האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

AN EXPERIMENT WITH ULTIMATUM BARGAINING IN A CHANGING ENVIRONMENT

by
EYAL WINTER
and
SHMUEL ZAMIR

Discussion Paper # 159 December 1997

מרכז לחקר הרציונליות
וההחלטות האינטראקטיביות
CENTER FOR RATIONALITY
AND INTERACTIVE DECISION THEORY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@vms.huji.ac.il
WWW: http://www.ma.huji.ac.il/~ranb

An Experiment with Ultimatum Bargaining in a Changing Environment*

Eyal Winter[1] and Shmuel Zamir[2]

Abstract

We have obtained experimental results on the ultimatum bargaining game that support an evolutionary explanation for subjects' behavior in the game. In these experiments we have created environments in which subjects interact with each other in addition to interacting with virtual players, i.e. computer programs with pre-specified strategies. Some of these virtual players were designed to play the equitable allocation, while others exhibited behavior closer to the subgame-perfect equilibrium, in which the proposer's share is much larger than that of the responder. We have observed significant differences in the behavior of real subjects depending on the type of "mutants" (virtual players) that were present in their environment.

* We thank Yishai Mor and Arnon Keren for their assistance with programming the experiments and with the statistical analysis. We are also grateful to the German Israeli Foundation for its financial support.

1 The Hebrew University of Jerusalem, Center for Rationality, and Washington University, Dept. of Economics.

2 The Hebrew University of Jerusalem, Center for Rationality, and Dept. of Statistics.

1. Introduction

A basic and elementary rationality assumption asserts that a person will prefer to receive any amount of money over receiving nothing. Suppose that person 1 is assigned the task of dividing a given stack of money between himself and person 2 by making a take-it-or-leave-it-offer to the latter person. If person 2 rejects the offer, none of the players receive a payoff. Following the above assumption, person 1 expects person 2 to accept any offer that yields a positive share of the stack, and therefore should offer the smallest possible bid, leaving virtually the whole stack for himself. This argument is the basic rationale underlying the notion of subgame-perfect equilibrium, which is a central notion in economic modelling. The very simple game described above is known in the literature as the *ultimatum bargaining game (UBG).* Player 1 can use his ultimatum power to reduce player 2's payoff to virtually zero.

Experimental results in a variety of designs and setups have shown that human subjects' behavior differs considerably from the argument presented above (see Roth (1995) for a survey of previous experiments on the UBG and some of its variants). Most offers fall slightly short of 50%, and offers that deviate substantially from an equal division are typically rejected.[3] These results have often been interpreted as an intriguing discrepancy between experimental results and game-theoretic predictions. The purpose of this paper is to report experimental results that, we believe, offer an explanation of the difference between real subject behavior in UBG and the subgame-perfect equilibrium solution of this game. These results also suggest that there is no real contradiction between the observed behavior in the UBG and the rationality postulates of game theory.

In trying to explain the apparently irrational behavior of subjects in the UBG one has to address two questions. First, why do proposers tend to offer relatively large shares, and second: why do responders reject low offers. We will discuss these two questions separately by studying the way by which subject behavior in the UBG responds to changes in the environment. Our approach is thus evolutionary in nature, and views the emergence of standard of behavior in the UBG as a process of mutual adaptation.

To study the effect of the environment we depart from the conventional setup of UBG experiments, extending it by an additional experimental tool. Our population of players includes both *real* subjects and *virtual* players; the latter are computer programs that play the roles of both proposers and responders by using fixed strategies specified at the beginning of the experiment. In each experiment the UBG was played over and over again for a large number of rounds. At the beginning of each round subjects were matched randomly either to another real player or to a virtual player. None of the real subjects knew

3 Gueth, Schmittberger & Schwarze (1982), who were the first to experiment with UBGs, have even obtained a modal offer of exactly 50%. Quantitatively similar results were later reported by many others. For a comprehensive survey of this and related experiments see Bolton and Zwick (1993).

about the presence of virtual players.[4] From their point of view, they were playing a regular UBG with a conventional design.

The objective of this design was to explore the way real subject behavior changes as a function of the *type* of virtual player in the experiment. One of the main questions this paper addresses is what elements determine individual behavior in the UBG. 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", 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.[5] 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%.

We will show that the presence and identity of virtual players dramatically affect real subject behavior in the UBG.

Section 2 presents a formal description of the UBG and an account of its game theoretic solutions. Section 3 describes the experimental design, while Section 4 presents the results. We defer most of the discussion to Section 5.

2. The Ultimatum Bargaining Game

Consider two players, 1 and 2, who have to share a cake of size K, where K is an integer number, according to the following rules which we call the UBG rule. Player 1, the proposer, has to make an offer to player 2, the responder. A proposal is simply a number from {1,2,...,K} which indicates the share of the responder. Player 2, upon hearing the offer, has to decide whether to accept the offer or reject it. Player 2's strategy is thus a function from {1,2...,K} to the set {Yes, No} which specifies the response. Let s_1 be a strategy for player 1, and s_2 a strategy for player 2. The payoff function of the UBG assigns the agreed payoff for the parties in case of acceptance, and zero in case of rejection. Formally, denoting by $A(s_2)$ the set of acceptable shares for player 2 according to s_2, we have:
$h(s_1,s_2) = (K-s_1, s_1)$ if $s_1 \in A(s_2)$ and $h(s_1,s_2) = (0,0)$ if $s_1 \notin A(s_2)$ (first coordinate is the payoff to the proposer and the second to the responder).

4 This is said for the main design in which subjects were informed about the presence of virtual players and about the objective of the experiment at the end of the project. We also run an alternative design in which subjects were told about the presence of virtual players in the instructions; see section 5.4.

5 It may seem strange to call such a responder "tough", but our terminology refers to the outcome rather than to the responder.

We analyze the game through the notion of Nash Equilibrium: a pair of strategies (s_1^*, s_2^*) is Nash Equilibrium if each of the two strategies (s_1^*, s_2^*) is a *best reply* to the other; that is, if we denote $h(s_1,s_2) = (h_1(s_1,s_2), h_2(s_1,s_2))$ then:

$h_1(s_1^*, s_2^*) \geq h_1(s_1, s_2^*)$ for all s_1

$h_2(s_1^*, s_2^*) \geq h_2(s_1^*, s_2)$ for all s_2

$h(s_1^*, s_2^*) = (h_1(s_1^*, s_2^*), h_2(s_1^*, s_2^*))$ is the equilibrium payoff corresponding to the equilibrium (s_1^*, s_2^*). The UBG has exactly K Nash equilibria. Each equilibrium sustains the share (i,K-i) (i= 1,2,...K) by having player 1 proposing this share and player 2 accepting at least this share and rejecting anything less. Among this set of Nash equilibria, two are *subgame-perfect* and can be worked out by simple backward induction.[6] The first corresponds to an allocation in which the proposer gets everything, i.e., (K,0) and the second is when the responder is granted only one unit. i.e., (K-1,1).[7] These last two (extreme) equilibria were often associated in the literature on the UBG as *the* game-theoretic prediction of the UBG. This is, of course, a misleading assertion: these equilibria are part of a larger set of Nash equilibria. It is true that the argument behind the notion of subgame-perfection is transparent when applied to the UBG because of the simplicity of this game, but it is wrong to claim that from a theoretical point of view other Nash equilibria are irrelevant or inconceivable. We will come back to this point later on, when discussing other works that emphasize the relevance of other Nash equilibria in the UBG. We now turn to the description of the exact design of the experiments and summarize their results.

3. The Experimental Design.

3.1 Virtual Players and Matching

Our design consists of two groups of sessions (experiments), which differ in terms of the revealed information concerning the presence of virtual players. The first group consists of 8 experiments, all involving a UBG in which a cake of 100 points was to be divided. All subjects were students at the Hebrew University of Jerusalem in a variety of academic stages and disciplines (most of them were undergraduate social sciences students). The experiments were all computerized and were conducted in the newly established experimental laboratory *Ratiolab,* at the Center for Rationality in the Hebrew University. The computer program heavily used the Ratimage software, developed in the University of Bonn, by Abbink & Sadrieh (1995). Before getting to the exact design of each session we will explain the nature of virtual players in the experiments. A virtual proposer is a

6 In this context a subgame-perfect equilibrium is a Nash equilibrium in which every positive offer made by the proposer is accepted by the responders (i.e. both on and off the equilibrium path).

7 In the continuous version of the UBG only the first allocation is sustainable by a subgame perfect equilibrium.

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 of 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 will 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.

