Survival of Altruistic Preferences in the Ultimatum Game - an Agent-based Approach

Zsombor Z. Méder

Maastricht University, Department of Economics
P.O. Box 610, 6200 Maastricht, the Netherlands
email: Z.Meder@maastrichtuniversity.nl

Abstract. We examine the survivability of altruistic preferences in the Ultimatum Game through two sets of agent-based simulations. We find that a self-centered, memory-based strategy updating provides a more plausible basis for altruism than classic imitate-your-neighbor learning. If memory of second-player acceptance thresholds is longer than memory of first-player offers, then our model behavior is consistent with the results of human experiments.

1 Introduction

The ultimatum game (UG) has been extensively studied by experimental economists, mainly because it clearly demonstrates the shortcomings of subgame-perfectness as far as predicting actual game behavior: the unique subgame-perfect equilibrium is almost never observed in experiments. This peculiarity places the UG into a whole web of games in which behavior should clearly be explained using common principles, including, but not limited to: Dictator Game, Public Goods Game, Trust Game, Market Games, and, obviously, other kinds of bargaining games. However, although many scholars have devoted efforts to understanding these, we are still very far from a sufficiently insightful framework.

We think that, since strategies for typical social situations are the results of long-run adaptation, an evolutionary approach is indispensable for the understanding of behavioral patterns of social interaction. In this study, we employ the evolutionary approach through agent-based simulation. Our research question focuses on the survivability of certain type of non-subgame-perfect preferences. Thus, we depart from the standard „evolution of cooperation"-literature, insofar as we are not chiefly concerned with the emergence of such strategies, but their long-term survival. As far as the grand project of understanding cooperation is concerned, we believe that these strategies can be interpreted to be the reflection of genuinely altruistic preferences. However, to find evidence for this, a more general approach would be needed, taking into account novel utility function types, meta-strategies and learning algorithms, drawing both from the behavioral literature and agent-based simulations. We make no such attempts in this work.

The paper is structured as follows: in Section 2, we describe the game and review the most significant results of the empirical literature. Next, we provide reasons for the agent-based approach employed. Sections 4 and 5 describe the simulation, and its main results. The last section concludes, hinting towards potential future experimental work.

2 Historical Background

The UG is a special, one-turn bargaining game. Let the value of the pie divided between the first players be normalized to 100. Player 1 proposes a division of the pie, asking $100-a$ for herself $(0 < a < 100)$, offering a for player 2. Player 2 either accepts or declines. In the former case, players divide the pie according to the proposal; in the latter, both leave empty-handed. Thus, the vector of monetary payoffs are $(100 - a, a)$ and $(0, 0)$ for acceptance and rejection respectively. It is easy to verify that, in subgame-perfect equilibrium, the player 1 gets at least $100-\epsilon$, where ϵ is the smallest monetary unit available, and player 2 does not get more than ϵ, accepting any positive offer.[1] So, in subgame-perfect equilibrium, player 1 fully exploits player 2. The set of Nash-equilibria is, however, much larger: it contains all strategy pairs where player 1's strategy is to offers t, and player 2 accepts t, and wouldn't accept any other $t' < t$ from player 1. However, for no such threshold $t > epsilon$ (we will call t player 2's „acceptance threshold") is player 2's threat of refusing lower offers credible: she should accept any positive offer.

The first experimental results with the Ultimatum Game were presented in Güth, Schmittberger and Schwarze (1982). It found that the modal offer for player 1 is half of the pie. Furthermore, if she offers less than one third of the pie, she faces a chance of more than 50% of being rejected. These results were replicated numerous times since 1982. A major change in approach came about in the 90's, when researchers started to extend the set of subjects, first by setting up experiments in undeveloped or non-western countries (Roth et al. (1991)), then gradually shifting attention from university students to other classes of the society. Heinrich et al. (2004a) reports more than a dozen experiments conducted by anthropologists and economists in a wide set of primitive societies. In terms of the first anomaly - that of the behavior of player 1 - they all confirm the results of Güth, Schmittberger and Schwarze (1982); pure exploitation is nowhere practiced. However, in some societies, acceptance thresholds of second movers are indeed very low or close to 0. The modal offer varies greatly, from 10% to over 60%. This indicates the existence of a culture-specific norm of fairness in terms of bargaining behavior.

