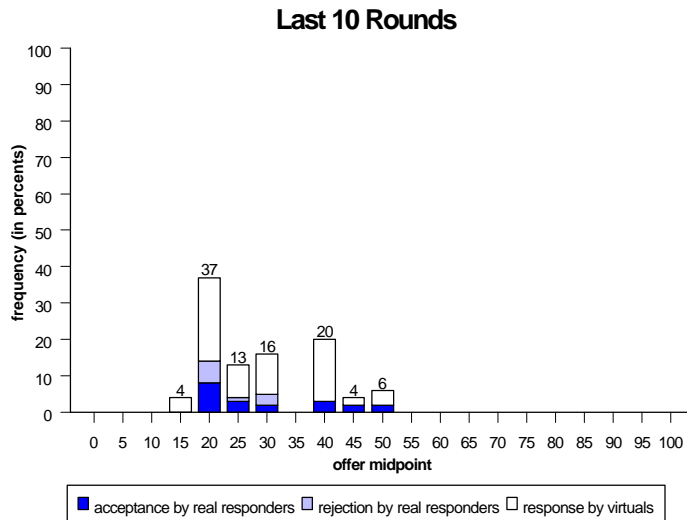
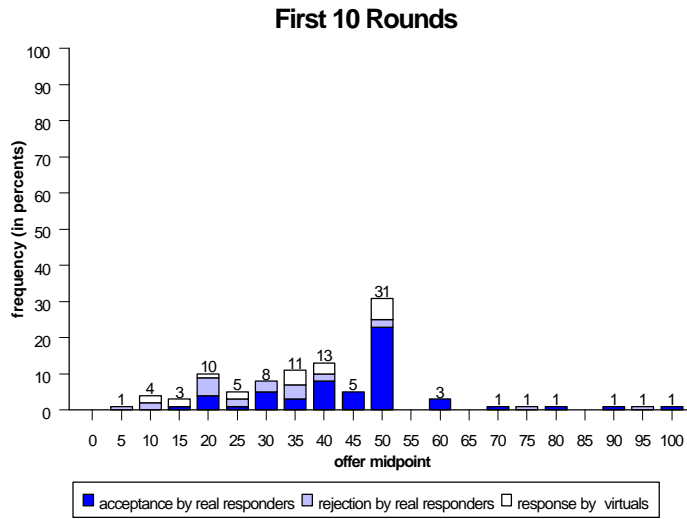
Source: Winter and Zamir, 1997: 6-7.

Figure 1.2
 Relative Distribution of Offers by Real Players
 no virtual players (20 players)



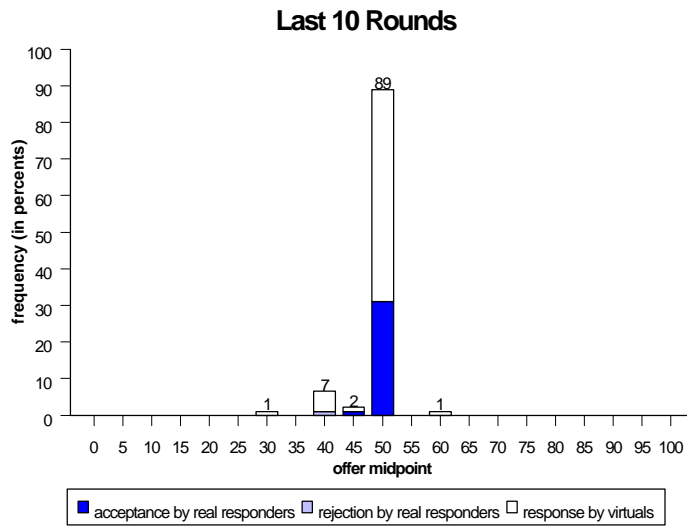
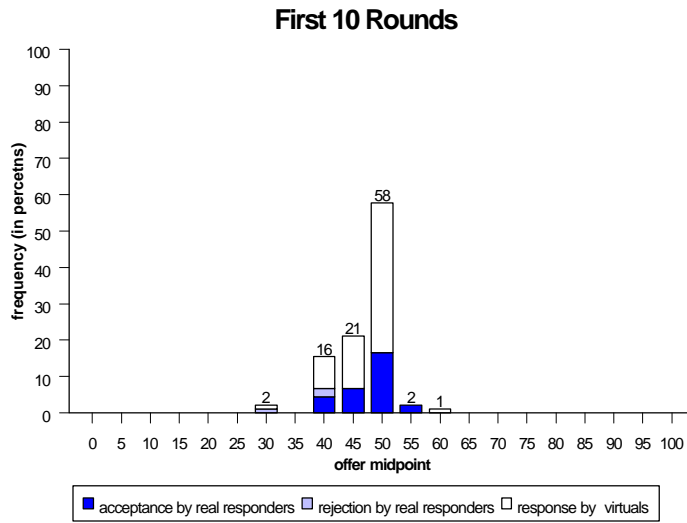
Source: Winter and Zamir, 1997: 92.

Figure 1.6
 Relative Distribution of Offers by Real Players
 virtual offer range = 13-16 (gradual, 20 players)



Source: Winter and Zamir, 1997: 13.

Figure 1.8
 Relative Distribution of Offers by Real Players
 virtual offer range = 46-49 (18 players)



Source: Winter and Zamir, 1997: 15.

Without virtual players, the distribution mode either shifts around 40 - 50 points or remains at 40. When introducing moderately tough virtual players ($P_{\{23,26\}}$ and $R_{\{23,26\}}$) the mode drops to 30 points and with extremely tough virtual players ($P_{\{13,16\}}$ and $R_{\{13,16\}}$) it sinks to 20, in spite of the fact that virtual players were introduced gradually. With fair virtual players the behavior is strikingly different: offers below 50 points vanish almost completely, and the distribution is unambiguously concentrated on the 50:50 offers. One observation which is consistent across all sessions is that the distribution of offers in the first 10 rounds is more widely dispersed than that in the last 10 rounds. This is due to the fact that the learning effect is stronger in early rounds of each session. Within this learning process, proposers “test” the reactions to various levels of offers.

Our interpretation of the difference between offer distributions across sessions is quite simple. There is nothing sacred about offers of 50% of the stack. Although they seem to be very popular in the environment without virtual players (as confirmed by so many other experiments of UB), such offers are not merely the result of proposers being concerned with equality or of the focal point of 50% having any special attraction. Our results show that an important element of proposers' behavior is driven by utilitarian considerations. In the long run, such offers remain attractive because they pay well; i.e. they respond best to the rejection patterns of the responders. If the rejection pattern of responders changes, which is indeed the case for environments with virtual players, proposers will change their behavior as well, to match the new environment. In an environment with tough (virtual) players (sessions 3, 4, 5 and 6), in which virtual responders accept low offers, *proposers learn that offering 50% is wasteful because lower offers have a high chance of being accepted*. In fact, the dynamics that shift the mode of the distribution towards lower offers is somewhat more complicated. There is a direct effect on proposers' behavior through the match to virtual players, as explained above. But there is also a weaker, indirect effect: the effect of persistent low offers by virtual proposers may lead real responders to expand their acceptance sets. These real responders, when meeting real proposers, will induce them to make low offers in the same way virtual responders do.

With "fair" virtual players the story is pretty much the same. However, here the behavior of virtual players is much closer to the initial patterns of real subject behavior. Proposers do not need to experiment long with low offers before realizing that they do not work well, and consequently the convergence to 50:50 offers is fast and unambiguous.