We can now describe the design of a typical session with virtual players. In each session a different group of subjects was received in the laboratory. Before commencing, a lottery determined the role of each subject (proposer or responder).[8] This role was fixed through-out the session. The subjects played the UBG 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. 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 through-out the session, but in two sessions (with virtuals $P_{13,16}$ and $R_{13,16}$) we increased the population of virtual players gradually. In all of sessions real subjects *did not know* that they were playing virtual players. They were not informed about virtual players, and believed that they were only matched among themselves.[9] We will discuss the aspect of this (admittedly unconventional) approach later. In addition to 6 sessions with virtual players (two sessions of different sized groups for each type) we ran two sessions with no virtual players at all.

8 The random choice of roles was used in order to guarantee that all subjects have the same ex-ante earning opportunities.

9 This was confirmed by a questionnaire that the subjects filled at the end of the experiment.

Table 1 describes the full specification of each of the 8 experiments that we have conducted.

Table 1: The experimetnal design of the 8 sessions.

Session	No of Rounds	No of Real Players	Distribution of Virtual Players	Type of Virtual Players
1	50	12	no virutals	-
2	50	20	no virtuals	-
3	50	12	8 all through the session	4 are $P_{23,26}$ 4 are $R_{23,26}$
4	50	20	14 all through the session	7 are $P_{23,26}$ 7 are $R_{23,26}$
5	70	12	gradual: 10 rounds with 2, next 10 with 4, next 10 with 6, and the rest with 8	at each round half are $P_{13,16}$ and half are $R_{13,16}$
6	70	20	gradual: 10 rounds with 4, next 10 with 6, next 10 with 10, and the rest with14	at each round half are $P_{13,16}$ and half are $R_{13,16}$
7	50	12	8 all through the session	4 are $P_{23,26}$ 4 are $R_{23,26}$
8	50	18	12 all through the session	6 are $P_{46,49}$ 6 are $R_{46,49}$

In addition to the 8 sessions reported above we have run a separate set of 6 sessions in which we have repeated twice the 3 small group sessions with virtual players (using new subject pools). In this design, and in contrast to the original one, 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 a real player. However, they were not told anything about the probability of this event or about the nature of these computer programs. We will discuss the results of this design at the discussion part of the paper.

3.2 Payoffs

Each player received payment according to the number of points accumulated in the course of the session, computed at a fixed exchange rate of NIS 1 per 100 points.[10] In addition, each player received a fixed payment NIS 10 for participating in the experiment.

4. Results

4.1 The distribution of offers

The most unambiguous result of this experiment is perhaps the effect of the presence of virtual players on the offers made by real players. When exploring this effect one has to make two types of comparisons. The first involves comparing different sessions at the same stage. The other concerns the evolution of the behavior in the same session over time.

Figures 1.1 to 1.8 show the distribution of offers made by *real* players in the first and the last 10 rounds. (In the numbering of the figures, we adopted the convention that the second digit is the number of the session, as is defined in table 1). This information is summarized in Table 2

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

Session	Total	First 10 Rounds	Last 10 Rounds
1	40 , 39.47 , 7.06	50 , 38.80 , 13.37	40 , 40.45 , 2.81
2	40 , 40.95 , 6.81	40 , 43.13 , 8.26	40 , 42.13 , 4.21
3	30 , 36.89 , 11.92	50 , 47.35 , 13.38	30 , 32.17 , 8.78
4	30 , 35.48 , 7.76	40 , 38.11 , 9.48	30 , 32.71 , 5.71
5	40 , 39.44 , 8.45	50 , 45.50 , 8.42	30 , 35.07 , 8.00
6	20 , 34.84 , 12.26	50 , 40.20 , 17.33	20 , 28.94 , 10.43
7	50 , 45.20 , 9.67	50 , 42.85 , 12.71	50 , 46.53 , 7.20
8	50 , 48.88 , 3.17	50 , 47.27 , 4.84	50 , 49.00 , 3.47

Without virtual players, the distribution mode either shifts around 40 - 50 points or remains at 40. When introducing moderately tough virtuals ($P_{23,26}$ and $R_{23,26}$) the mode drops to 30 points and with extremely tough virtuals ($P_{13,16}$ and $R_{13,16}$) it sinks to 20, in spite of the fact that virtuals 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

10 At the time of the experiment the exchange rate was around NIS 3 per one US dollar.

across all sessions is that the distribution of offers in the first 10 rounds are 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.

Figure 1.1
Relative Distribution of Offers by Real Players
no virtual players (12 players)

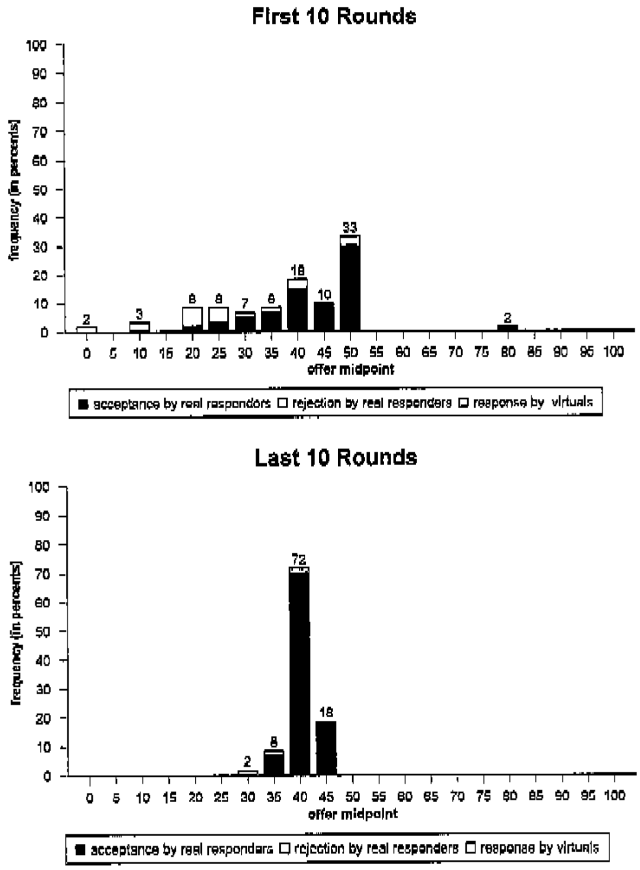

Figure 1.2
Relative Distribution of Offers by Real Players
no virtual players (20 players)

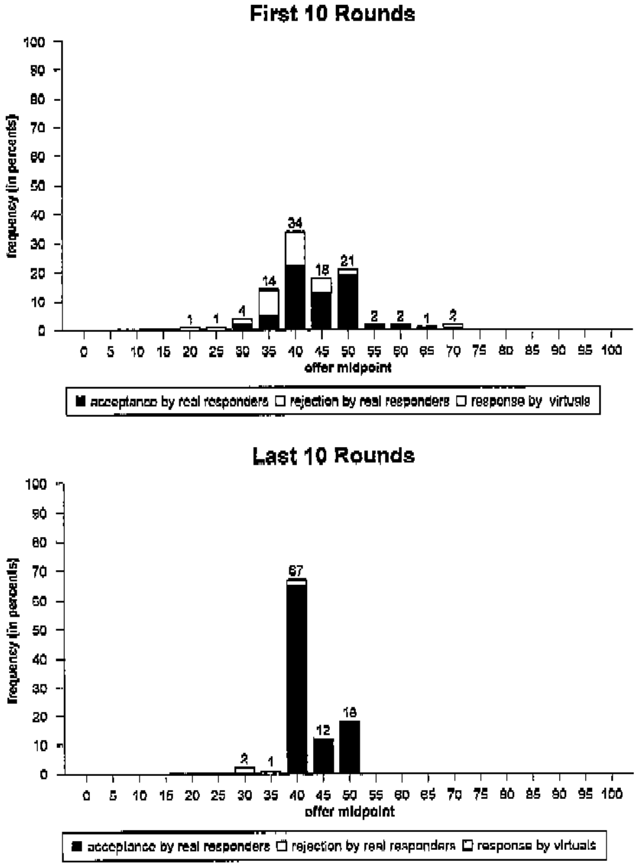

Figure 1.3
Relative Distribution of Offers by Real Players
virtual offer range = 23-26 (12 players)