A great deal of attention has been given to these results. A possible way to attack their was to point out that the monetary incentives typically used are simply too low to induce subgame-perfect behavior. However, in some experiments the pie was increased to three months' average way for the concerned community, and offers moved even *closer* to the fairness norm. It could also be suggested that subjects are unable to calculate subgame-perfect equilibria. Unfortunately, this

[1] In the case of continuous offer set, the equilibrium payoff to player one is either 1, or converges to it.

argument will not work, since even in Güth's original experiment, subjects were also presented with a more complex division problem, where they could clearly solve for subgame-perfect equilibrium. Binmore 2006 and others have claimed that the experimental scenario is too remote from everyday problems (subjects might not comprehend its one-shot nature and such), and so they would need time to learn how to play most efficiently and adapt their strategies. However, Roth et al. (1991) have compared learning in the UG to one where there was a competition between first-movers (the best offer would be accepted), calling this the market game. Whereas in the market-game subgame-perfectness (and radical exploitation) emerged after a couple of rounds, behavior in the UG actually *drifted away* from the game-theoretical prediction, the distribution of offers concentrating around the - culture-specific - norm. As a partial explanation, one might say that player 1 is afraid of an „irrational" player 2, and is offering high shares of the pie as a strategic move, to avoid player 2's veto. This, combined with the risk aversion, might offer a partial explanation, but Heinrich et al. (2004b) shows that these factors combined would still result in lower offers.

Attempting to synthesize these results, it has been suggested that players, in fact, behave as if maximizing an inter-dependent utility function, where the payoff of each player depends on the monetary payoff of each other. This is in line with results obtained from the Dictator Game, where the second player has no effect on the outcome: here, player 1 typically still gives a positive amount to the second player. Other experiments, where, for example, reputation effects are minimized, seem to counter this hypothesis, however, for the time being, we accept the existence of utility functions depending on both player's payoffs. We will call these „altruistic" preferences.

3 Preliminary analysis

As a methodological question, we must first consider the relationship between monetary payoffs and agent utility. One critique to *all* of the above-mentioned experimental results would be that they make the trivial mistake of confusing monetary payoffs with subjective utility. Since we cannot peak into the minds of our subjects, it is impossible to tell whether their behavior actual conflicts with the predictions of subgame-perfectness. This is a legitimate point: but the new task then becomes the construction of utility functions that conform to *all* of the experimental findings. This of course, seems a quite hopeless quest at the moment.

In this paper we examine the survivability of utility functions where the resulting strategy of players can be described by a rational number between 0 and 100 in each position. Although some equilibrium concepts like the „psychological Nash equilibrium" are incompatible with this approach, it can be argued that for player 1, this is not a major simplification, when players are - and will be - unknown to each other, which is true in most experimental designs. For player 2, the strategy set of the original games is all the functions from the $[0, 100]$ interval to $\{0, 1\}$. The condensation of player 2's strategy into a single number, the acceptance threshold, essentially allows only for monotonous functions. Our approach is consistent with the linear function form, proposed by Costa-Gomez and Zauner (2001) and utility functions where players have a conception of „fair" division, and lose utility linearly for every deviation from that division, like in Fehr and Schmidt (1999).

We employ agent-based simulation, since it can easily be shown that, in a traditional evolutionary game setting, the evolutionary process with the standard replicator dynamics (or any other dynamic strictly monotonous on payoffs) drives strategies to the subgame-perfect equilibrium, the resulting ESS once again creating full exploitation for player 2's. To this effect, suppose that a monomorphic population of $\{a, a\}$ is stable with $a > 0$. Since all observed offers are a, mutants that have acceptance levels lower than a can appear. Thus, with probability 1, we will arrive at a stage where, for some $\delta > 0$ mutants that offer $a - \delta$ will do better those offering a. Thus, we decided to use agent-based simulation, where strategy updating is myopic and local.

4 Simulation setup

Our agent-based simulation, implemented using the NetLogo platform, uses the following algorithm outline:

0. Determine the initial strategies for each agent.
1. Each agents plays UG with each of his neighbors both as player 1 and 2.
2. Agents update their strategies according to their fitness levels.
3. Collected fitness is discounted.
4. Repeat from 1 (unless stop condition is true).

Agents are located on a two-dimensional grid that connects at the edges, a torus. The distributions of initial strategies are independent for the two player positions, and follow a normal distribution with the mean and standard deviation as parameters.

