The Japanese Economic Review
Vol. 56, No. 3, September 2005

AN EXPERIMENT WITH ULTIMATUM BARGAINING IN A CHANGING ENVIRONMENT*

By EYAL WINTER and SHMUEL ZAMIR
Hebrew University of Jerusalem

We present experimental results on the ultimatum bargaining game which support an evolutionary explanation of subjects' behaviour in the game. In these experiments subjects interacted with each other and also with virtual players, i.e. computer programs with prespecified strategies. Some of these virtual players were designed to play the equitable allocation, while others exhibited behaviour 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 behaviour of real subjects depending on the type of "mutants" (virtual players) that were present in their environment.
JEL Classification Numbers: C72, C78, C79, Z13.

1. Introduction

A basic and elementary rationality assumption asserts that a person will prefer to receive any amount of money to 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. These arguments of utility maximization and sequential rationality are crucial components of 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' behaviour differs considerably from the argument presented above. (See Thaler, 1988, and Roth and Erev, 1995, for surveys 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.[1] 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 behaviour in UBG and the subgame-perfect equilibrium solution of this game. These results also suggest that there is no real contradiction between the observed behaviour in the UBG and the rationality postulates of game theory.

In trying to explain the apparently irrational behaviour of subjects in the UBG, one has to address two questions. First, why do proposers tend to offer relatively large shares; and

---

* 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] Gueth *et al*. (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).

second, why do responders reject low offers? We will consider these two questions separately by studying the way by which subject behaviour in the UBG responds to changes in the environment. Our approach is thus evolutionary in nature, and views the emergence of a standard of behaviour 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 about the presence of virtual players;[2] 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 behaviour 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 behaviour in the UBG. Should one ascribe differences in behaviour 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 proposer and responder. 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.[3] The second type, which we call "fair", involves proposers and responders who form an equilibrium outcome that 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 behaviour 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 $\{1, 2, \ldots, K\}$ 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 $\{1, 2, \ldots, K\}$ 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

[2] 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 ran an alternative design, in which subjects were told about the presence of virtual players in the instructions; see Section 5.4.

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

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

$$h(s_1, s_2) = (K - s_1, s_1) \qquad \text{if } s_1 \in A(s_2)$$

and

$$h(s_1, s_2) = (0, 0) \qquad \text{if } s_1 \notin A(s_2)$$

(The first coordinate is the payoff to the proposer and the second, the payoff to the responder.)

We analyse the game through the notion of Nash equilibrium: a pair of strategies (s_1^*, s_2^*) is Nash a 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

$$h_1(s_1^*, s_2^*) \geq h_1(s_1, s_2^*) \qquad \text{for all } s_1,$$

$$h_2(s_1^*, s_2^*) \geq h_2(s_1^*, s_2) \qquad \text{for all } s_2,$$

and $h(s_1^*, s_2^*) = (h_1(s_1^*, s_2^*), h_2(s_1^*, s_2^*))$ 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) ($i = 1, 2, \ldots, K$) by having player 1 propose this share and player 2 accept at least this share and reject anything less. Among this set of Nash equilibria, two are *subgame-perfect* and can be worked out by simple backward induction.[4] The first corresponds to an allocation in which the proposer gets everything, i.e. (K, 0), and the second to when the responder is granted only one unit. i.e. ($K - 1$, 1).[5] 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

[4] 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).

[5] In the continuous version of the UBG, only the first allocation is sustainable by a subgame-perfect equilibrium.

laboratory Ratiolab*,* at the Centre for Rationality at Hebrew University. The computer program heavily used the Ratimage software, developed in the University of Bonn by Abbink and 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 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 would accept any offer of more than 14 points and would 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.

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).[6] This role was fixed throughout the session. The subjects played the UBG for either 50 or 70 rounds, depending on the session, and were informed about the length of the session and the fact that matching in each period is random. 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 throughout 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 matched only among themselves.[7] We will discuss the aspect of this (admittedly unconventional) approach later. In addition to six sessions with virtual players (two sessions of different sized groups for each type), we ran two sessions with no virtual players at all. Each subject participated only in a single session.

Table 1 describes the full specification of each of the eight experiments that we have conducted. In addition to the eight sessions reported above, we ran a separate set of six sessions in which we twice repeated the three 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 might 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.

[6] The random choice of roles was used in order to guarantee that all subjects had the same *ex ante* earning opportunities.

[7] This was confirmed by a questionnaire that the subjects filled at the end of the experiment.

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 virtuals	–
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 with 14	at each round half are $P_{13,16}$ and half are $R_{13,16}$
7	50	12	8 all through the session	4 are $P_{46,49}$ 4 are $R_{46,49}$
8	50	18	12 all through the session	6 are $P_{46,49}$ 6 are $R_{46,49}$

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 NIS1 per 100 points. In addition, each player received a fixed payment NIS10 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 comparison. The first involves comparing different sessions at the same stage. The other concerns the evolution of the behaviour in the same session over time.

Figures 1(a)–(h) show the distribution of offers made by *real* players in the first and the last ten 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.

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 behaviour is strikingly different: offers below 50 points vanish almost completely, and the distribution is unambiguously concentrated on the 50 : 50 offers. One observation that is consistent across all sessions is that the distribution of offers in the first ten rounds is more widely dispersed than that in the last ten 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.

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

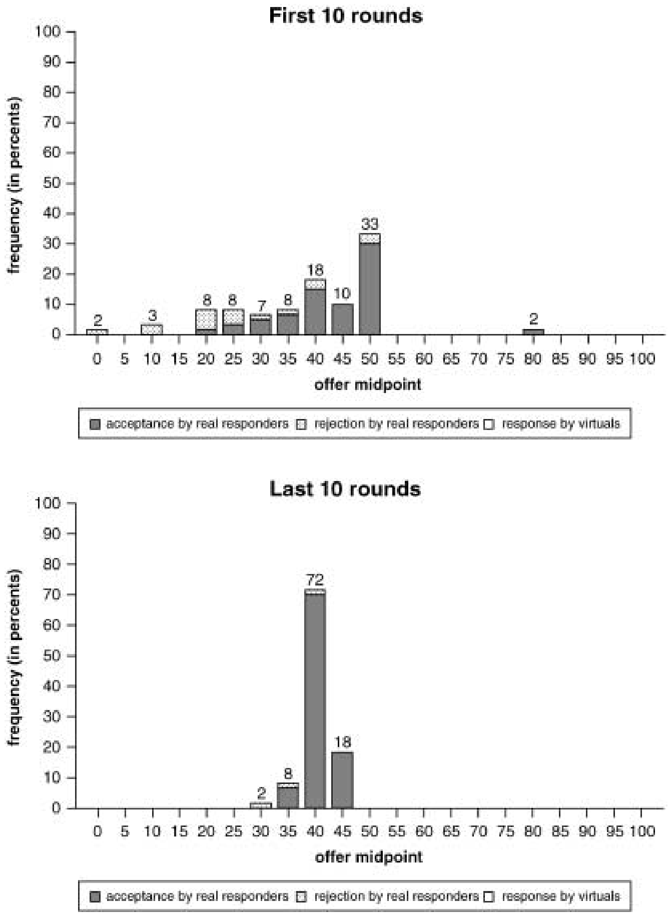

FIGURE 1(a) Relative distribution of offers by real players
No virtual players (12 players)