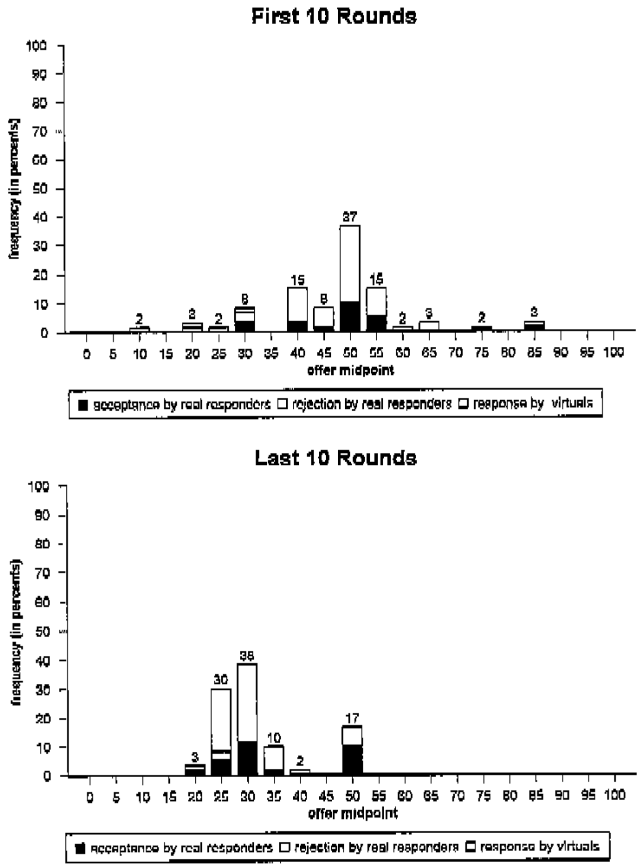

Figure 1.4
Relative Distribution of Offers by Real Players
virtual offer range = 23-26 (20 players)

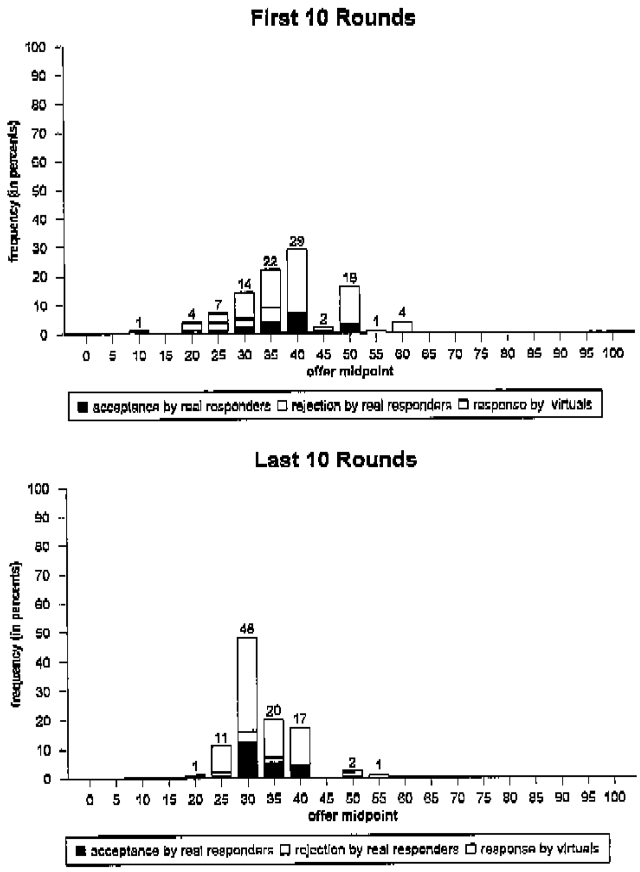

Figure 1.5
Relative Distribution of Offers by Real Players
virtual offer range = 13-16 (gradual, 12 players)

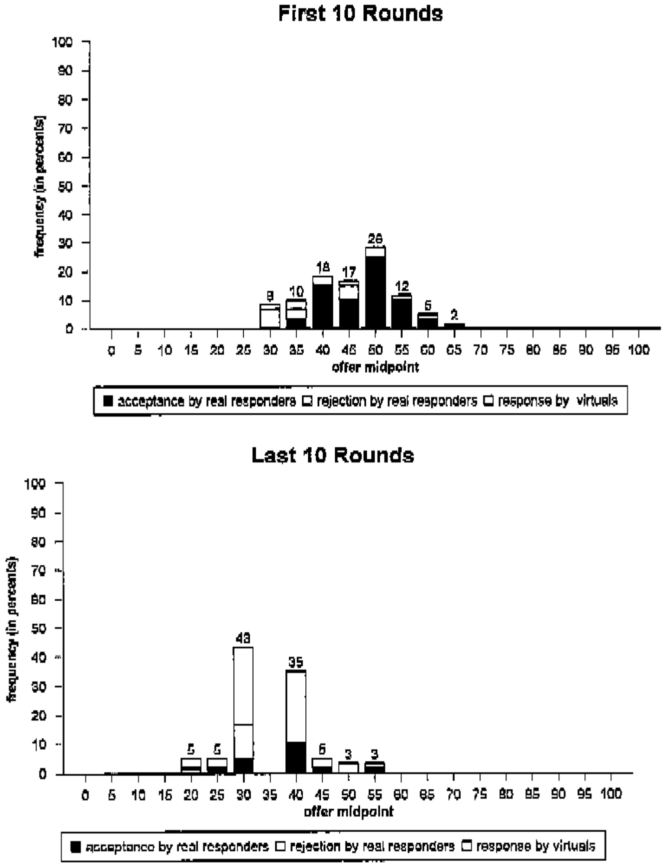

Figure 1.6
Relative Distribution of Offers by Real Players
virtual offer range = 13-16 (gradual, 20 players)

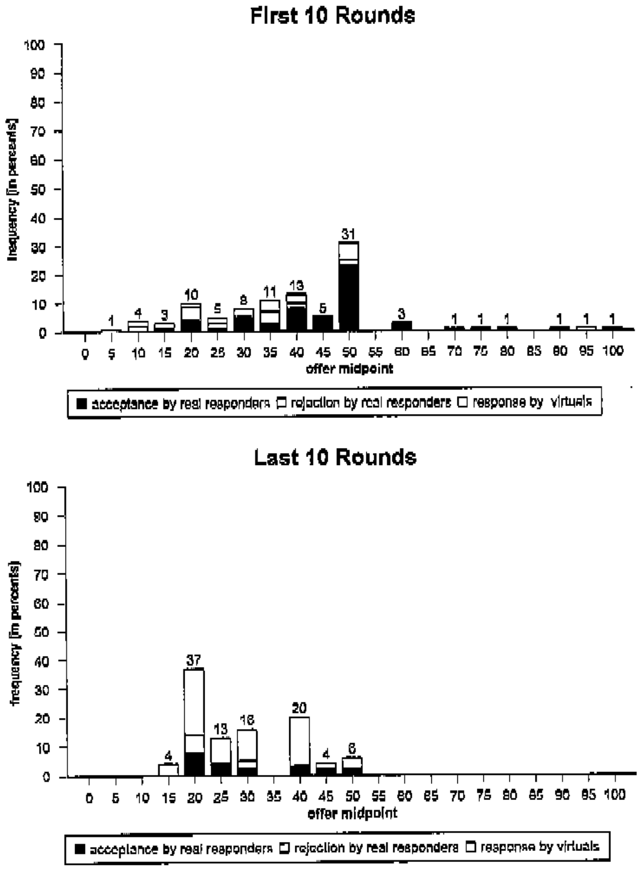

Figure 1.7
Relative Distribution of Offers by Real Players
virtual offer range = 46-49 (12 players)

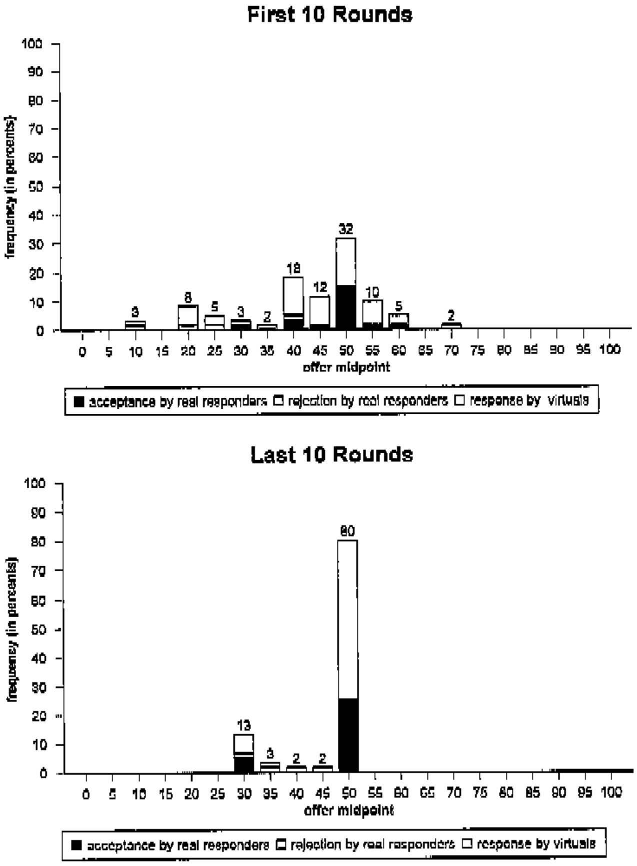

Figure 1.8
Relative Distribution of Offers by Real Players
virtual offer range = 46-49 (18 players)