The dynamic of the change in offer patterns is seen in the following figures which show the way the modes and average offers evolve in time in various environments.

Figure 2.1
Mean and Mode of Offers by Real Players as a Function of Time
 virtual offer range=no virtuals (12 players)

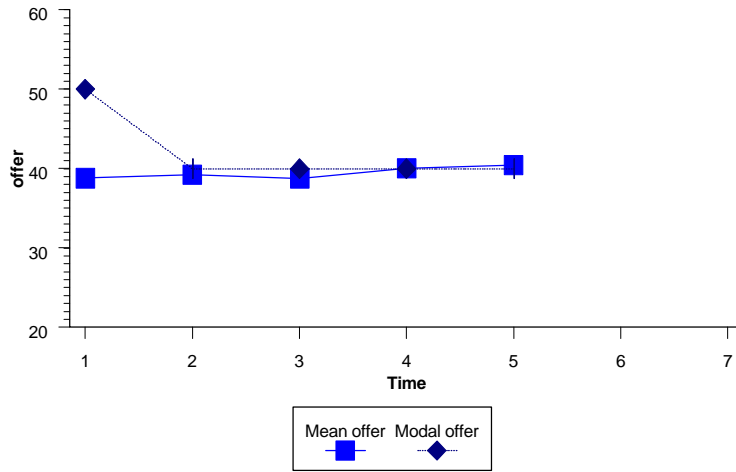
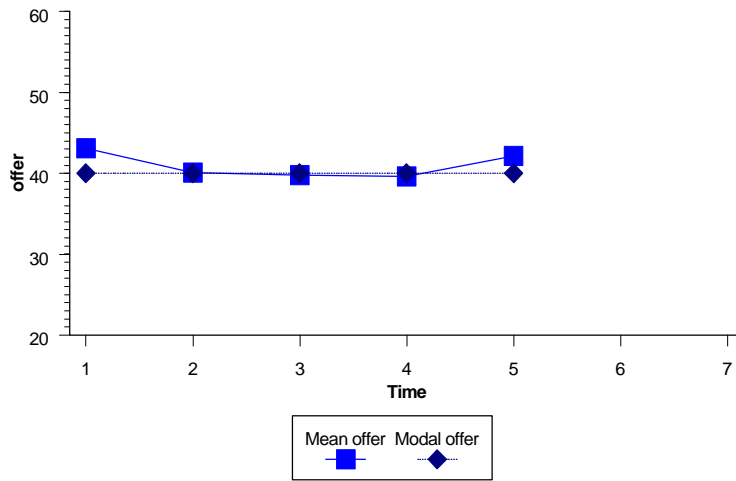


Figure 2.2
Mean and Mode of Offers by Real Proposers as a Function of Time
 virtual offer range=no virtuals (20 players)



Source: Winter and Zamir, 1997: 17.

Figure 2.3
 Mean and Mode of Offers by Real Players as a Function of Time
 virtual offer range = 23-26 (12 players)

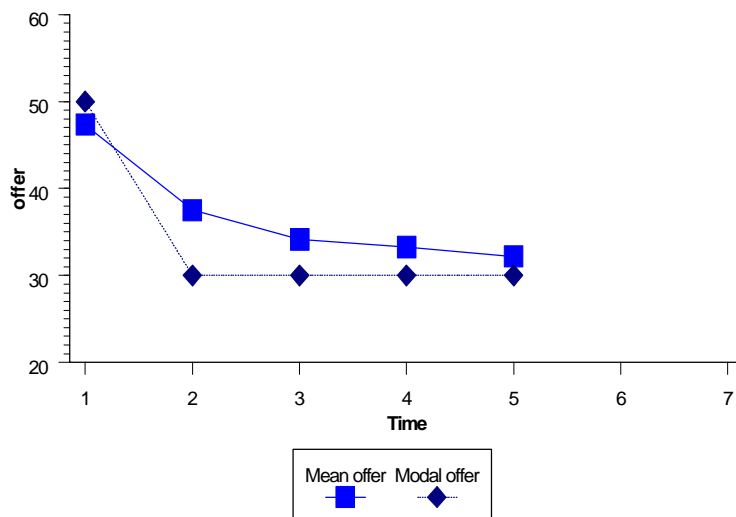
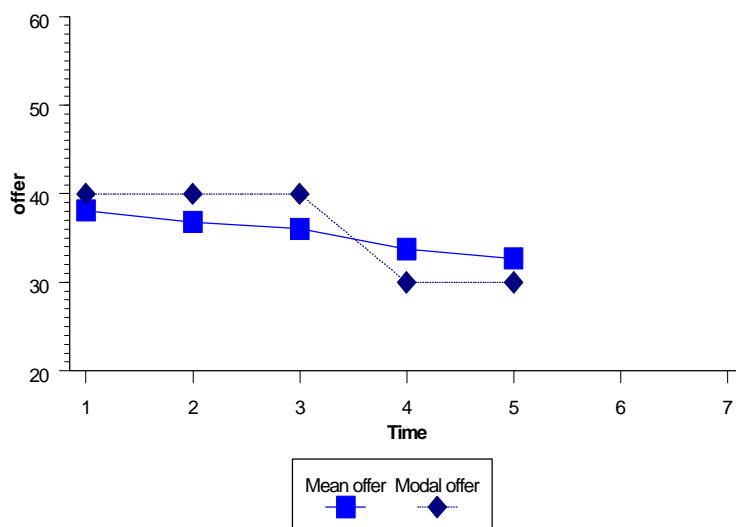


Figure 2.4
 Mean and Mode of Offers by Real Proposers as a Function of Time
 virtual offer range = 23-26 (20 players)



Source: Winter and Zamir, 1997: 18.

Figure 2.5
 Mean and Mode of Offers by Real Players as a Function of Time
 virtual offer range = 13-16 (12 players)

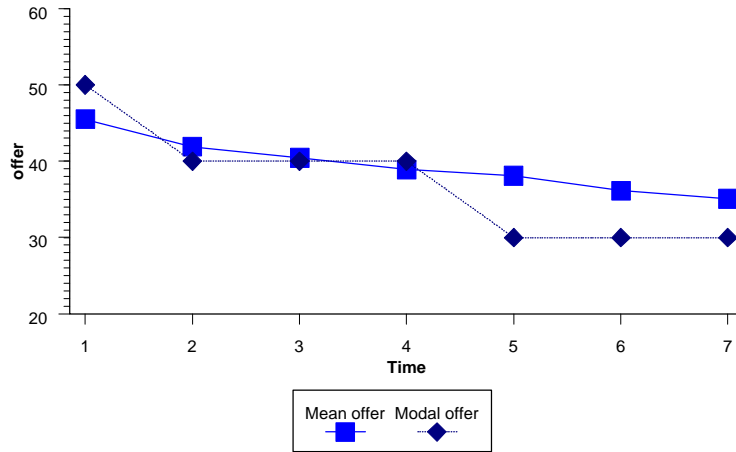
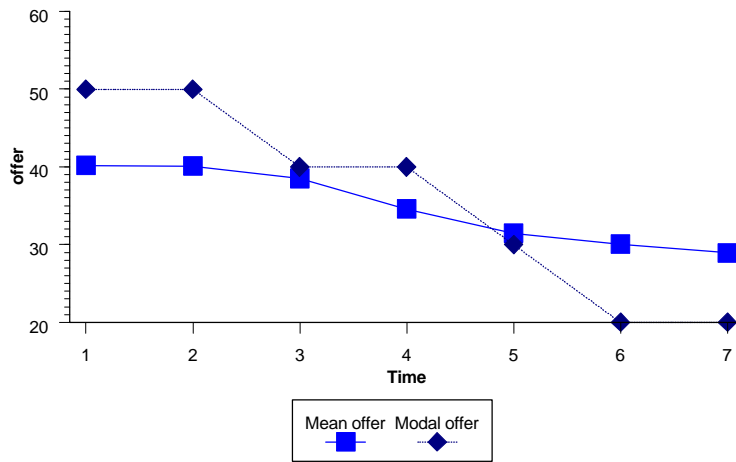


Figure 2.6
 Mean and Mode of Offers by Real Proposers as a Function of Time
 virtual offer range = 13-16 (20 players)



Source: Winter and Zamir, 1997: 19.