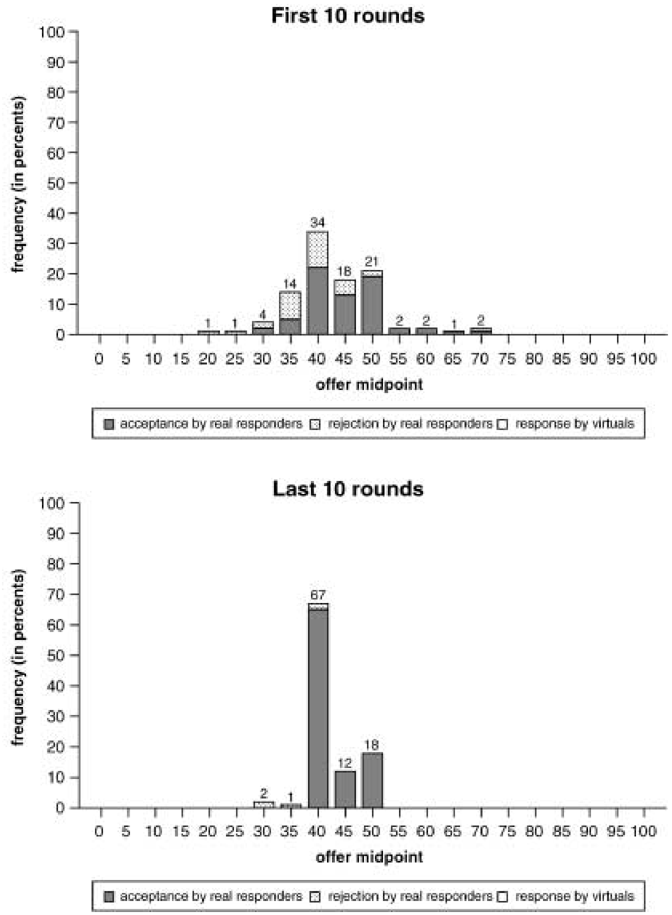

FIGURE 1(b) Relative distribution of offers by real players
No virtual players (20 players)

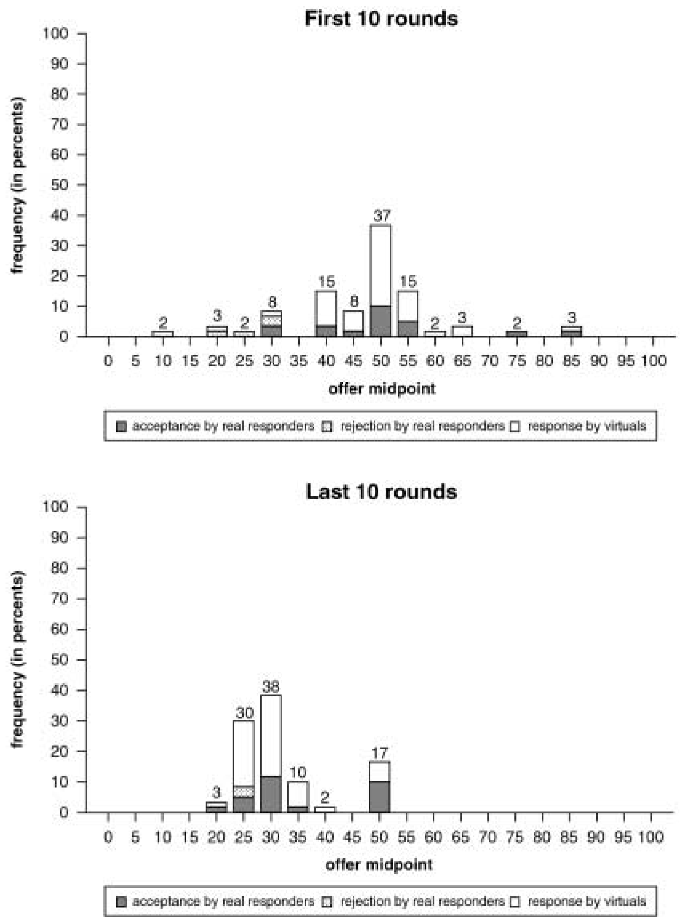

FIGURE 1(c) Relative distribution of offers by real players
Virtual offer range: 23–26 (12 players)

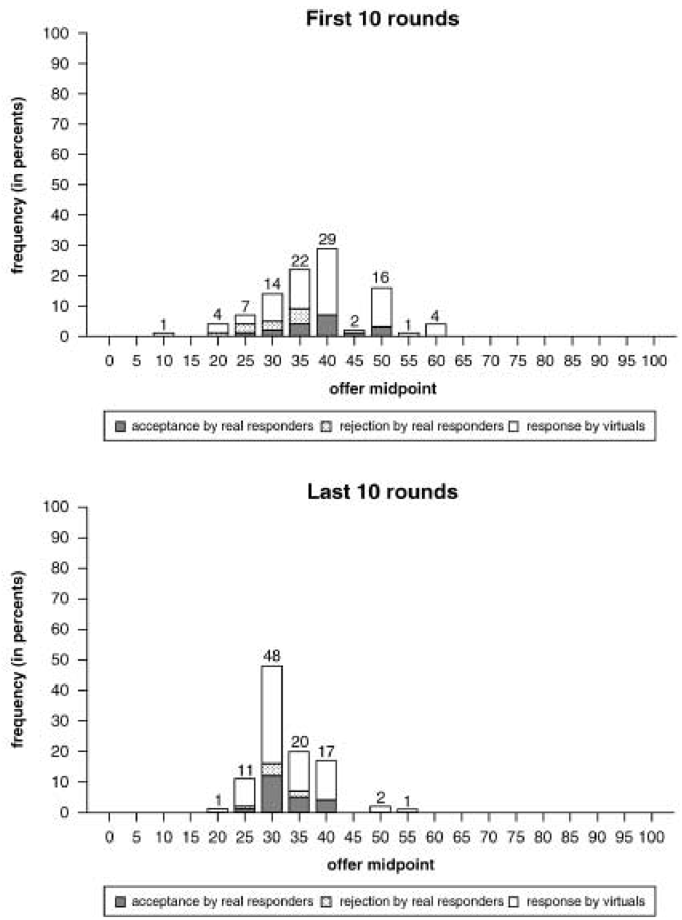

FIGURE 1(d) Relative distribution of offers by real players
Virtual offer range: 23–26 (20 players)

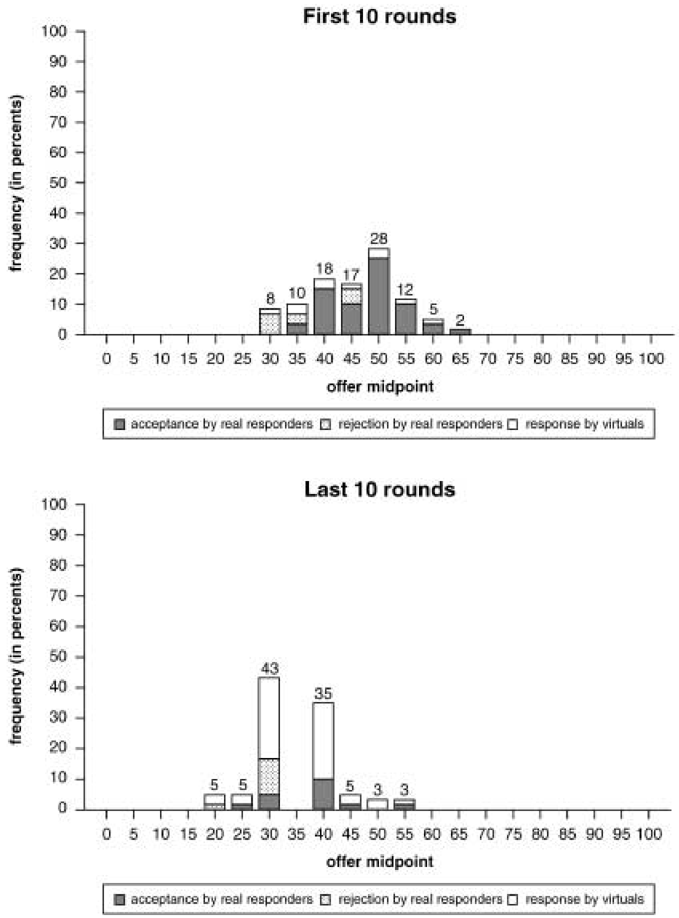

FIGURE 1(e) Relative distribution of offers by real players
Virtual offer range: 13–16 (gradual, 12 players)