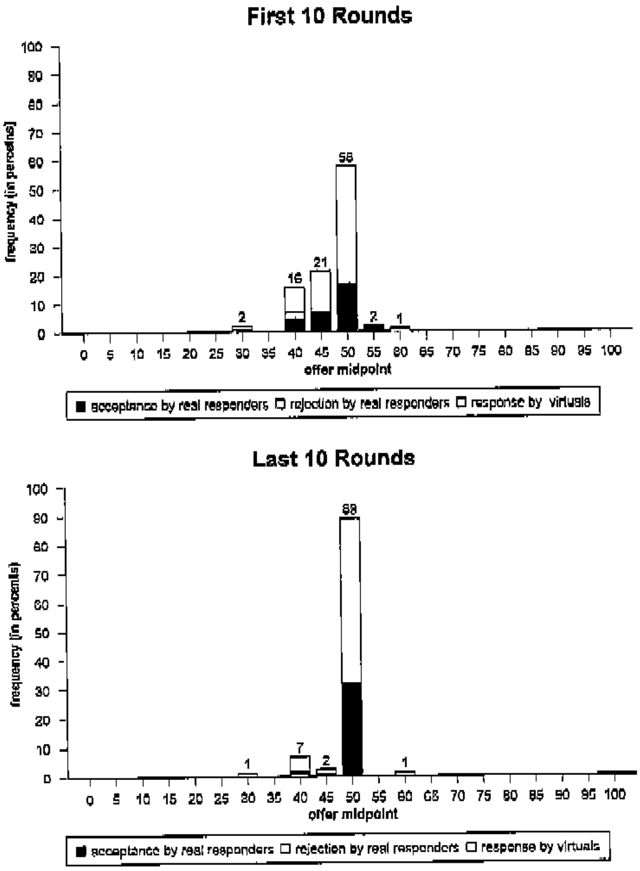

To further illustrate the effect of virtual players on the distribution of offers we compare offers by computing the probability that a randomly sampled offer from one group exceeds a randomly sampled offer from another group (table 3). Apart from the comparison of [23-26] against [13,16], which is distorted by the fact that in the latter group virtual players were introduced gradually, the other comparisons fit the intuition. Offers made in the environment with fair virtuals are higher than in the one without virtuals. And in these two environments offers are higher than in those with tough virtuals.

Table 3: Probability that an offer sampled randomly from experimental condition A will be greater than an offer sampled randomly from experimental condition B

		experimental condition B		
		13 -16	23 - 26	no virtuals
experimental condition A	13 -16	-		
	23 - 26	0.410	-	
	no virtuals	0.545	0.640	-
	46 - 49	0.768	0.791	0.789

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 the UBG), such offers are not the result of proposers being concerned with equality or of focal point of 50% having any special attraction. 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 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" virtuals 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.

Does the change in offer patterns occur instantly or is it a gradual process? Figures 2.1 to 2.8 show the way the modes and average offers evolve in time in each of the environments.[11]

11 In Figures 2.1 - 2.8 and 4.1 - 4.8 the points, 1,2,... on the horizontal axis represents a segment of 10 rounds (i.e., 1 represents rounds 1-10, 2 represents rounds 11-20, etc.)

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

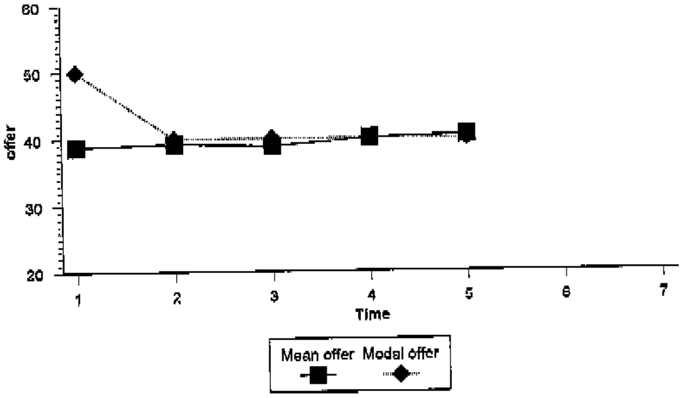

Figure 2.2
Mean and Mode of Offers by Real Proposers as a Function of Time
no virtual players (20 players)

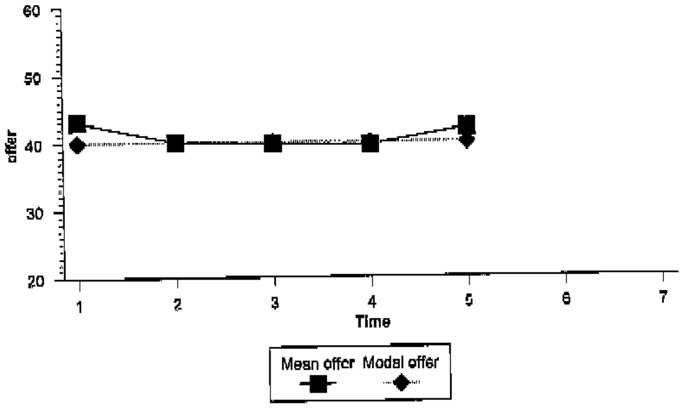

Figure 2.3

Mean and Mode of Offers by Real Players as a Function of Time

virtual offer range = 23-26 (12 players)

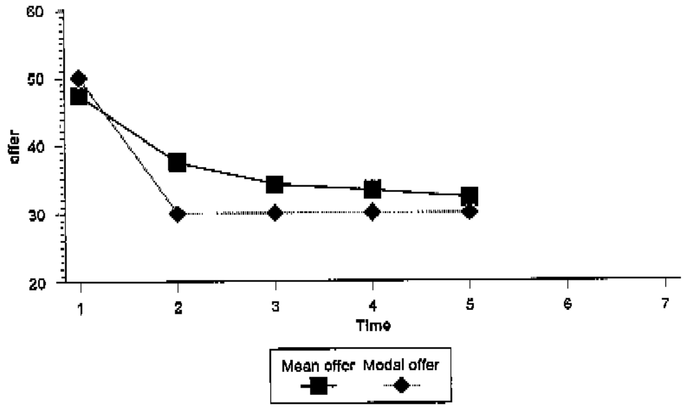

Figure 2.4

Mean and Mode of Offers by Real Proposers as a Function of Time

virtual offer range = 23-26 (20 players)

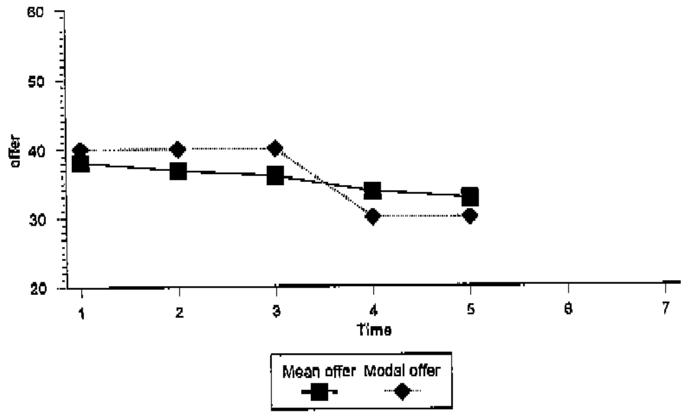

Figure 2.5
Mean and Mode of Offers by Real Players as a Function of Time
virtual offer range = 13-16 (12 players)