Figure 2.7
Mean and Mode of Offers by Real Players as a Function of Time
 virtual offer range = 46-49 (12 players)

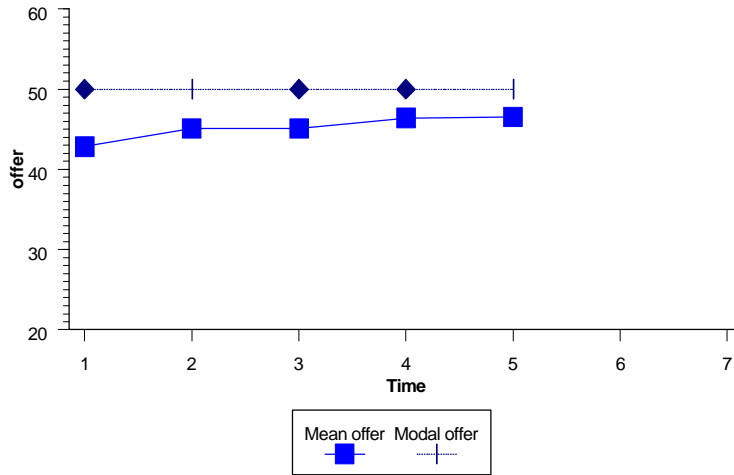
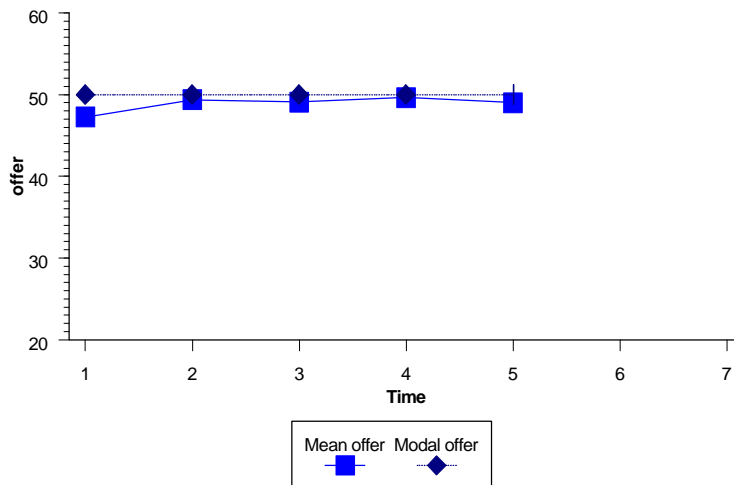


Figure 2.8
Mean and Mode of Offers by Real Proposers as a Function of Time
 virtual offer range = 46-49 (18 players)

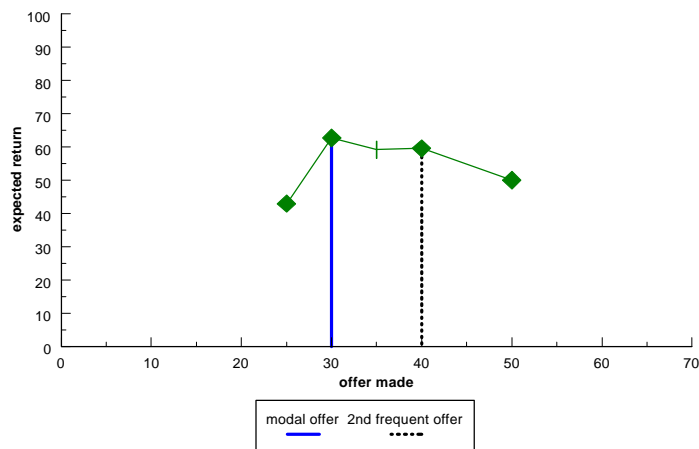


Source: Winter and Zamir, 1997: 20.

It is apparent from these figures that the environment changes gradually. With tough virtual players, these indicators constantly decrease and with fair virtual players they increase gradually.

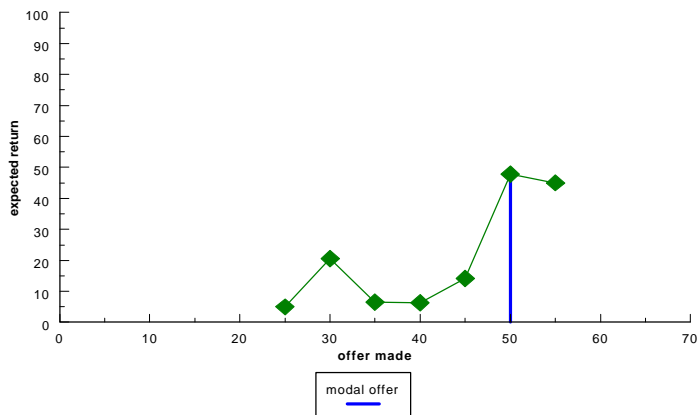
To check again the existence of a ‘best reply’ component in the behavior of the proposers, we produced, as we did in the four country paper, the ‘average earning’ curves in relation to the observed modal offer. The following are several examples:

Figure 3.4
Expected Returns for Offers Made by Real Players
 based on virtual and real responses
 virtual offer range = 23-26 (20 players)



Source: Winter and Zamir, 1997: 23.

Figure 3.7
Expected Returns for Offers Made by Real Players
 based on virtual and real responses
 virtual offer range = 46-49 (12 players)



Source: Winter and Zamir, 1997: 27.

Is there any rationality in responders' behavior?

Responders' observed behavior in the Ultimatum Bargaining game is considered to be very hard to rationalize. Why do responders reject when this just means leaving money on the table? Such behavior is quite obviously not rational, if subjects are only concerned about their own payoffs and the game is only played once. Thus, one approach to explain rejections in the Ultimatum Game is to assume that responders prefer receiving nothing to receiving little, because the rejection creates an allocation in which the proposer also receives nothing. This approach is related to a fairness argument: Responders are willing to pay the price of receiving nothing because of some inherent resistance to unfairness (Thaler 1988). While I do not deny the relevance of this argument and possibly the existence of such a 'natural' attribute of 'revolt against unfairness and injustice', I will argue that there are, most likely, other motives behind responders' behavior, some of which have a strong flavor of (even short-term) rationality. Let me already express my viewpoint that incorporating 'fairness' as a variable in the utility functions does not *explain* the phenomenon, rather, at best, describes it, which may be useful in an econometric type of analysis.