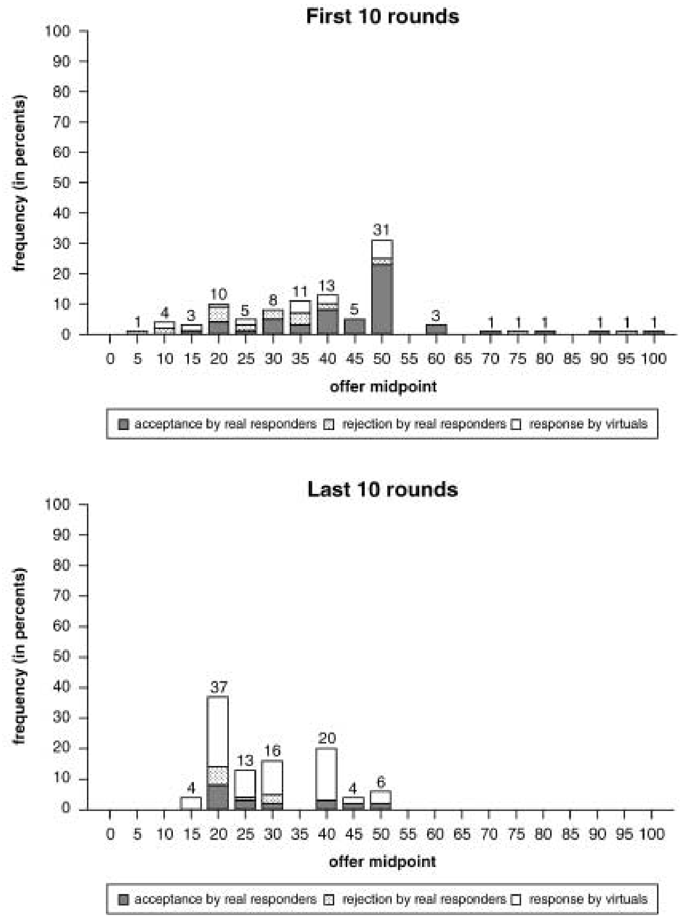

FIGURE 1(f) Relative distribution of offers by real players
Virtual offer range: 13–16 (gradual, 20 players)

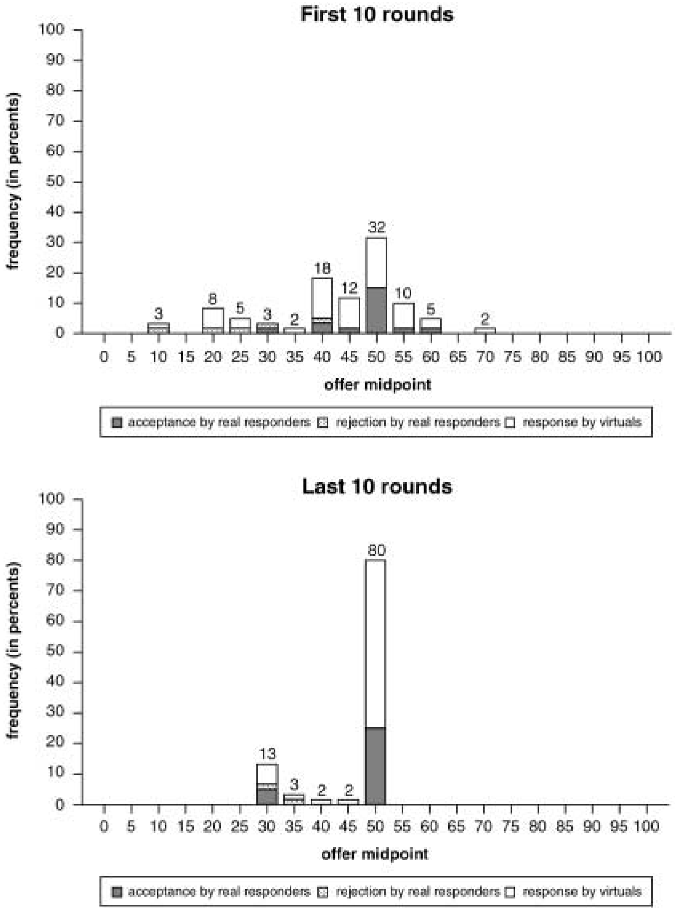

FIGURE 1(g) Relative distribution of offers by real players
Virtual offer range: 46–49 (12 players)

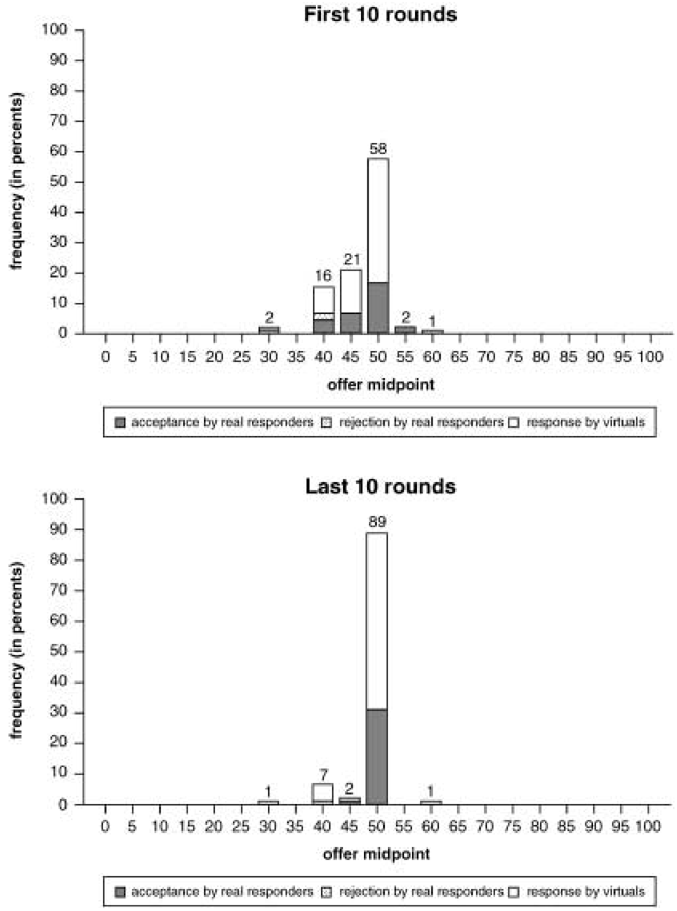

FIGURE 1(h) Relative distribution of offers by real players
Virtual offer range: 46–49 (18 players)

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

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		
Experimental condition A	13–16	23–26	No virtuals
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 too change their behaviour, 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' behaviour 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 that virtual responders do.