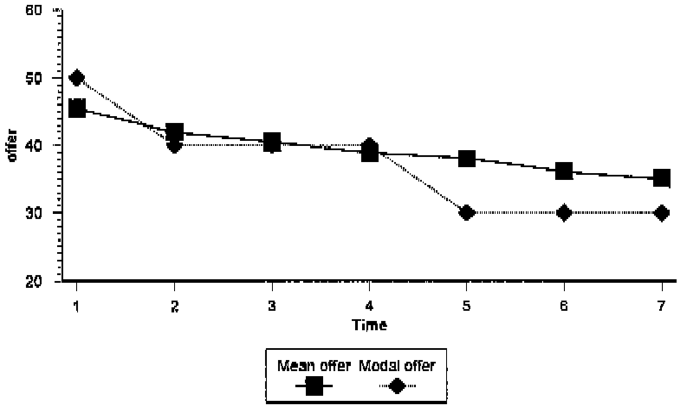

Figure 2.6
Mean and Mode of Offers by Real Proposers as a Function of Time
virtual offer range = 13-16 (20 players)

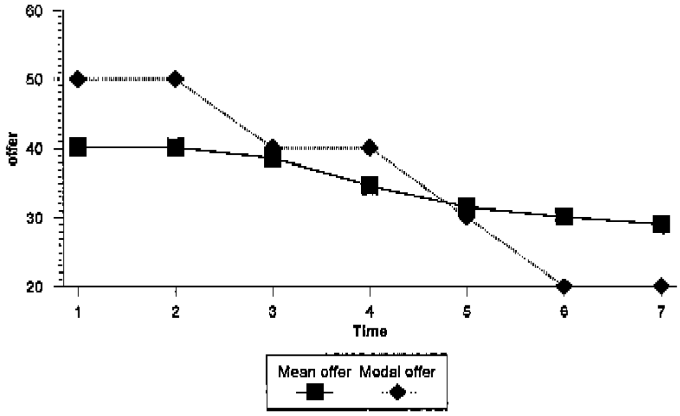

It is apparent from these figures that the environment changes gradually. With tough virtuals these indicators constantly decrease and with fair virtuals they increase gradually.

We have argued earlier that in the long run proposers submit offers that are best response with respect to the environment in which they "live". We check to what extent the proposers are "expected utility maximizers". To do that we take the empirical frequencies of responses' as estimates of the probabilities for a certain proposal to be accepted.[12] This enables us to plot the "expected" return for offers in each environment. For each environment we partition the range of offers into intervals of 5 points each, and disregard those intervals with less than 10 offers. The expected return for each interval T, is then defined as: $f_T*(100-m_T)$, where f_T is the proportion of accepted offers out of all offers within the interval T, and m_T is the midpoint of the interval T. This is done for the 8 sessions in Figures 3.1 to 3.8.

These figures show that the mode offers are strikingly close to the maximizers of the expected return, which means that (real) proposers typically match their offers to the response patterns of their environment. Note that the expected profit is computed with respect to data aggregated over all periods. Thus, even with the typical noise due to learning at the early periods of a session, players' offers respond to their environments quite accurately.

12 For Technical reasons, some of the virtual responses were not recorded. In such cases we calculated the acceptance probability using virtual responders' characteristics.

Figure 3.3
Expected Returns for Offers Made by Real Players
based on virtual and real responses
virtual offer range = 23-26 (12 players)

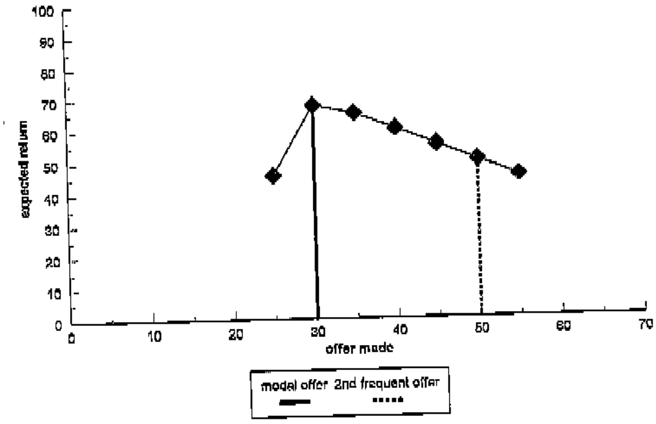

Figure 3.4
Expected Returns for Offers Made by Real Players
based on virtual and real responses
virtual offer range = 23-26 (20 players)

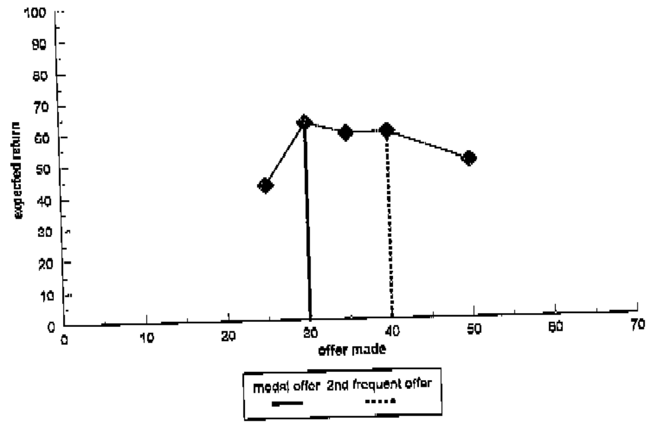

Figure 3.5
Expected Returns for Offers Made by Real Players
based on virtual and real responses
virtual offer range = 13-16 (12 players)

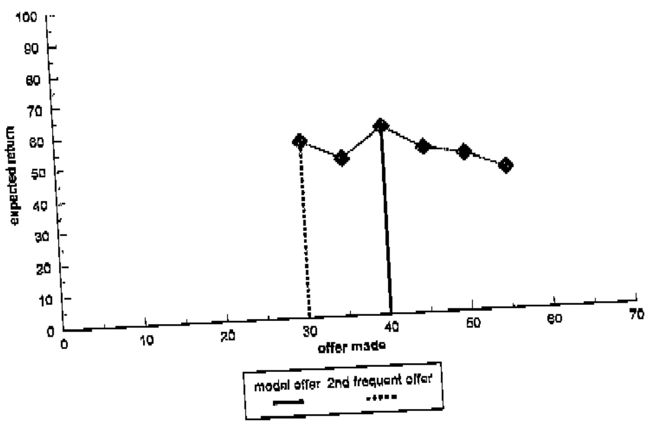

Figure 3.6
Expected Returns for Offers Made by Real Players
based on virtual and real responses
virtual offer range = 13-16 (20 players)

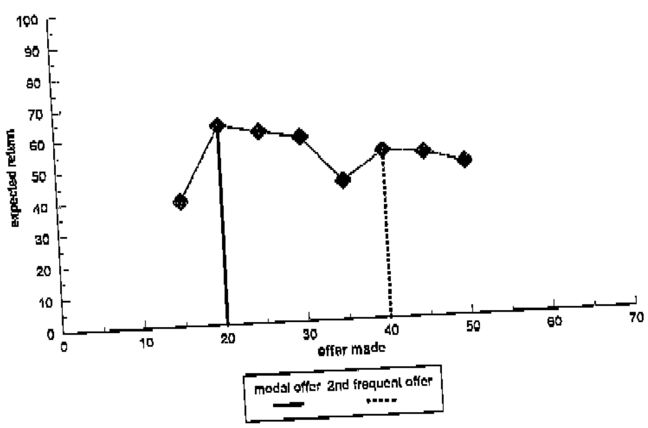

Figure 3.7
Expected Returns for Offers Made by Real Players
based on virtual and real responses
virtual offer range = 46-49 (12 players)

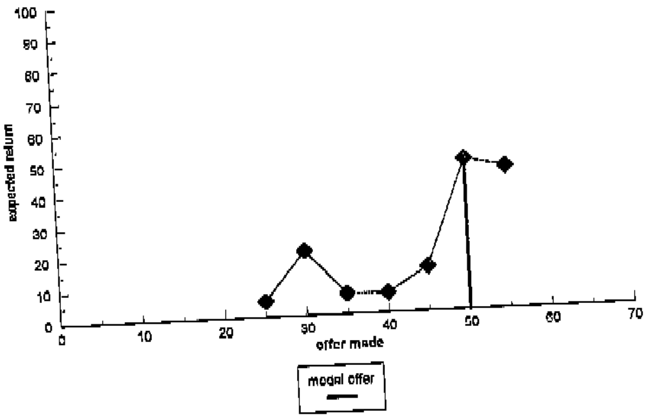

Figure 3.8
Expected Returns for Offers Made by Real Players
based on virtual and real responses
virtual offer range = 46-49 (18 players)