Let me start with a perhaps provocative remark, saying that the 'notion of fairness', as it manifests itself in these environments, is often used strategically and hence is not really free from 'strategic flavor'. Practically all experiments in this area confirm that this explanatory variable called a 'notion of fairness', is *context dependent*. In the same subject pool of university undergraduate students, a student that happened to be a responder employed the 'notion of fairness' differently than a student who happened to be a proposer and even more dramatically so, than a proposer in a dictator game. A responder in an ordinary Ultimatum Game who considers 50-50 to be *the* fair offer, rejects it (and rejects even a 60% offer) when facing three competing proposers (in the competitive Ultimatum Game I will speak about shortly). If this is not enough to suggest that 'fairness is often used strategically', then it is certainly enough to sustain the claim that using 'resistance to unfairness' as the main explanation for responders' behavior is not quite satisfactory. This also indicates that we cannot expect much help from two variable utility functions (depending on both bargainers' payoffs), unless we are willing to accept that these functions strongly depend on the strategic role in the game...

Now, there is something disturbing in saying that 'rejections are irrational' in view of the striking experimental evidence that, although responders leave some money on the table by rejecting low offers, on average they receive greater payoffs than in subgame perfect equilibrium play. Does this mean that it is not only fairness, but rather some kind of monetary self-interest that drives rejection behavior *even in the one-shot game*? There is something even more disturbing, logically, in saying that the proposer is to some extent rational, taking into account the responders' behavior which includes (non income-maximizing) rejection of positive offers (see e.g. Roth, p. 287 in the discussion of the four country experiment): When saying that the 'non-rational' rejections make the proposer offer more, how far is this from saying that these rejections are, after all, rational, even under the strict definition of rationality as profit maximizing?

How to formalize these arguments?

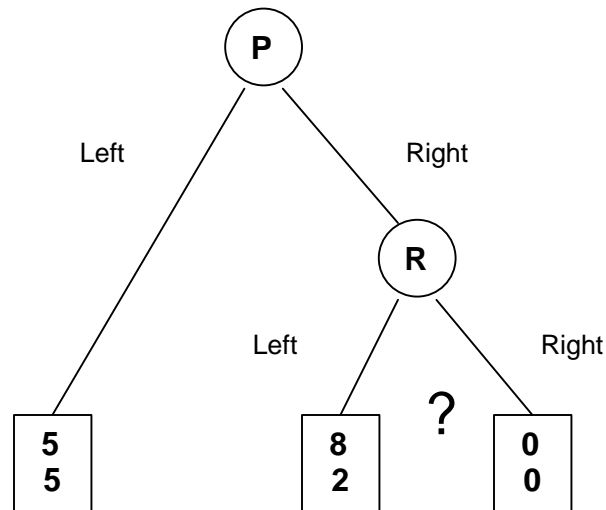
One way is to resort to a multistage or super-game setting, where there are possibilities for reputation building, punishment etc. Then one could argue that although subjects fully understand the rules of the game and its payoff structure, their behavior is influenced by an unconscious perception that the situation they are facing is part of a much more extended game of similar real-life interactions. It may even be that it is practically impossible to create laboratory conditions that would cancel out this effect and induce subjects to act as if they were facing an anonymous, one-shot UB situation. Real life UB situations are typically not anonymous and often involve repeated play between the same two players. In such environments, rejections play an important and rational role of reputation building, acting as a ‘rule of behavior’ that seems to work well, evolutionarily speaking. Such a rule “has not been consciously chosen and will not be consciously abandoned” (Aumann, 1997). A responder who nods at every offer will easily teach proposers to make him low offers, and his overall stream of payoffs will be pretty poor. Real life interaction in conflicts similar to UB, gives rise to certain conventions to which players adhere, without giving them much conscious thought. They are internalized as ‘rule rationality’ in contrast with ‘act rationality’. That is, there is a meaningful notion of equilibrium among ‘rules of behavior’ (Aumann, 1997; see also van Damme, 1998).

Is there any short-term rationality in responders’ behavior?

As long as we refrain from extending the framework to multistage or evolutionary games where rejection (or even ‘resistance to unfairness’) can be rationalized, we still ask whether there is any evidence of rationality, even in the short duration laboratory interaction? While reputation on an individual level cannot explain the success of responders in receiving payoffs that by far exceed the subgame perfect equilibrium outcome in the Ultimatum Game, perhaps reputation on a group level can? If the population of responders has a group reputation for being “tough” (i.e. rejecting low offers), then proposers fearing rejections may increase their offers, thus driving responders’ payoffs up. In this sense, the group reputation may be viewed as a *public good* for the population of responders and *each individual responder can contribute to this public good by rejecting low offers*. Hence, rejections in the Ultimatum Game may be due to some form of *population rationality*. Can we observe this kind of (short-term) population rationality? Can we distinguish it from ‘resistance to unfairness’ as two separate motives for rejections in the Ultimatum Bargaining game? In other words, can we answer the question: ‘*Do responders happen to earn more because (‘by nature’) they resist unfair offers, or do they choose to reject unfair offers in order to earn more?*’ This was the objective of the experiment of Abbink, Sadrieh and Zamir (1999).

In our experimental design we used a discrete simplified version of the Ultimatum Game which we found more convenient when focusing only on the issue of ‘fair’ versus ‘unfair’ offers.

The Covered Response Ultimatum Game



Source: Abbink, Sadrieh, and Zamir (February 1999): 5

The simple modification we applied to the Ultimatum Game was that in one treatment (the covered response treatment) the responder's decision was not reported to the proposer immediately, while in a second treatment (the open response treatment), it was. Rejections observed in the covered response treatment cannot be interpreted as contributions to group reputation since there is no possibility to contribute to this reputation when the proposers do not observe the responses to their offers before the end of the game. On the other hand, fairness utility models cannot explain differences across the two treatments because there are neither strategic nor payoff differences across treatments that would allow divergent predictions by fairness utility models.

Each subject played the same role, proposer or responder, during the entire session. The experiments were run using a revolving (or round-robin) matching scheme, such that each proposer met each responder only once. This was particularly important for the control condition with open responses; each subject was sure that he would never be matched again with the same partner, leaving no room for repeated game effects such as individual reputation. To ensure compatibility, we used the same matching scheme in the covered response treatment as well. Eight rounds were played with eight proposers and eight responders.

The experimental results in this study strongly suggest that responders' behavior is motivated both by resistance to unfairness and by contribution to group reputation, but neither explanation is sufficient on its own. This is readily seen from the following table.