With "fair" virtuals, the story is pretty much the same. However, here the behaviour of virtual players is much closer to the initial patterns of real subject behaviour. 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(a)–(h) show the way in which the modes and average offers evolve in time in each of the environments.[8] It is apparent from these charts that the environment changes gradually.

[8] In Figures 2(a)–(h) and 4(a)–(h) the points, 1, 2, . . . on the horizontal axis represent segments of 10 rounds (i.e., 1 represents rounds 1–10, 2 represents rounds 11–20, etc.)

With tough virtuals the indicators decrease constantly, and with fair virtuals they increase gradually.

We argued earlier that in the long run proposers submit offers that are their best response with respect to the environment in which they "live". We checked the extent to which the proposers were "expected utility-maximizers". To do that, we took the empirical frequencies of responses as estimates of the probabilities that a certain proposal would be accepted.[9] This enabled us to plot the "expected" return for offers in each environment. For each environment, we partitioned the range of offers into intervals of 5 points, and disregarded those intervals with less than ten offers. The expected return for each interval T was 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 was done for the eight sessions, and the results are shown in Figures 3(a)–(h).

These charts 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.

4.2 Responses

Interpreting responders' behaviour 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 offers are not 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 were documented using two methods. First, the histograms of offers (Figures 1(a)–(h)) include the proportion of acceptances and rejections by *real* players for offers made by real players. Second, Figures 4(a)–(h) plot the rejection rates as a function of time across all offers by real and virtual players (the bold 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.

In almost all sessions, there is some 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 ten sessions was already very low and it remains virtually unchanged throughout the session. In the environment without virtuals, the rate of rejection drops to virtually zero towards the end of the session. However, in the other two environments of tough virtuals (i.e. moderate and extreme tough virtuals) the results are ambiguous. The rate of rejection depends on the mode of the distribution of offers: In small group sessions (Figures 4(c) and (e)), where the modal offer is relatively high, the rate of rejection of offers around the mode is very low towards the end of the session, but it remains high in the large group sessions (Figures 4(d) and (f)), which have a lower modal offer.

[9] For technical reasons, some of the virtual responses were not recorded. In such cases we calculated the acceptance probability using virtual responders' characteristics.

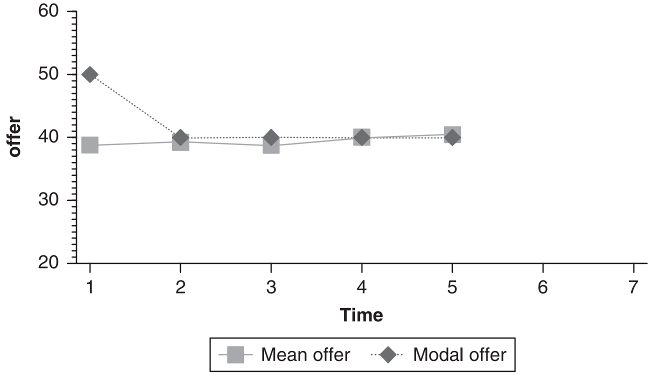

FIGURE 2(a) Mean and mode of offers by real players as a function of time
No virtual players (12 players)

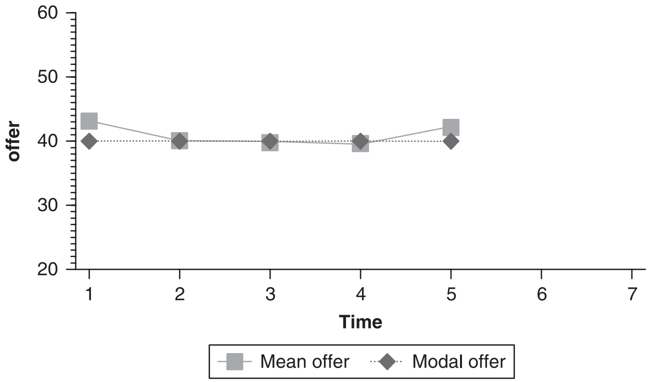

FIGURE 2(b) Mean and mode of offers by real proposers as a function of time
No virtual players (20 players)

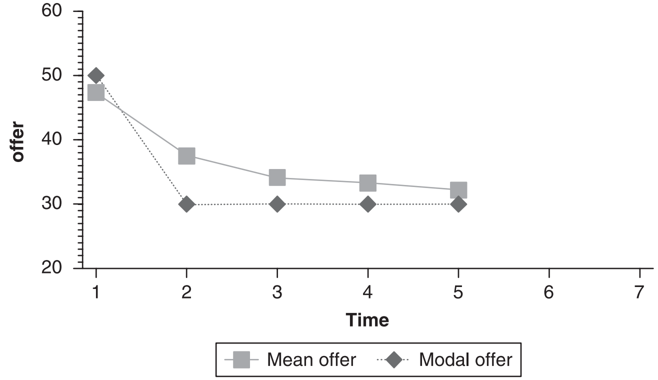

FIGURE 2(c) Mean and mode of offers by real players as a function of time
Virtual offer range: 23–26 (12 players)

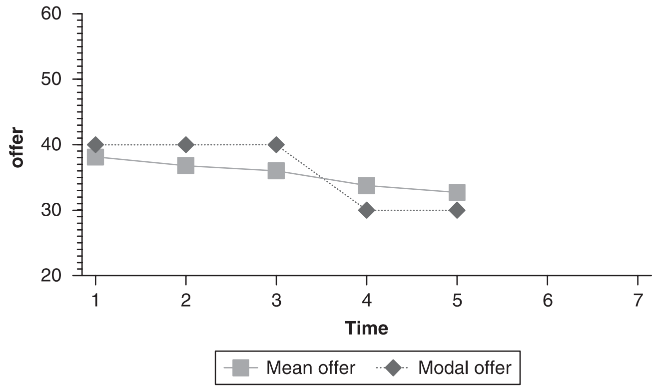

FIGURE 2(d) Mean and mode of offers by real proposers as a function of time
Virtual offer range: 23–26 (20 players)

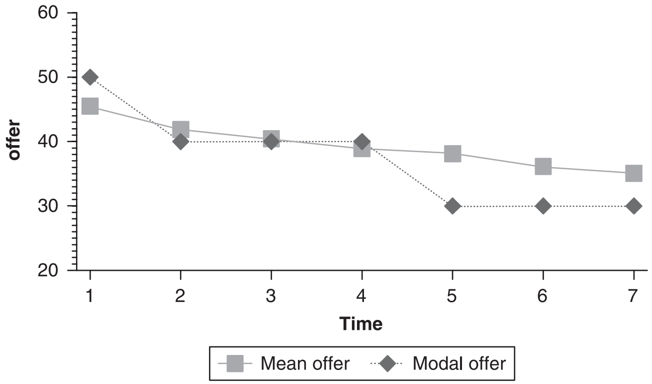

FIGURE 2(e) Mean and mode of offers by real players as a function of time
Virtual offer range: 13–16 (12 players)

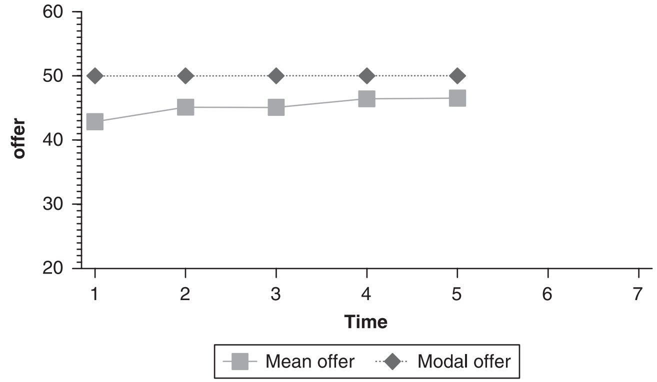

FIGURE 2(g) Mean and mode of offers by real players as a function of time
Virtual offer range: 46–49 (12 players)