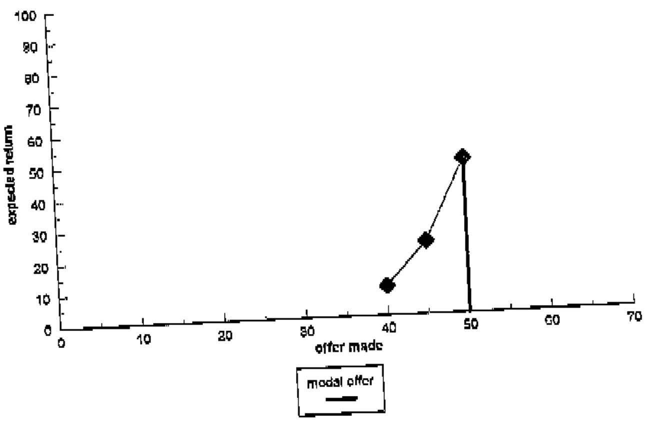

4.2 Responses

Interpreting responders' behavior in the UBG has always been the trickiest part of any analysis of experimental results of the UBG. We have seen that proposers perform part of the job of playing a Nash equilibrium pretty well, by best responding to their environment. But how closely do responders adhere to equilibrium guide lines? For a Nash equilibrium to be played it is necessary that no positive offer be rejected. The Nash equilibrium solution concept makes no prediction whatsoever about what the proposals should be, but it has an unambiguous prediction about what the responses should be. Given that proposers submit offers that are their best response to the environment, the frequency of rejections by responders is a good estimate of how far we are from a Nash equilibrium. The patterns of responses are documented using two methods. First, the histograms of offers (Figures 1.1 to 1.8) include the proportion of acceptances and rejections by *real* players for offers made by real players. Secondly, Figures 4.1 to 4.8 plot the rejection rates as a function of time, across all offers by real and virtual players (the circle dots) as well as for offers within the vicinity of the modal offer (± 4 points from the mode), which we interpret as the "equilibrium" offer. This information is summarized in Table 4.

Table 4: Rate of rejection by real players to offers made by both real and virtual platers.

Session	Total	First 10 Rounds	Last 10 Rounds
1	0.19	0.27	0.05
2	0.24	0.33	0.05
3	0.37	0.48	0.23
4	0.55	0.56	0.53
5	0.39	0.27	0.47
6	0.61	0.42	0.70
7	0.15	0.15	0.10
8	0.07	0.08	0.03

Figure 4.1

Rate of Rejection by Real Responders as a Function of Time

no virtual players (12 players)

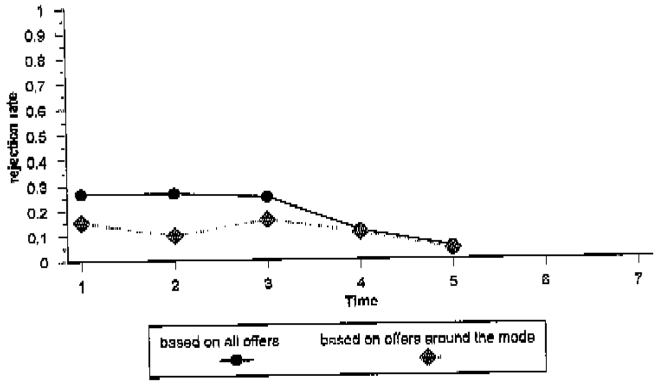

Figure 4.2

Rate of Rejection by Real Responders as a Function of Time

no virtual players (20 players)

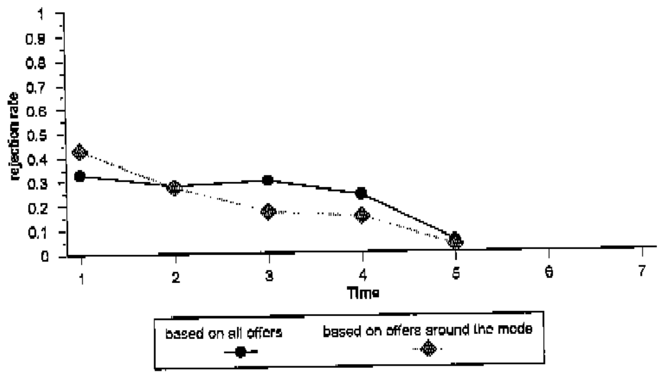

Figure 4.3
Rate of Rejection by Real Responders as a Function of Time
virtual offer range = 23-26 (12 players)

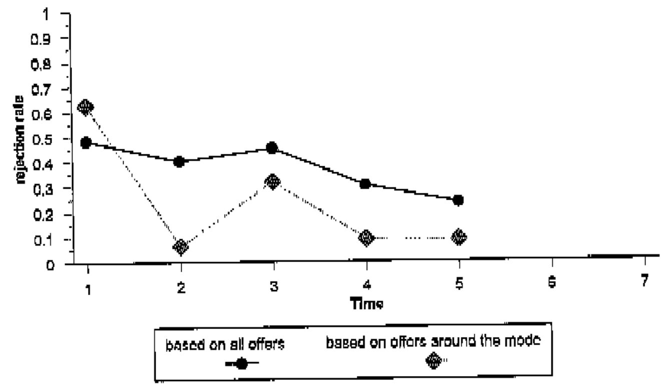

Figure 4.4
Rate of Rejection by Real Responders as a Function of Time
virtual offer range = 23-26 (20 players)

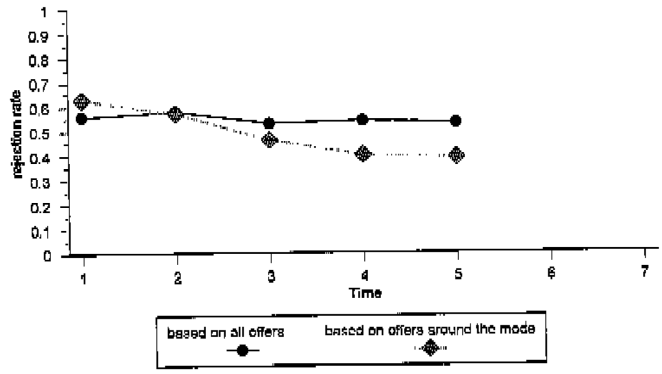

Figure 4.5

Rate of Rejection by Real Responders as a Function of Time

virtual offer range = 13-16 (gradual, 12 players)

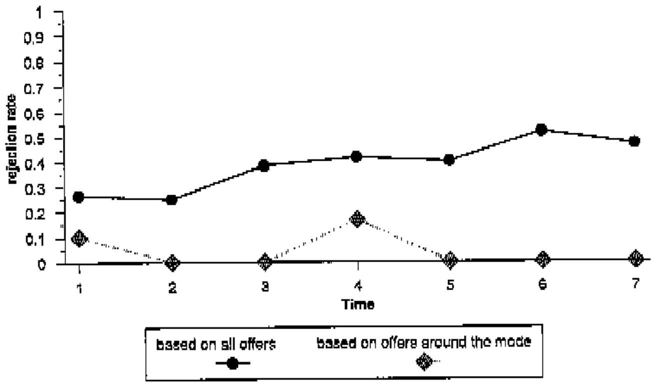

Figure 4.6

Rate of Rejection by Real Responders as a Function of Time

virtual offer range = 13-16 (gradual, 20 players)

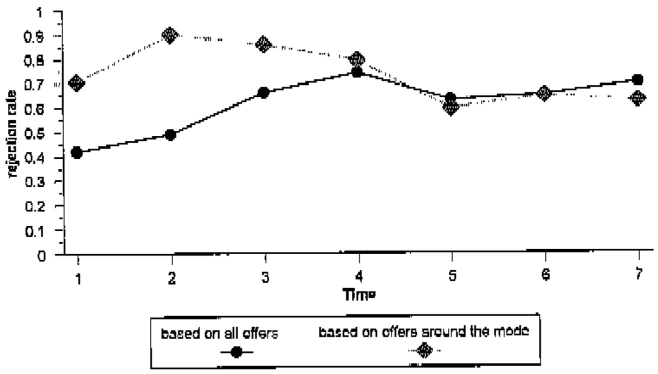

Figure 4.7
Rate of Rejection by Real Responders as a Function of Time
virtual offer range = 46-49 (12 players)