Average rates of rejected unequal offers in each session (in percent)							
covered	6.5 ⁺	17.6 [*]	22.7 ⁺	26.1 [*]	31.0 [*]	31.3 ⁺	Ø=22.5
open	10.0 ⁺	20.0 [*]	40.0 [*]	43.5 [*]	50.0 ⁺	80.0 ⁺	Ø=40.6

Source: Abbink, Sadrieh, and Zamir, 1999: 7. ⁺=Bonn session; ^{*}=Jerusalem session

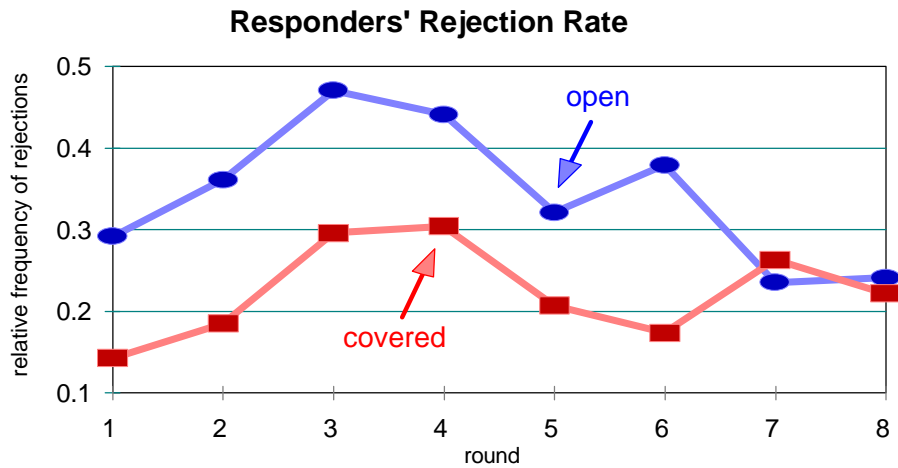
Observation 1: *In the covered response treatment, substantially positive rejection rates are observed.*

Even when the response is not reported to the proposer, almost one quarter of all unequal offers are turned down. This is clear evidence for responders' resistance to unfairness that is entirely independent of all considerations of short-run monetary payoff maximization. Responders cannot openly punish proposers in the covered response treatment - not even as a group - in order to receive higher payoffs later. Thus, the relatively high average rejection rate (22.5%) is evidently motivated by negative emotions towards unfair actions or distributions. Responders are willing to pay a price solely to soothe their anger concerning the proposer's greed. The negative utility of unfair outcomes that in some way or another is contained in all fairness utility models can be interpreted as a formalization of such an emotional component.

Observation 2: *In the open treatment, unequal offers are rejected significantly more often than in the covered treatment.*

Fairness utility models do not explain a different aspect of the data. There is a significant difference in the rejection rates of the two treatments. The average rejection rates of the open response condition are about 75% higher than those of the covered treatment (The difference is significant with a p-value of 0.051, according to the Mann-Whitney U-test applied to the average rates of rejection in sessions). The extent to which more rejections are observed in the open response treatment *must be attributed to the visibility of the rejection*. The difference caused by visibility could be based on the fact that visibility turns rejection into an act of reciprocal punishment. Since visibility allows for the reciprocation to be uniquely ascribed to the reciprocating subject, the connection between punisher, the reason for punishment, and the addressee of the punishment becomes unambiguous. This - in a second step of reasoning - allows for an immediate perception of the educational goal of punishment by the proposers. In this sense, the visibility of punishment enables responders to educate proposers with each rejection of an unfair offer and, thus, to build up a group reputation of being "tough". Obviously, proposers facing a group of "tough" responders will tend to switch to more equal split offers.

Other results of this study were:

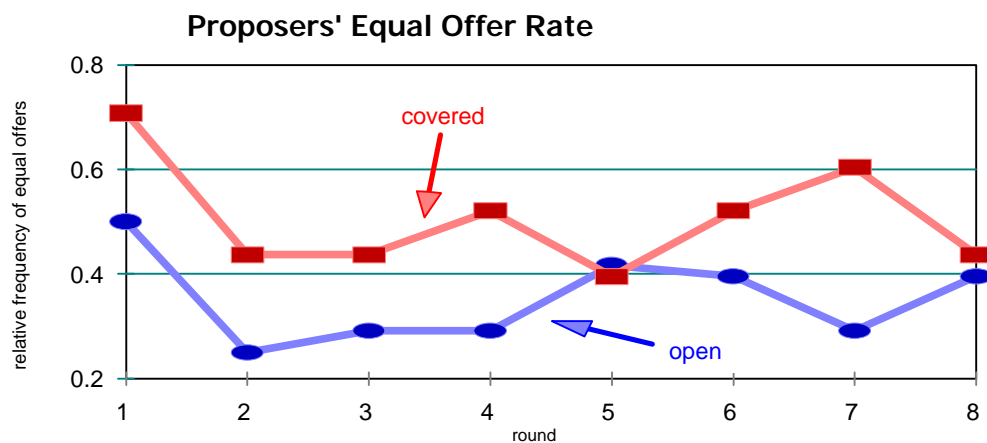


Source: Abbink, Sadrieh, and Zamir, 1999: 9.

Observation 3: *In the open treatment, responder rejection rates fall significantly from the first to the second half of the experiment. In the covered treatment, no trend can be detected.*

It is consistent with the group reputation hypothesis that the differences between the rejection rates across treatments diminish towards the end of the session: in the last two rounds, when contributing to group reputation makes little or no sense, no significant difference between the treatments is observed.

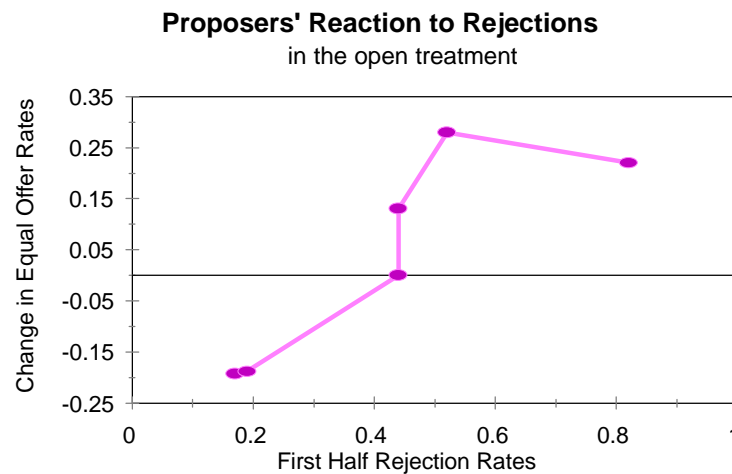
Observation 4: *The equal offer rates are significantly higher in the covered response treatment than in the open response treatment.*



Source: Abbink, Sadrieh, and Zamir, 1999: 10.

In the covered response treatment, we observe equal offers in half of the cases. This rate is about 40% higher than in the open response treatment. Aggregated over all eight rounds, the difference across treatments is significant at $p = 0.026$ (one-tailed), according to the Mann-Whitney U-test, applied to average equal offer rates in the sessions.

Observation 5: *Group reputation is effective: In the open treatment, proposers tend to switch to the equal offer significantly more often after having observed a rejection than after acceptance of unequal offers.*



Source: Abbink, Sadrieh, and Zamir, 1999: 11.

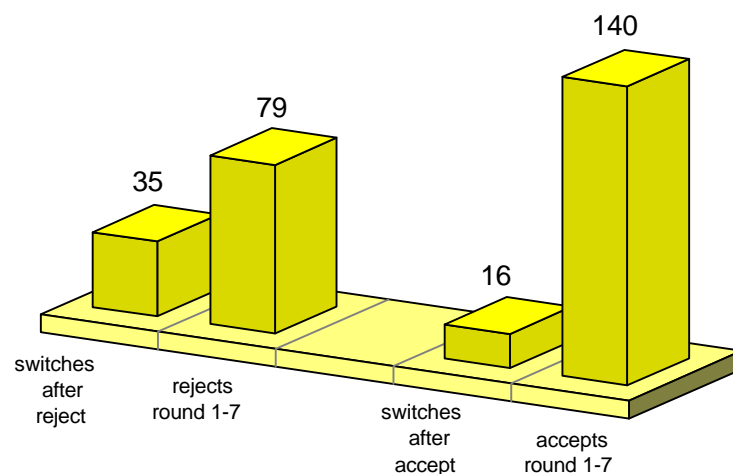
High frequencies of rejections (in the first half of the session) are strongly correlated to higher increases in equal offer rates (in the second half of the session).