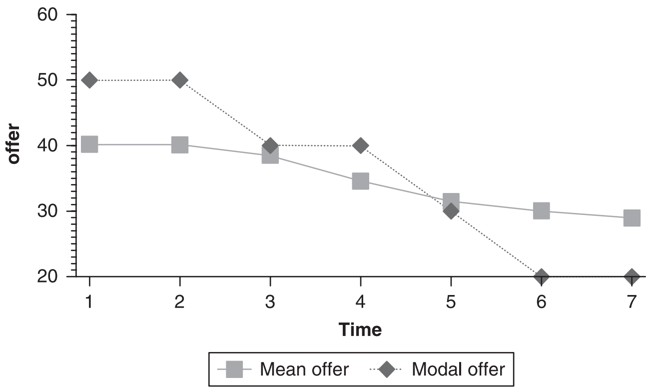

FIGURE 2(f) Mean and mode of offers by real proposers as a function of time
Virtual offer range: 13–16 (20 players)

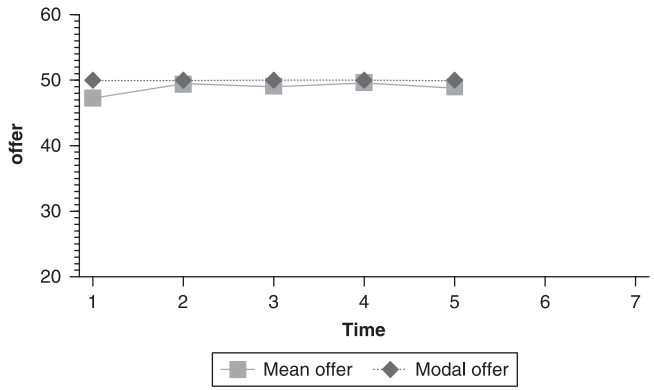

FIGURE 2(h) Mean and mode of offers by real proposers as a function of time
Virtual offer range: 46–49 (18 players)

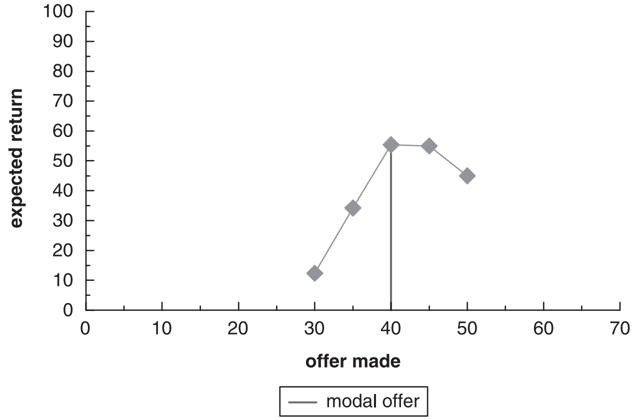

FIGURE 3(a) Expected returns for offers made by real players, based on virtual and real responses
No virtual players (12 players)

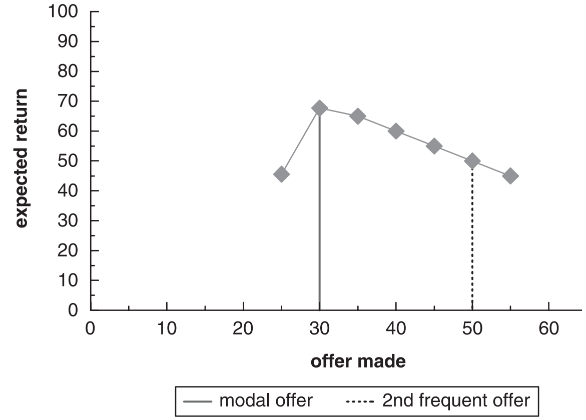

FIGURE 3(c) Expected returns for offers made by real players, based on virtual and real responses
Virtual offer range: 23–26 (12 players)

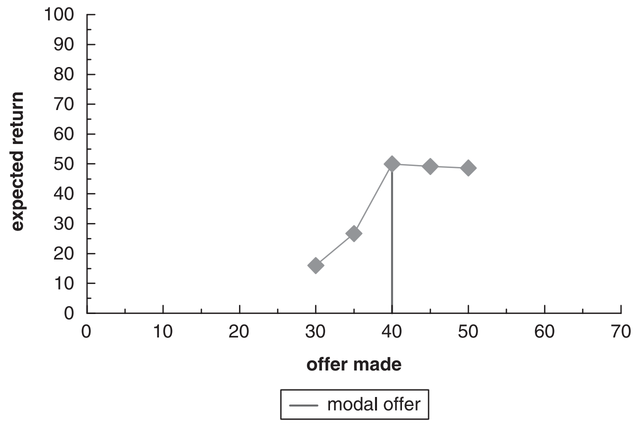

FIGURE 3(b) Expected returns for offers made by real players, based on virtual and real responses
No virtual players (20 players)

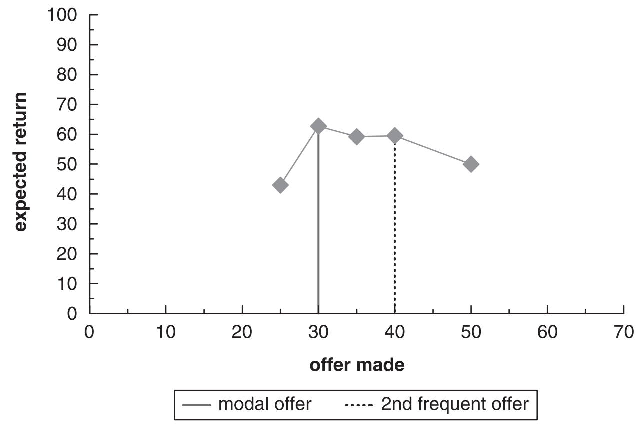

FIGURE 3(d) Expected returns for offers made by real players, based on virtual and real responses
Virtual offer range: 23–26 (20 players)

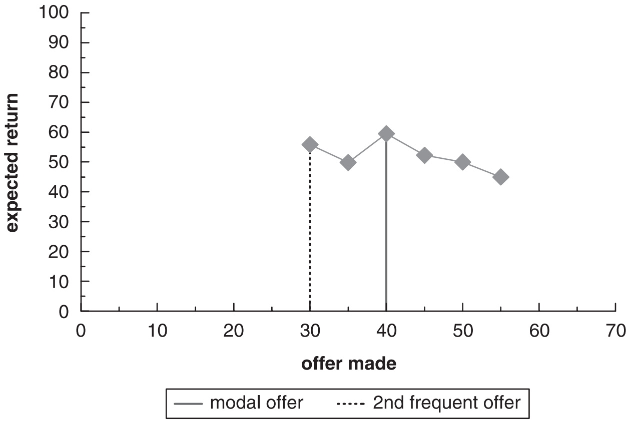

FIGURE 3(e) Expected returns for offers made by real players, based on virtual and real responses
Virtual offer range: 13–16 (12 players)