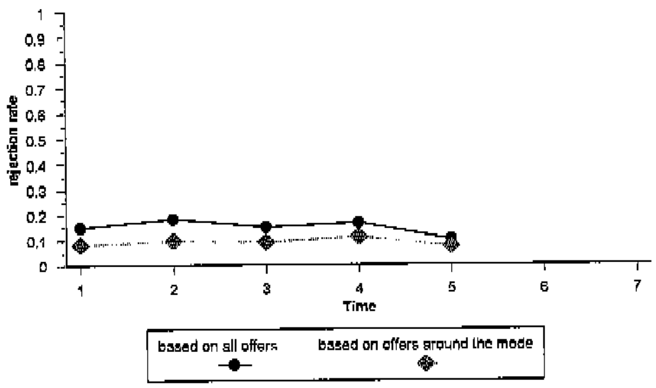

Figure 4.8
Rate of Rejection by Real Responders as a Function of Time
virtual offer range = 46-49 (18 players)

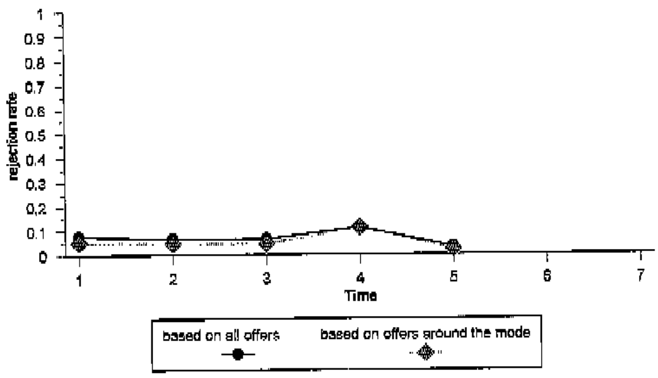

In almost all sessions there is a tendency towards a declining rate of rejection for proposals around the modal offer. In the environment of fair players the rate of rejection in the initial phase of the first 10 sessions is already very low and remains virtually unchanged through-out the session. In the environment without virtuals, the rate of rejection drops to virtually zero towards the the end of the session. However, in the other two environments (i.e., with moderately tough virtuals and with extremely tough virtuals) the results are more ambiguous. In small group sessions (Figures 4.3 and 4.5) the rate of rejection of offers around the mode is very low toward the end of the session, but it remains high in the large group sessions (Figures 4.4 and 4.6). The discrepancy between large and small groups in these environments lies in the fact that low offers are somewhat more frequent in the large groups of these environments. In particular, the modal offer of the large 13-16 group is significantly lower than the modal offer in the corresponding small group. This may be explained by the fact that large sessions are perceived as more anonymous, which encourages players to make lower offers. Note that the overall rate of rejection is quite similar in the two sessions.

4.3 Why do responders reject low offers?

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. We believe 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 <u>one shot</u> UBG. Real life UBGs 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. A responder who nods at every offer will easily teach proposers to make low offers, and his overall stream of payoffs may be pretty poor. Real life interaction in conflicts similar to the UBG develop a certain convention to which players adhere. A certain rejection rate (of low offers) is "needed" to sustain an equilibrium (convention) involving an "acceptable" (reasonably high) level of offers. When entering the laboratory subjects are endowed with this convention, which directs their behavior in the early stages of the session. In the course of the experiment players learn about their environment and adapt their behavior accordingly. A process of adaptive learning similar to the one proposed by Roth and Erev (1995) governs their behavior in later stages of the interaction. A new convention then emerges, which is environment-dependent. This results in proposers adapting their offers to the patterns of responses, and in responders decreasing their rate of rejection. The effect of outside-the-laboratory experience diminishes the longer they play (and the higher the learning stimulus is, as argued by Gale Binmore and Samuelson, 1995) but it can never be wiped out completely.

What accounts for the relative high rate of rejections in the environment with tough virtuals? The initial convention with which players are endowed at the beginning of the experiment induces a threshold for rejections, which far exceeds the offers made by tough virtuals. These offers, therefore, will almost always be rejected in the early stages of the session. But the rejection rate is not only high to begin with in these environments, it also remains relatively high as time elapses. This is so because the process of building a new convention requires adaptive behavior by *all* the participants in the environment. The virtual players do not take part in this process directly. Their "stubborn" behavior prevents them from raising their offers in response to continuous rejections. In environments with fair virtuals (whose offers are close to the initial convention), or without any virtuals, this obstacle to the emergence of a new convention does not exist. Note however, that even when this obstacle does exist, the rate of rejection decreases towards the end of the session (except for the sessions in which virtual players are introduced gradually).

5. Conclusions

5.1 Rationality

The extreme equilibrium outcome of the UBG, in which proposers get almost all and responders get virtually nothing, is often regarded by experimental economists as the only outcome that is compatible with the assumption that subjects behave rationally. We believe that strong evidence against rational behavior in the UBG is the failure of players to respond efficiently to the environment in which they play. Our experimental results indicate no such failure. Overall, and depending on the environment, players do move towards some Nash equilibrium of the UBG, but how close to Nash equilibrium this process leads depends very much on the environment. With fair virtuals the process comes very close to the 50:50 Nash equilibrium, but with tough virtuals, whose offers are very remote from the initial conventions of real subjects, the process remains relatively far from an equilibrium outcome (especially in large groups). The suggestion that the subgame perfect equilibrium is the wrong notion to focus on in analyzing behavior in the UBG is also supported by the simulation results reported by Gale, Binmore and Samuelson (1995). Gale et al., designed a replicative dynamic model to simulate plays of the UBG. They have shown that in the absence of mutations the process can move and stay at any Nash equilibrium of the game, but when mutations are introduced the dynamic process moves towards a specific Nash-equilibrium which is not the subgame perfect one.

5.2 Multi-National UBG

We have shown how behavior in the UBG is affected by the introduction of virtual players. In environments with tough virtuals real proposers make remarkably low offers but they rarely deviate from 50:50 offers in environments with fair virtuals. This strengthens the view that the notion of "reasonable" (or "decent") offers is environment-dependent, a view supported by Roth, Prasnikar, Fujiwara and Zamir (1991) in their multi-national experiment. Running the UBG in the U.S., Yugoslavia, Japan and Israel, they observe significant differences in the distribution of offers across these countries. Our results indicate that the different conventions prevailing in each country should not necessarily be attributed to some deep cultural or educational characteristics of the participants; these conventions are actually quite fragile. In small groups they may change within minutes and indeed may even dramatically approach the subgame-perfect outcome.

5.3 The Tool

This paper is about the ultimatum game, and we hope it will be judged as such. However, the paper also uses an exceptional experimental technique based on creating environments that combine real subjects with virtual players whose behavior can be fully controlled in the course of the experiment. This technique was previously used by Roth and Schoumaker (1983) to study whether manipulating players' expectations affects their behavior in binary lottery games.[13] We acknowledge the fact that this approach may appear controversial and we hope that this paper will initiate a discussion about this methodology. For our present purposes, an effective and meaningful control of the environment could only be attained by not revealing the presence of virtual players. However, we believe that after completing the project the presence of virtual players should be revealed to the subjects, in addition to a detailed explanation of the objective of the experiment and the choice of design.

We propose that the tool be judged by the way it is applied and by what it achieves. Its general usefulness for other experiments should be determined by a cost-benefit analysis. In terms of benefits, we believe that in some interactive situations moderate manipulation of the experimental environment can shed light on aspects that simply cannot be studied otherwise, especially when one aspires not only to predict the outcome of a certain interaction but also the process by which this outcome is reached .

Game Theory is often described as the Theory of Counterfactuals. Many game theoretic solutions are menu-type solutions: they specify a list of counterfactual statements such as, "If player 1's behavior is x_1, then player 2's behavior is y_2". "If Player 2 chooses z_2, then player 1 chooses w_1" etc. Unfortunately, such statements cannot be verified unless their premise is satisfied. In order to examine such statements in the laboratory, one needs at least partial control of the environment.

As for the costs, the most significant price of using this method is probably the risk of long-term contamination of future subject groups if the manipulation is revealed to the subjects after the experiment. We received no indication of credibility damage in any subsequent experiments carried out at the laboratory (either from talking to the subjects or from their questionnaires). There are two straightforward ways in which this risk can be significantly reduced. The first is to establish a large and diverse pool of potential subjects, and the second is to limit the number of experiments with manipulation conducted in one

13 Unlike our setup, in which only part of the environment involved interactions between real subjects and virtual players, in Roth and Schoumaker all real subjects played against computers in earlier trials and then switched to play against one another.

academic year at the same laboratory. This raises the issue of how to design the right agenda for a laboratory across experiments, but this issue is beyond the scope of this paper.