Another manifestation of the efficiency of group reputation is seen in the individual patterns of switches from an unfair offer to a fair offer following a rejection or acceptance (of an unfair offer):

Proposers' Switches to the Equal Offer

switch to equal offer significantly more likely after rejection

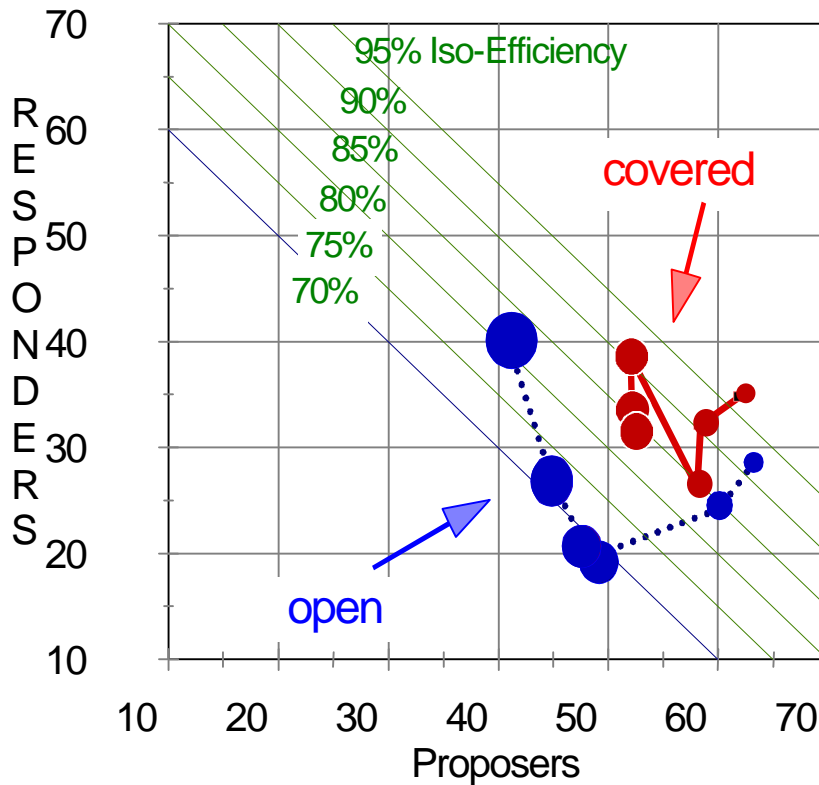
responders can effectively build up a group reputation for being "tough"



Source: Abbink, Sadrieh, and Zamir, 1999: 12.

Proposers in fact tend to react to punishment in a manner that is favorable to the responders. If responders anticipate such behavior, it can be reasonable for them to contribute to the group reputation by rejecting unequal offers.

Observation 6: *The highest responder payoffs and the highest efficiency are observed in the sessions in which rejection rates are either very low or very high.*



Comparison of Proposers' and Responders' Average Payoffs (in % of the Cake)

The size of the ● or ■ corresponds to the rejection rate of the session

Source: Abbink, Sadrieh, and Zamir, 1999: 15.

Notice that efficiency in the covered treatment is greater than in the open treatment. This is due to the lower rejection rates coinciding with higher equal offer rates.

Competitive Ultimatum Game: What is the theoretical prediction?

The last work I would like to talk about here will give me the opportunity to come back to the issue of ‘*the* theoretical prediction’ tested in an experiment. This is an experimental work in progress, involving ten co-authors (including myself⁶).

Consider an Ultimatum Game with three (potential) proposers P_1 , P_2 , and P_3 and one responder R . The three proposers sequentially propose an allocation of the cake C to the responder. First, P_1 proposes an allocation $a_1=(x_1,C-x_1)$ of C to R , where x_1 denotes the proposed payoff for the responder R and $C-x_1$ denotes the payoff desired by P_1 for himself. The game ends if R accepts the proposal of P_1 . Otherwise, if P_1 's proposal is rejected by R , P_2 proposes an allocation $a_2=(x_2,C-x_2)$ of C to R . If R accepts P_2 's proposal, the game ends. Otherwise, if R also turns down P_2 's proposal, it is P_3 's turn to propose an allocation $a_3=(x_3,C-x_3)$ of C to R . If R accepts the proposal of proposer P_3 , R receives x_3 , P_3 receives $C-x_3$, and each of the other two proposers receives 0. If R rejects all three proposals, all four players receive 0.

Our design focused on two variants of this game which differ in their informational structure. The first is the *perfect information* version in which each proposer, when called to propose, is completely informed about the (rejected) proposal(s) of the proposer(s) who preceded him. The second version is the *competitive Ultimatum Game with imperfect information*, in which none of the proposers is informed about the proposal(s) of the proposer(s) who preceded him. This means that the second and third proposer, when called to make a proposal, can only infer that the responder rejected the previous proposal(s). However, they do not know which amount was actually proposed by the previous proposer(s).

What is the theoretical prediction for this game?

Since in the standard Ultimatum Game, the prediction is the backward induction subgame perfect equilibrium, let us check this concept for our game. In fact, the perfect information game has subgame perfect equilibria. (There is no unique subgame perfect equilibrium due to the existence of a minimal monetary unit μ .) In all of these equilibria, the first proposer P_1 offers 2μ or 3μ to the responder and the responder accepts it. The reasoning is straightforward, as in standard Ultimatum Bargaining. If P_3 proposes at least μ to R , it is strictly dominant for R to accept. Anticipating this, it is strictly dominant for R to accept any offer from P_2 , of 2μ or more. Therefore, it is strictly dominant for R to accept each proposal of P_1 , offering him 3μ or more. (We just outlined the argument to establish the

⁶This work was part of a joint German-Israeli Foundation (GIF) project, and was co-authored by all members of the two laboratories involved in the experiment : Klaus Abbink, Ron Darziv, Zohar Gilula, Harel Goren, Bernd Irlenbusch, Arnon Keren, Bettina Rockenbach, Abdolkarim Sadrieh, Reinhard Selten, and Shmuel Zamir

bound for what R can get in subgame perfect equilibrium; there are also subgame perfect equilibria in which P_1 offers 0 , μ , or 2μ to the responder.)

For the version with imperfect information, a similar, but more careful, argument has to be made. To begin with, this game has no subgames. Hence, each Nash equilibrium is trivially subgame perfect. The appropriate analogue to the subgame perfect equilibrium here is the *sequential equilibrium*. In any such equilibrium, the beliefs have to be consistent with the (common knowledge) fact that R is rational and hence will not reject an offer of μ or more from P_3 . Consequently, in any sequential equilibrium (with ‘rationalizeable beliefs’), P_1 proposes $a_1=(3\mu, C-3\mu)$, and the responder accepts it.

Is that ‘the theoretical prediction’? Is that the only possible theoretical analysis?