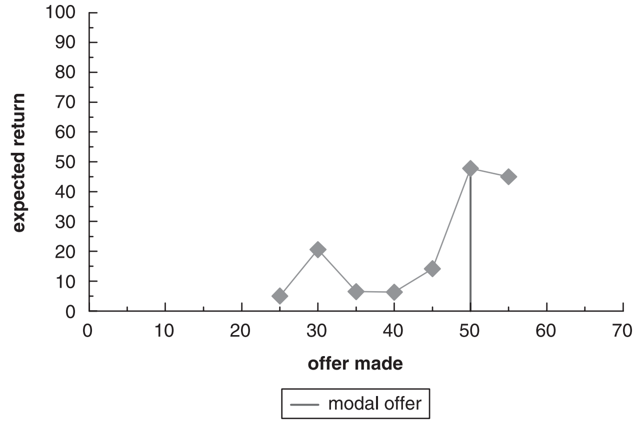

FIGURE 3(g) Expected returns for offers made by real players, based on virtual and real responses
Virtual offer range: 46–49 (12 players)

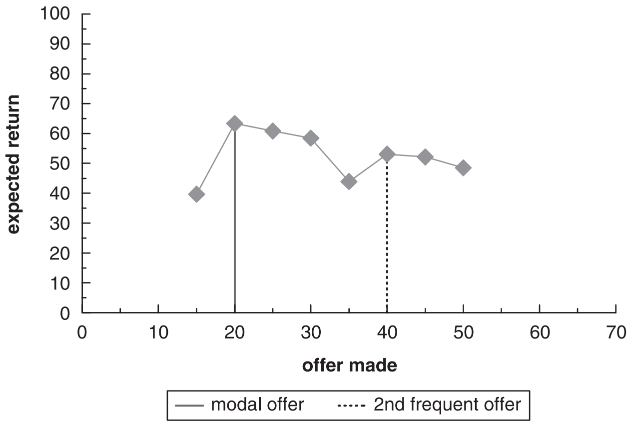

FIGURE 3(f) Expected returns for offers made by real players, based on virtual and real responses
Virtual offer range: 13–16 (20 players)

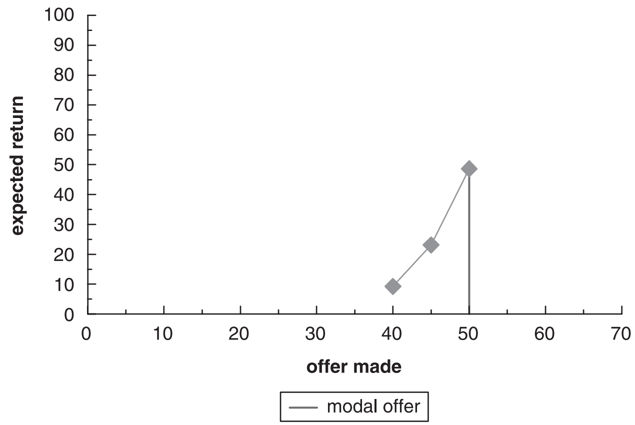

FIGURE 3(h) Expected returns for offers made by real players, based on virtual and real responses
Virtual offer range: 46–49 (18 players)

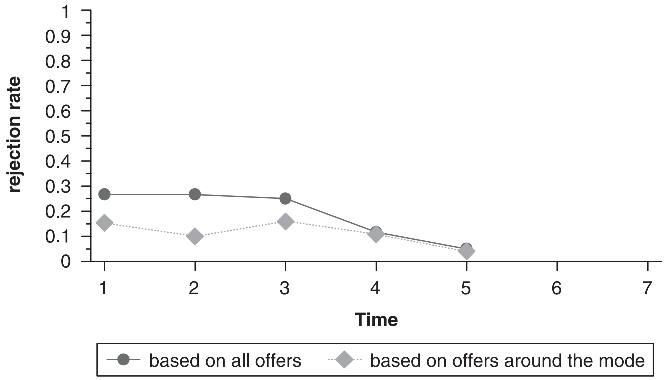

FIGURE 4(a) Rate of rejection by real responders as a function of time
No virtual players (12 players)

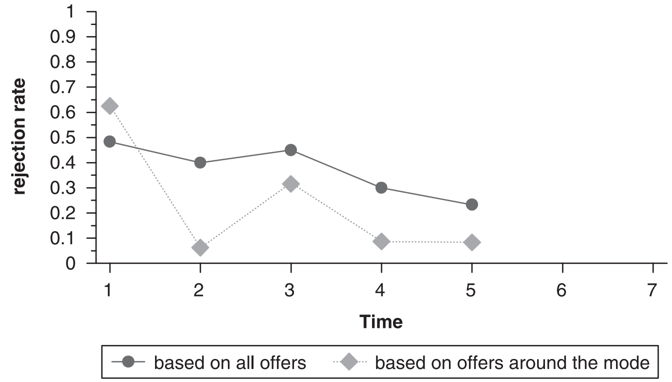

FIGURE 4(c) Rate of rejection by real responders as a function of time
Virtual offer range: 23–26 (12 players)

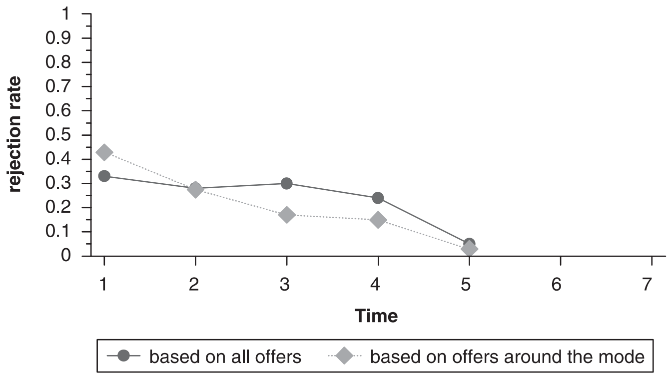

FIGURE 4(b) Rate of rejection by real responders as a function of time
No virtual players (20 players)

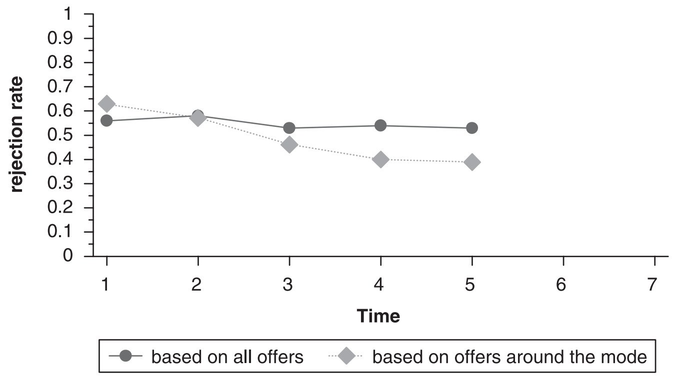

FIGURE 4(d) Rate of rejection by real responders as a function of time
Virtual offer range: 23–26 (20 players)

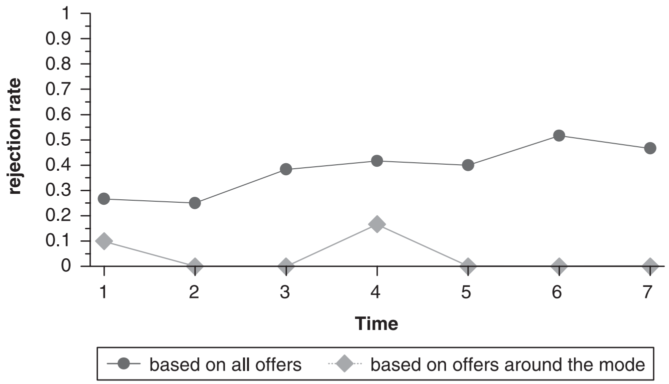

FIGURE 4(e) Rate of rejection by real responders as a function of time
Virtual offer range: 13–16 (gradual, 12 players)