Payoffs or fitness values are stored separately for each player position. Past fitness is discounted according to the discount rate parameter.

$$U(n) = \Sigma_{i=1}^{n} r^{n-i} u(i)$$

$U(n)$ and $u(n)$ representing total and current fitness in period n, the discount factor being r.

There is no memory of neighbors, that is, no agent realizes it is repeatedly playing with the same agents; thus, they can not differentiate their strategies depending on the agent they are playing with.

We employ two types of strategy updating. In case of the first, agents try to imitate their best-performing neighbor. They look at the collected utility for both player roles of each of their neighbors on the grid each turn. If the collected utility of their best neighbor for a certain role exceeds their own by at least as much as the adjustment cost parameter, they pay the adjustment cost, and move their strategies closer to that of the best neighbor's. The adjustment cost parameter thus serves a dual role: it makes intentional strategy change costly and discourages imitation when the expected utility gain would be too low.

The second type of strategy updating is rather different: agents have two „memory" vectors for the two roles, containing their used strategy, and the utility gained *in that turn*. Agents look at the maximum gained utility as player 1 or 2; if this exceeds their currently gained utility by the adjustment cost, they switch to their old strategy, otherwise, they stick with their current.

The strategies in both roles are also affected my random drift, or mutation. Mutations have an expected value of 0 and a fixed standard deviation, and are independent for strategies as player 1 or 2.

5 Results

5.1 Imitate-best-neighbor

When there are no adjustment costs, the behavior of the system is quite simple for all combinations of the relevant parameters, i.e. starting distributions, discount rate, and mutations. Any initial strategy diversity is quickly wiped out, average offers and average thresholds move within 1 percentage point to each other, and slowly but consistently random mutation drives them both down to complete exploitation. This result is hardly surprising, the dynamics being similar to the evolutionary game-theoretic setting mentioned in Section 3.

With the introduction of adjustment costs (c), the systems gains some impetus to move away from the subgame-perfect equilibrium. As Fig. 1 shows, adjustment costs create a lower boundary for acceptance thresholds, and a buffer between offers and these. Very high adjustment costs (up from 20% of the total maximum gain in one game) produce results that are not dissimilar to those in behavioral experiments, according to three proxies, long-term average offers, acceptance thresholds and percentage of accepted offers. Table 1 demonstrates that in these scenarios, the discount rate might have a non-monotonous effect on behavior, though very high discount rates indeed induce more short-sighted behavior - lower offers - , while very low ones increase average offers. Thresholds are affected inversely, very high discount rates generating lower thresholds ceteris paribus.

On the whole, while the results are promising, the necessary adjustment costs (min. 20% of the one-shot game, 30% too achieve the $50-50$ split norm of western societies)to create behaviorally plausible behavior can be regarded as being too high.

5.2 Memory

In this scenario, agents no longer look at their best-performing neighbor, instead, they choose one of their strategies that they remember to have been successful. As in the case of imitation, we see that the overall behavior of the system does not

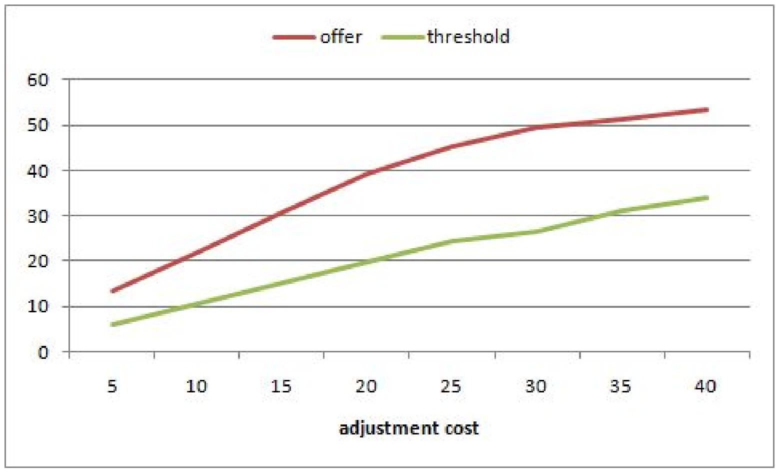

Fig. 1. Effect of adjustment costs on the average offers and acceptance thresholds in the long run.