Of course NOT! What about the competitive aspect of the situation which was completely ignored in the previous analysis? We could as well view the situation as that of three buyers (the proposers) and a seller (the responder) who holds an indivisible object which is worth C for each buyer and 0 for the seller. Any reasonable model of price competition based on supply and demand would result in price C ; namely, the responder’s payoff is C and each proposer’s payoff is 0 . This may be called the *competitive solution* of the given situation. If this does not look game theoretic enough, we can make it a formal, game theoretical solution. Let us model the situation as a *cooperative game with transferable utility*: The players are $\{R, P_1, P_2, P_3\}$ and the *characteristic function* is $v(S) = C$ if the coalition S contains R and at least one of the players P_i and $v(S)=0$ otherwise. This is known as *the gloves game* since one way to think of it is that R owns a right glove and each of P_1, P_2 and P_3 owns a left glove. Any single glove is worth zero but a pair of gloves is worth C . So the whole ‘society’ can produce C . How will it be allocated? The relevant solution concept here is the *CORE*; that is, the set of allocations that cannot be improved upon by any coalition. So we ask:

What is the CORE of this game?

Well, the CORE of this game consists of a single allocation which is $(C, 0, 0, 0)$. This is the competitive solution according to which R gets the whole cake C leaving nothing to the three P players.⁷

So, what is the theoretical prediction for this game?

Is it the backward induction solution which gives R practically zero, or the CORE which gives him the whole cake C ? Or maybe some other solution concept, for example the *Shapley value* which suggests the allocation:

$$\left(\frac{3C}{4}, \frac{C}{12}, \frac{C}{12}, \frac{C}{12} \right)$$

⁷ Another interpretation of this cooperative game, more suited to political science, is to view it as a simple (or committee) game with one *veto player*. Again, in such a game, the only point in the core is that in which the veto player has all the power.

None of these seems to have the claim for the title ‘the theoretical prediction’ although – and maybe because – all are important and respectable game theoretical solution concepts.

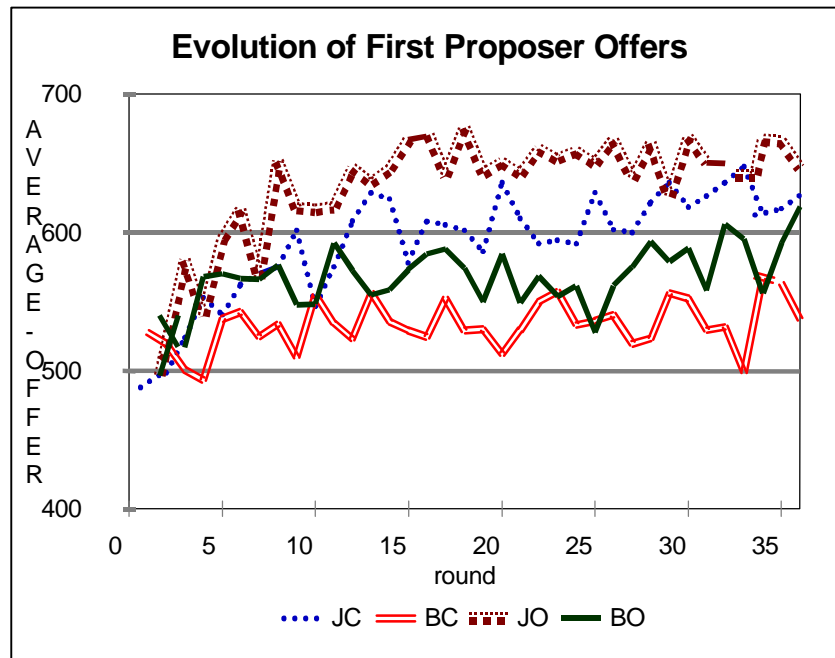
Let us first look at the main results of our experiment. Most of first and second proposals were accepted as can be seen from the following table:

Table 1: Frequency of acceptance at the three stages of the game

Accepted offer	Jerusalem covered	Bonn Covered	Jerusalem Open	Bonn open	Total
1 st proposer	448 (69.1%)	463 (71.5%)	440 (67.9%)	409 (63.1%)	1760 (67.9%)
2 nd proposer	132 (20.4%)	139 (21.4%)	140 (21.6%)	147 (22.7%)	558 (21.5%)
3 rd proposer	52 (8.0%)	32 (4.9%)	45 (6.9%)	60 (9.3%)	189 (7.3%)
none	16 (2.5%)	14 (2.2%)	23 (3.6%)	32 (4.9%)	85 (3.3%)
Total	648 (100%)	648 (100%)	648 (100%)	648 (100%)	2592 (100%)

Source: Abbink, *et al.*, 2000: 6.

For a cake of size $C = 1000$, how high were the first proposers’ offers ? This is given in the following figure showing the evolution of the average first offers:



Source: Abbink, *et al.*, 2000: 9.

So first offers are well above 50% of the cake and often above 60%. The graphs do not show a tendency of the first proposer offers to fall; on the contrary, they even rise slightly. In other words, there is no tendency to move toward the backward induction solution, but rather a (slight) upward trend, toward the competitive solution. Finally, complicating the analysis even further, we realize that as a matter of fact, from all three concepts mentioned, the Shapley value was the closest to the observed data.

The data of this experiment is still being analyzed further, but it is quite clear that we are observing neither the backward induction solution nor the competitive solution. Nor do we see convergence to either of them. We do see evidence of both effects. The competitive aspect is evident from the high offers at all levels, especially by first and second proposers. The responder is well aware of his strong position and often rejects 50% of the cake, which in a standard Ultimatum Game, is the highest conceivable offer. Nevertheless, backward induction also seems to have an effect. In 21 of 24 sessions, the average offers have a strictly declining pattern from the first to the third proposer, such that the second proposers' offers are lower than the first proposers', and the third proposers' offers are lowest. Only in three sessions, the average second proposer offer is higher than the average first proposer offer. Obviously, the first proposers' fear of their offers being rejected is greater than the second proposers' fear, which in turn is greater than the third proposers' fear.

An interesting observation is that the most frequently made offers by subjects in all three proposer positions are from the interval [451..500]. This holds true, even though the average offers made by first and second proposers are well above 500, whereas those made by third proposers are well below 500. As I said before, these 'equitable' offers from first and second proposers were often rejected. Is it now the (first and second) proposers who are appealing to the notion of fairness?

There have been other experiments of competitive bargaining situations. Both Prasnikar and Roth (1992) and Roth, Prasnikar, Okuno-Fujiwara and Zamir (1991) studied situations of Ultimatum Bargaining with nine proposers (buyers) and one responder (seller). However, the design of the bargaining was such that the predictions of the two prominent solution concepts coincided. Both the backward induction outcome and the competitive (or CORE) outcome were that the responder (seller) gets the entire cake. In our design, we confront (for the first time, to the best of my knowledge) the two solution concepts, as they provide totally opposed predictions. Prasnikar and Roth, whose experimental results clearly confirmed the (extreme) game theoretically predicted outcome, concluded that in their game, unlike in the standard Ultimatum Game, the strategic aspects are prominent (compared to other aspects of the situation such as 'fairness', 'emotions' etc.). I would like to rephrase this conclusion by saying that in their game, there was a *clear game theoretical prediction*; all relevant solution concepts yield the same outcome. This is a situation which *we know how to analyze with game theory*. This is not the case for our competitive Ultimatum Game, where there is *no clear game theoretical prediction*. This is so because, whether we like it or not, we do not know the right theoretical model for this game. Before running the experiment, we noticed that there are several ways to analyze the game, none of which is satisfactory, because each model captures only one aspect of the situation. The experiment confirmed this by showing that neither of the two predictions alone could explain the data, though there seems to be evidence supporting both and, rather surprisingly, the Shapley value is the best in approaching the data. This may indicate that an appropriate game theoretical model might include ideas and elements from both

cooperative and non cooperative game theory. Developing this theory may be a difficult and prolonged venture, if at all possible, but until we do so, we cannot say that the results are not in agreement with the game theoretical predictions, since there is no game theoretical prediction.