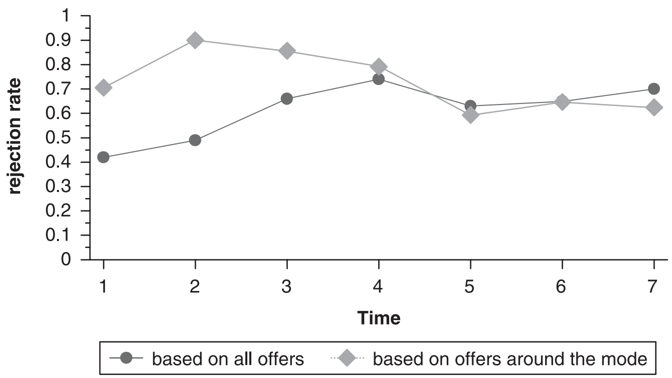

FIGURE 4(f) Rate of rejection by real responders as a function of time
Virtual offer range: 13–16 (gradual, 20 players)

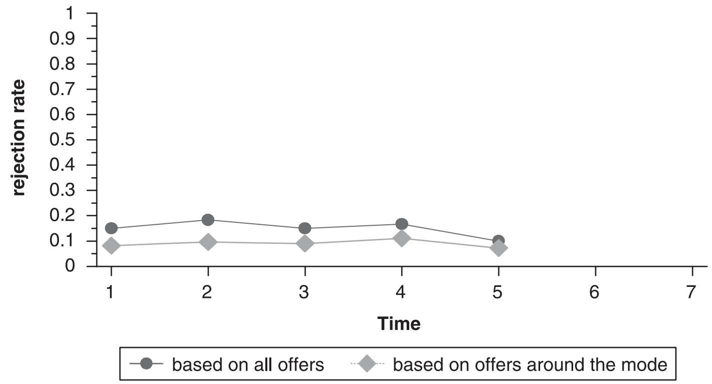

FIGURE 4(g) Rate of rejection by real responders as a function of time
Virtual offer range: 46–49 (12 players)

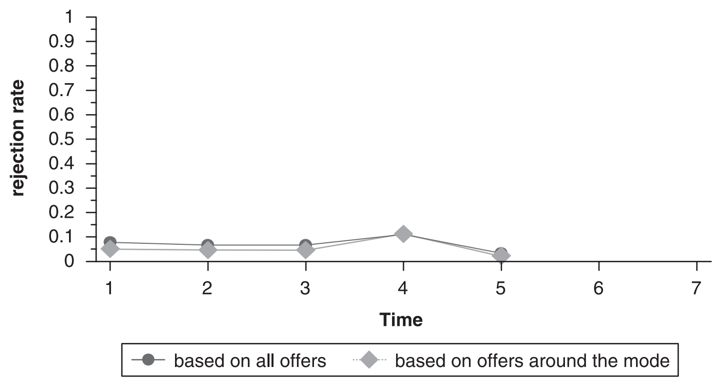

FIGURE 4(h) Rate of rejection by real responders as a function of time
Virtual offer range: 46–49 (18 players)

TABLE 4
Rate of Rejection by Real Players to Offers Made by Both Real and Virtual Players

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

The differences between small and large sessions in terms of modal offers may be explained by the fact that large sessions are perceived as more anonymous, which encourages players to make lower offers. When looking at the overall rate of rejection in the two "tough" environments, there is very little evidence of a declining rate of rejection. In the environment of 13–16 virtuals the rate of rejection even increases in time, which is due partly to the gradual introduction of the virtual players. In both environments, most of these offers are rejected even towards the end of the session. This leads us to conclude that, in contrast to the proposers, who adapt well to the environment in which they operate, responders do not perform that well in responding optimally to their environment.

4.3 Why do responders reject low offers?

It is more difficult to explain the behaviour of responders in the ultimatum game than to explain proposers' behaviour, and our design does not offer a complete picture for this issue. The considerable rate of rejection that we observe in response to low offers can be explained by negative reciprocity or inequality aversion (see e.g. Bolton and Ockenfels, 2000, and Fehr and Schmidt, 1999). These behaviours, while irrational in a one-shot environment, are rational in real-life environments similar to the UBG, where the interaction is not anonymous and agents may get to play with each other more than once. In such environments rejections play an important role in 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.

Although subjects fully understand the rules of the game and its payoff structure, their behaviour may be 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 in which players completely disregard their real-life experience in similar situations. When entering the laboratory, subjects are endowed with conventions and standards of behaviour that work well outside the lab. In the course of the experiment, they learn about their new environment and gradually adapt to it. A process of adaptive learning similar to that proposed by Roth and Erev (1995) governs their behaviour in later stages of the interaction.

However, our results give a strong indication that proposers learn to adapt to their environment much better than responders. To understand this observation, one should focus on the environments with tough virtuals (Less learning takes place in the "fair" environment because virtuals' offers do not differ much from the initial convention.) Why do responders persist in rejecting low offers in these environments at the same time that proposers adapt extremely well to their environment? One rough answer to this question is that proposers learn "good news" in these environments, i.e. the fact that low offers are accepted, while responders face "bad news", i.e. the fact that they have to put up with low offers. But there is a more intrinsic reason for this difference.

We believe that there are two effects that govern subjects' behaviour at each period of the game. The first is the adaptation effect, which drives subjects to respond optimally to their environment in order to maximize the current payoff. The other is a future effect, by which subjects try to affect the future behaviour of their partners. With their actions they try to induce the convention in future to move in their preferred direction. As far as proposers are concerned, the two effects work in the same direction, i.e. towards their making lower offers. For responders, however, they work in opposing directions. Adaptation instructs them to accept low offers, but the future effect tells them to reject low offers in order to induce higher ones in the future. Thus, while the future effect interferes with adaptive learning for responders, it enhances adaptation for proposers. We believe that the future effect exists even when the rules of the game prescribe that no two individuals interact more than once, because of players' perception that the game is part of a larger repeated interaction. The effect of outside-the-laboratory experience diminishes the longer they play (and the higher the learning stimulus is, as argued by Gale *et al.* 1995), but it can never be wiped out completely.

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 the failure of players to respond efficiently to the environment in which they play is strong evidence against rational behaviour in the UBG. Our experimental results indicate no such failure as far as proposers are concerned. However, responders, while acting differently in different environments, persistently reject offers that are exceedingly low. Overall, and depending on the environment, players do move towards some Nash equilibrium of the UBG, but how close to Nash equilibrium the 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 analysing behaviour in the UBG is also supported by the simulation results reported by Gale *et al.* (1995). These authors designed a model of replicative dynamics to simulate plays of the UBG. They show 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 Multinational UBG

We have shown how behaviour in the UBG is affected by the introduction of virtual players. In environments with tough virtuals real proposers make remarkably low offers, but