c	r=1	r=3	r=5	r=10	r=25
5	15.1	13.4	14.5	12.2	11.1
10	23.4	21.2	22.4	22.0	20.0
15	29.7	31.4	35.0	28.6	28.7
20	40.5	40.6	37.4	39.8	38.4
25	45.8	45.4	46.1	44.3	44.3
30	52.9	48.4	52.4	46.9	47.7
35	49.6	50.6	51.0	51.5	54.3
40	55.4	52.7	54.9	51.8	51.7

Table 1. Non-monotonous effect of the discount rate on offers.

essentially depend on the initial conditions. After initial searching, the average of offers is consistently higher than the average threshold.

We have found that the dynamics of the system show remarkable robustness for a variety of parameter values. Once adjustment costs are introduced, the mutation parameter, the discount factor and, in fact, adjustment costs themselves have hardly any impact the system.

This does not, however, mean that under memory-based updating, system behavior is simple. On the contrary: not only is it polymorphic, it is much more stable over time than imitate-best-neighbor updating. Visually identifiable segments of the grid make aggressive offers over long intervals (several hundred or thousand periods), others are very generous. In each region, offers and acceptance thresholds adapt to each other. Nonetheless, in our experiments, the system does stabilize, showing gradual but continual changes for up to $2 * 10^6$ periods. The standard deviations of the strategies and fitness values do not diminish over time.

This behavior is difficult to characterize mathematically. The length of memory vectors, however, had clearly an important effect. It seems that our three proxies - long-run average offers, thresholds and acceptance rates - fall into intervals that conform with most experimental studies (i.e. $40-50\%$, $15-25\%$, $75-90\%$ if the memory vector for offers is about one order of magnitude shorter than the one for acceptance thresholds.

6 Conclusion

Our findings have two principal aspects. Firstly, memory-based strategy updating seems to outperform imitative learning for plausible values of the adjustment cost parameter. We interpret this an *internal* emergence of norms: altruistic preferences can thus be supported by intrinsic, rather than extrinsic reasons, which might be an important point for developmental psychology. This result could be corroborated by designing experiments where subjects could pay some of their endowments to observe the UG behavior of their peers (reported acceptance thresholds and acquired payoff). However, such experiments might face serious methodological problems of their own: the importance or lack thereof of imitation will be hard to measure, since it is even subjectively difficult to differentiate between external and internal reasons for updating a strategy. Furthermore, paying for observing others can only serve as a rough proxy for adjustment costs, since expectations of other player's rationality might very greatly (and rightly so). Another type of UG experiments might force players to stick to their offers and acceptance thresholds, unless they paid a pre-defined sum. However, this might just increase the prevalence of culturally focal offers and acceptance rates.

The other point worth to be emphasized is the proposed difference in the length of the memory vector for first- and second-player roles. This is in agreement with our intuition that first-player offers are prone to be defined more by strategic calculations, whereas moves as player 2 are chiefly determined by moral considerations. It might be impossible to test these hypotheses directly, but indicates that introducing the analysis of *qualitative* responses of subjects in behavioral experiments might be necessary.

References

Binmore, Ken: Does Game Theory work? The Bargaining Challenge, MIT Press, 2006

Costa-Gomez, M., Zauner, K.G.: Ultimatum Bargaining Behavior in Israel, Japan, Slovenia, and the United States: A Social Utility Analysis. Games and Economic Behavior 34 (2001), 238–270.

Fehr, E., Schmidt, K. M.: A Theory of Fairness, Competition, and Cooperation. The Quarterly Journal of Economics 114 (1999), 817–868.

Güth, W., Schmittberger, R., Schwarze B.: An Experimental Analysis of Ultimatum Bargaining. Journal of Behaviour and Organization 3 (1982), 367-388.

Heinrich, J., Boyd, R., Bowles, S., Camerer, C. F., Fehr, E., Gintis, H.: Foundations of Human Sociality, Oxford University Press, 2004.

Heinrich, J., Boyd, R., Bowles, S., Camerer, C. F., Fehr, E., Gintis, H., McElreath, R.: Overview and Synthesis, in: Heinrich et al. (2004).

Roth, A. E., Vesna, P., Okuno-Fujiware, M., Zamir, S.: Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study, American Economic Review 81 (1991), 1068–95.