And what about the standard Ultimatum Game?

I argue that the situation is very similar to that described above, except that, for some reason, it is widely believed that *we know* the game theoretical prediction in the Ultimatum Game, mainly because we know only one candidate.⁸ I think that the issue is not the rationality of the players but rather our modeling power. Having observed a fair bit of evidence attesting to the rationality of players, even in the Ultimatum Game, maybe the main problem is that we still do not have a good model for this situation; a model that will have a richer interpretation of the notion of rationality. Such a model should enhance the descriptive power of game theory without having to assume, ad hoc, that players are occasionally irrational – whenever we fail to predict their behavior with simplistic solution concepts.

References

Abbink, Klaus, Gary E Bolton, Abdolkarim Sadrieh, and Fang-Fang Tang (1996). “Adaptive Learning versus Punishment in Ultimatum Bargaining.” SFB Discussion Paper B-381, University of Bonn.

Abbink, K., R. Darziv, Z. Gilula, H. Goren, B. Irlenbusch, A. Keren, B. Rockenbach, A. Sadrieh, R. Selten, and S. Zamir (2000). “The Competitive Ultimatum Game.” Working paper.

Abbink, Klaus, and Abdolkarim Sadrieh (1995). “RatImage - Research Assistance Toolbox for Computer-Aided Human Behavior Experiments.” SFB Discussion Paper B-325, University of Bonn.

Abbink, Klaus, and Abdolkarim Sadrieh and Shmuel Zamir (1999). “Fairness and Public Good Aspects of Punishment Behavior.” The Hebrew University, Center for Rationality and Interactive Decision Theory, DP #191 (Previously titled: “The Covered Response Ultimatum Game”).

Aumann, Robert J. (1997). “Rationality and Bounded Rationality.” *Games and Economic Behavior* 21, 2-14.

Bolton, Gary E (1991). “A Comparative Model of Bargaining: Theory and Evidence.” *American Economic Review* 81(5), 1096-1135.

⁸ Actually, why not take the cooperative model of the ‘Unanimity Game’, in which: every allocation is in the CORE (analogous to every allocation is a NE) and the ‘minimal CORE’ coincides with the Shapley Value and the Nucleolus, namely (50, 50)? This solution is very much in line with the data. That is, perhaps the prominent issue (apart from rationality) is cooperation.

Bolton, Gary E, and Axel Ockenfels (2000). "ERC - A Theory of Equity, Reciprocity and Competition." *American Economic Review* 90(1), 166-193.

Bolton, Gary E, and Rami Zwick (1995). "Anonymity versus Punishment in Ultimatum Bargaining." *Games and Economic Behavior* 10(1), 95-121.

Brandts, Jordi, and Charles A. Holt (1992). "An Experimental Test of Equilibrium Dominance in Signaling Games." *American Economic Review* 82(5), 1350-1365.

Camerer, Colin, and Teck-Hua Ho (1999). "Experience-weighted Attraction Learning in games: A Unifying Approach." Working paper, California Institute of Technology.

Camerer, Colin, and Richard H. Thaler (1995). "Anomalies - Ultimatums, Dictators and Manners." *Journal of Economic Perspectives* 9(2), 209-219.

Camerer, Colin, and Keith Weigelt (1988). "Experimental Tests of a Sequential Equilibrium Reputation Model." *Econometrica* 56(1), 1-36.

Cameron, Lisa (1995). "Raising the stakes in the Ultimatum Game: Experimental Evidence from Indonesia." Discussion Paper, Princeton University.

Erev, Ido, and Alvin E. Roth (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88(4), 848-882.

Fehr, Ernst, and Klaus M. Schmidt (1999). "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114(3), 817-868.

Forsythe, Robert, Joel L. Horowitz, N. E. Savin and Martin Sefton (1994). "Fairness in simple bargaining experiments." *Games and Economic Behavior* 6(3), 347-369.

Gale, John, Kenneth G Binmore, and Larry Samuelson (1995). "Learning to be Imperfect: The Ultimatum Game." *Games and Economic Behavior* 8(1), 56-90.

Gneezy, Uri, and Avraham Stoler (1998). "Why Punish? Norms and Revenge in an Experimental Game." Working paper, University of Haifa and University of Tel Aviv.

Güth, Werner (1995). "On Ultimatum Bargaining Experiments - A Personal Review." *Journal of Economic Behavior and Organization* 27(3), 329-344.

Güth, Werner, R. Schmittberger, and B. Schwarze (1982). "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization* 3(4), 367-388.

Güth, Werner, and Reinhard Tietz (1990). "Ultimatum Bargaining Behavior : A survey and and comparison of experimental results." *Journal of Economic Psychology* 11(3), 417-449.

Hoffman, Elizabeth, Kevin McCabe, and Vernon Smith (1996). "On Expectations and the Monetary Stakes in Ultimatum Games." *International Journal of Game Theory* 25(3), 289-301.

Jacobsen, Eva, and Abdolkarim Sadrieh (1996). "Experimental Proof for the Motivational Importance of Reciprocity." SFB Discussion Paper B-386, University of Bonn.

Jung, Yun Joo, John H. Kagel, and Dan Levin (1994). "On the existence of predatory pricing: an experimental study of reputation and entry deterrence in the chain-store game" *RAND Journal of Economics* 25(1), 72-93.

Kagel, John, Chng Kim and Donald Moser (1996). "Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs." *Games and Economic Behavior* 13(1), 100-110.

McKelvey, Richard D., and Thomas R. Palfrey (1992). "An Experimental Study of the Centipede Game." *Econometrica* 60(4), 803-836.

Ochs, Jack and Alvin E. Roth (1989). "An experimental study of sequential bargaining." *American Economic Review* 79(3), 355-384.

Prasnikar, Vesna and Alvin Roth (1992). "Considerations of fairness and strategy: Experimental data from sequential games." *Quarterly Journal of Economics* 107(3), 865-888.

Rabin, Matthew (1993). "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83(5), 1281-1302.

Roth, Alvin E. (1995). "Bargaining Experiments." In John Kagel and Alvin E. Roth (eds): *Handbook of Experimental Economics*. Princeton University Press.

Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir (1991). "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo." *American Economic Review* 81(5), 1068-1095.

Roth, Alvin E., and Ido Erev (1995). "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in Intermediate Term." *Games and Economic Behavior*, 8(1), 164-212.

Roth, A. E., and Schoumaker, F. (1983). "Expectations and Reputations in Bargaining: An Experimental Study", *American Economic Review* 73(3), 362-372.

Thaler, Richard H. (1988). "The Ultimatum Game." *Journal of Economic Perspectives* 2(4), 195-206.

Van Damme, Eric (1998). "On the State of the Art in Game Theory: An Interview with Robert Aumann." *Games and Economic Behavior* 24, 181-210.

Winter, Eyal, and Shmuel Zamir (1997). "An Experiment with Ultimatum Bargaining in a Changing Environment." The Hebrew University, Center for Rationality and Interactive Decision Theory, DP # 159.