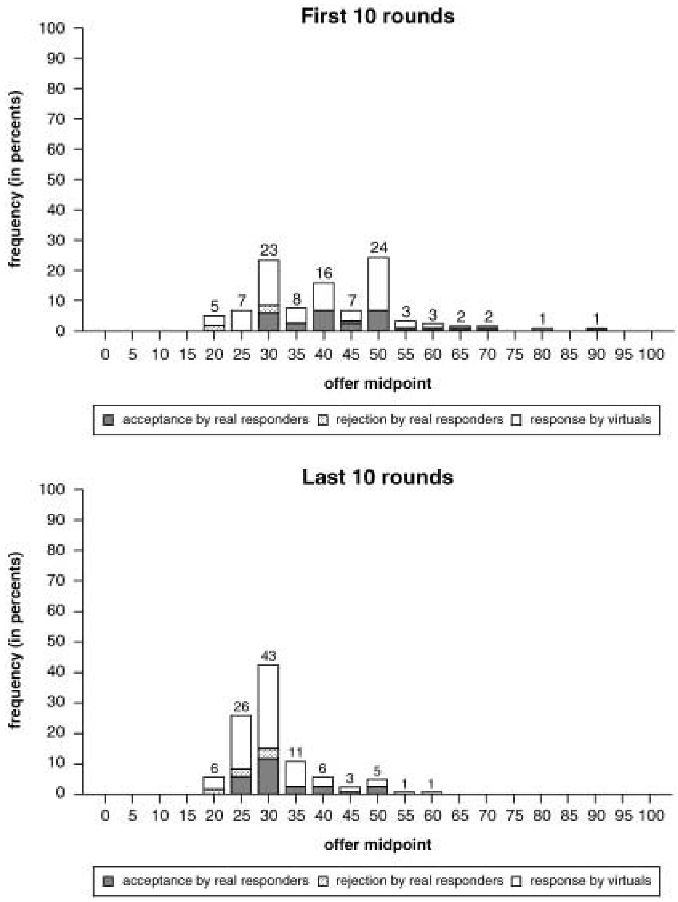

FIGURE 1(c)* Relative distribution of offers by real players, based on sessions where subjects were informed about the presence of virtual players
Virtual offer range: 23–26 (12 players)

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 *et al.* (1991) in their multinational experiment. Running the UBG in the United States, 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 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 conducted a series of sessions in which the presence of virtuals was

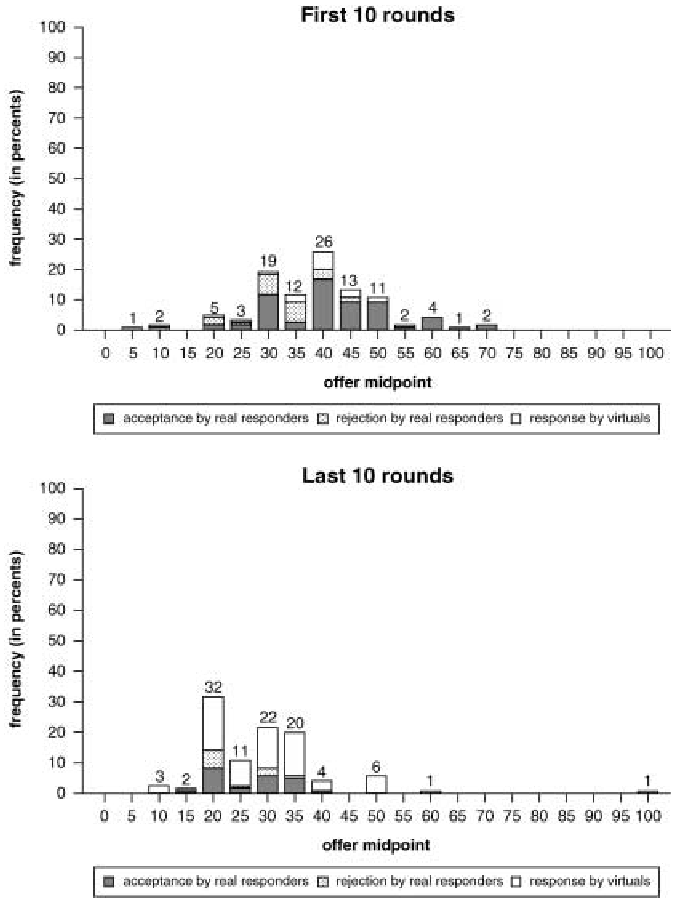

FIGURE 1(e)* Relative distribution of offers by real players, based on sessions where subjects were informed about the presence of virtual players
Virtual offer range: 13–16 (gradual, 12 players)

revealed. Specifically, subjects were told that during the course of the session they might 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 ran 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 six 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(c)*, 1(e)* and 1(g)*; note that we have aggregated the data of each two sessions of the same environment.)[10]

[10] 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.

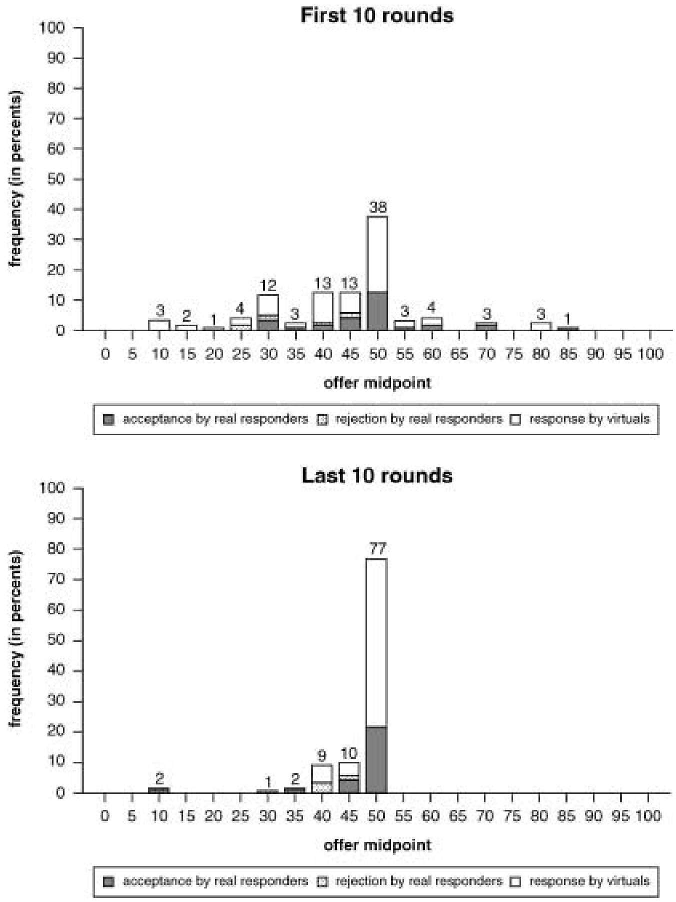

FIGURE 1(g)* Relative distribution of offers by real players (based on sessions where subjects were informed about the presence of virtual players)
Virtual offer range: 46–49 (12 players)

These observations seem to support the claim that offers in the game result from players attempting to study their environment and respond to it as well as possible. Subjects faced similar environments in the two designs, which resulted in similar behaviour, especially towards the end after some learning had taken place. The additional information concerning the presence of virtual players had a minor effect on their behaviour. 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, however, as it is clear that one cannot simply assume these similarities without verifying them by running the two designs.

Final version accepted on 26 November 2004.

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", mimeo, Pennsylvania State University.

Bolton, G. and A. Ockenfels (2000) "ERC: A Theory of Equity, Reciprocity and Competition", *American Economic Review*, Vol. 90, pp. 166–193.

Fehr, E. and K. Schmidt (1999) "A Theory of Fairness Competition and Cooperation", *Quarterly Journal of Economics*, Vol. 114, pp. 817–868.

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

Gueth, W., R. Schmittberger and B. Schwarze (1982) "An Experimental Analysis of Ultimatum Bargaining", *Journal of Economic Behavior and Organization*, Vol. 3, pp. 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, Vol. 8, pp. 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*, Vol. 81, pp. 1068–1095.

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

Thaler, R. H. (1988) "The Ultimatum Game", *Journal of Economic Perspective*, Vol. 4, pp. 195–206.