5.4 An Alternative Design where the Presence of Virtuals is Revealed

In addition to our main setup in which the presence of virtual players was concealed from the subjects, we have conducted a series of sessions in which the presence of virtuals was revealed. Specifically, subjects were told that during the course of the session they may be matched to a computer program instead of a real player. They were not told anything about the likelihood of this event or about the nature of these computer programs. We have run the three environments with virtual players and groups of 12 (real) subjects. These environments were run twice with the new setup each with a new subject pool, i.e., a total of additional 6 sessions. The results did not exhibit substantial differences with respect to the original design. The distributions of offers made by real players are slightly skewed towards lower offers compared with the original design (especially in the environment with $P_{13,16}$ and $R_{13,16}$) and these differences are more apparent at the beginning of a session than towards its end (see Figures 1.3a, 1.5a, 1.7a. note that we have aggregated the data of each two sessions of the same environment)[14]

14 We have established a complete analysis for this group of sessions, including all the figures obtained for the first group of sessions. These are available by request from the authors.

Figure 1.3a
Relative Distribution of Offers by Real Players
virtual offer range = 23-26 (12 players)
based on sessions where subjects were informed about the presence of virtual players

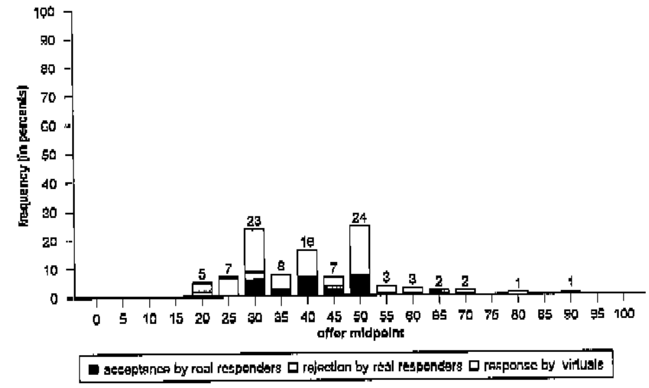

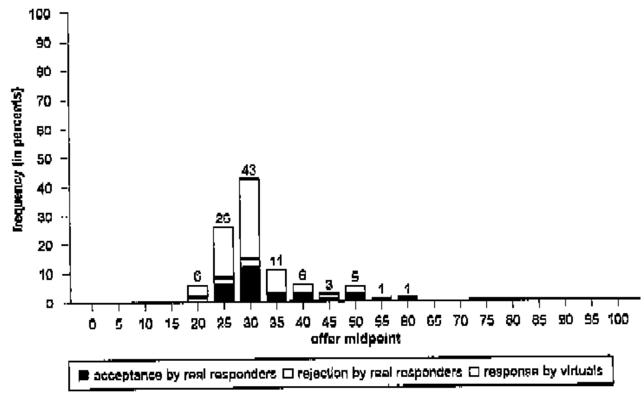

Figure 1.5a
Relative Distribution of Offers by Real Players
virtual offer range = 13-16 (12 players, aggregated data)
based on sessions where subjects were informed about the presence of virtual players

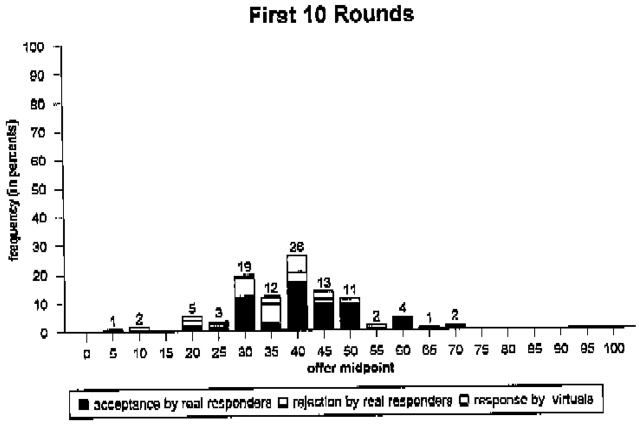

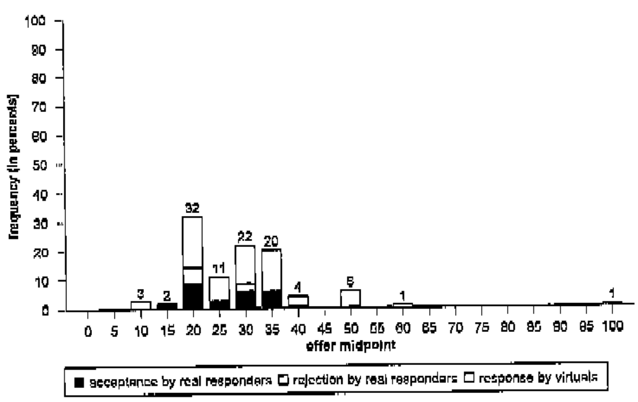

Figure 1.7a
Relative Distribution of Offers by Real Players
virtual offer range = 46-49 (12 players)
based on sessions where subjects were informed about the presence of virtual players

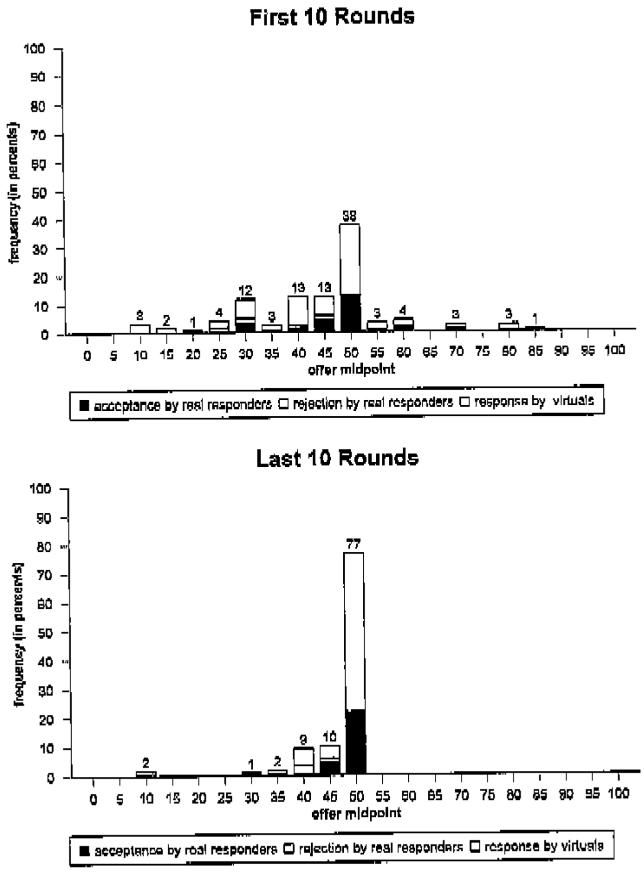

These observations seem to support the claim that offers in the game result from players attempting to study their environment and best respond to it. Subjects faced similar environments in the two designs, which resulted in a similar behavior, especially towards the end after some learning has taken place. The additional information concerning the presence of virtual players had a minor effect on their behavior. The information acquired by learning was far more relevant.

One might argue that in view of the similarity in the results we could have based our whole analysis on the second design in which the presence of virtual players was not concealed. We would find such an assertion problematic as it is clear that one cannot simply assume these similarities without verifying them by running the two designs.

References

Abbink, K., and A. Sadrieh, (1995) "Ratimage: Research Assistance Toolbox for Computer-Aided Human Behavior Experiments" University of Bonn, Discussion Paper No. B-325.

Bolton G. E., and R. Zwick, (1993) "Anonymity versus Punishment in Ultimatum Bargaining" Mimo Pennsylvania State University .

Gale J., K. G. Binmore and L. Samuelson (1995) "Learning to Be Imperfect: The Ultimatum Game." *Games and Economic Behavior* 8, 56-90.

Gueth, W., R. Schmittberger, and B. Schwarze, (1982). "An Experimental Analysis of Ultimatum Bargaining" *Journal of Economic Behavior and Organization* 3, 367-368.

Roth, A.E., and I. Erev, (1995) "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term", *Games and Economic Behavior*, Special Issue: Nobel Symposium, 8, 164-212.

Roth, A. E., V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir, (1991) "Bargaining and Market Power in Jerusalem, Ljubljana, Pittsburgh and Tokyo: An Experimental Study." *American Economic Review*. 81, 1068-1095.

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