


“No Switchbacks: Rethinking Aspiration-Based Dynamics
in the Ultimatum Game”

by

Jeffrey Paul Carpenter and Peter Hans Matthews

June, 2002

MIDDLEBURY COLLEGE ECONOMICS DISCUSSION PAPER NO. 02-18



DEPARTMENT OF ECONOMICS
MIDDLEBURY COLLEGE
MIDDLEBURY, VERMONT 05753

<http://www.middlebury.edu/~econ>

**NO SWITCHBACKS: RETHINKING ASPIRATION-BASED
DYNAMICS IN THE ULTIMATUM GAME**

Jeffrey Paul Carpenter

Peter Hans Matthews

Department of Economics
Middlebury College
Middlebury, Vermont 05753

Department of Economics
Middlebury College
Middlebury, Vermont 05753

jpc@middlebury.edu

peter.h.matthews@middlebury.edu

First Draft: 13 May 2001
Current Version: 14 June 2001
Comments Welcome

Abstract: The stylized facts of ultimatum bargaining in the experimental lab are that offers tend to be near an equal split of the surplus and low, near perfect, offers are routinely rejected. Binmore *et al* (1995) use aspiration-based evolutionary dynamics to model the evolution of fair play in a binary choice version of this game, and show that incredible threats to reject low offers persist in equilibrium. We focus on two possible extensions of this analysis: (1) the model makes assumptions about agent motivations (aspiration levels) and the structure of the game (binary strategy space) that have not yet been tested experimentally, and (2) the standard dynamic is based on the problematic assumption that unhappy agents who switch strategies may end up using the same strategy that was just rejected. To examine the implications of not allowing agents to “switch back” to their original strategy, we develop a “no switchback dynamic” and run a new, binary choice, experiment with induced aspirations. We find that the resulting dynamic predicts the evolution of play better than the standard dynamic and that aspirations are a significant motivator for our participants.

JEL Classification Numbers: C78, C91

Keywords: ultimatum game, learning, aspirations, switchbacks, replicator dynamics

No Switchbacks: Rethinking Aspiration-Based
Dynamics in the Ultimatum Game*

1. Introduction

Almost two decades have passed since Güth et al [1982] first documented a now familiar pattern in Ultimatum Game experiments - “fair” offers are more common, and unfair ones rejected more often, than is consistent with subgame perfection.¹ Evolutionary game theorists would later find this pattern to be less anomalous than their predecessors, however. In an influential paper, Binmore, Gale, and Samuelson [1995] (BGS) would show that when the shares of proposers and responders committed to pure strategies in a “miniature Ultimatum Game” (MUG) evolve on the basis of “replicator dynamics” (RD), there are *two* stable outcomes.² The first of these corresponds to the subgame perfect equilibrium - no proposers are fair, and all of their offers, fair or not, are accepted - but in the second, all proposers are fair, and a substantial (but indeterminate) number of responders would reject unfair offers. No less important, BGS were able to rationalize RD as a form of *social*

* The first author’s research was supported by NSF SES-0092953. We thank Larry Samuelson, Herb Gintis, Carolyn Craven and Corinna Noelke for their comments on a previous draft.

¹ In the Ultimatum Game, the first mover or “proposer” offers a division of some finite “pie” to the second mover or “responder,” who either accepts or rejects this offer. An accepted division is then implemented, but a rejected one leaves both with nothing.

² The project of explaining laboratory behavior using evolutionary and other (best response, for example) dynamics has also been taken up by Van Huyck, Cook and Battalio [1994] and Friedman [1996], among others.

evolution based on “aspiration-based learning.”³

Our own contribution follows from three observations about these results. First, while the experimental evidence is consistent with the presence of considerable fairness, there is less fairness than the “fair” RD outcome implies, with or without noise.⁴ This echoes the previous work of Van Huyck et al [1995], who found that the RD did not predict the observed behavior in two person “divide the dollar” games. Second, and on a related note, the binary choice version of the Ultimatum Game in BGS differs from that which experimental subjects play. And third, there is a possible lacuna in the BGS treatment of “disenchanted” players, who are sometimes assumed to “switch back” to their original strategies, no matter how disappointing these have proven. We find that these observations are connected: the amended RD described in the the next section are more consistent with the new evidence presented in our third section, based on an experimental design in which aspiration levels are induced. Furthermore, our empirical results support the use of simple aspiration-based learning as a plausible basis for social evolution, in contrast to the recent emphasis on rules-based approaches - see, for example, Stahl [2001] or Costa-Gomes and Weiszacker [2001].⁵

It will be useful, however, to first review the treatment of MUG in BGS. There

³ BGS [69] caution readers not to “place too much significance on the particular value of the equilibrium offer ... [since] ... different specifications ... can give different results.” Despite this, their rationalization for the RD remains both an appealing, and influential, one.

⁴ With more or less comparable noise in the two populations, the outcome in which all proposers are selfish, and no responder turns down a selfish offer becomes the unique rest point. When responders are noisier, there is a second stable rest point in which “almost all” proposers are fair. For more details, see BGS.

⁵ This said, the aspirations we induce are, by current theoretical standards, sim-

are two populations, proposers and responders, the members of which are matched at random each period to play the normal form game:

	<i>Accept</i>	<i>Reject</i>
<i>Fair</i>	2, 2	2, 2
<i>Selfish</i>	3, 1	0, 0

in which proposers must decide whether to offer a fair (equal) division of a pie of size 4 or demand most (3) of it, and it is assumed that fair offers are never rejected.⁶ Let the shares of fair and selfish proposers be denoted s_F^P and s_S^P , the shares of responders who accept and reject unfair offers s_A^R and s_R^R , and suppose that time is marked in discrete intervals of length Δ . Suppose, too, that each period, a fraction Δ of proposers and responders evaluate their current performance, and that this evaluation is based on a comparison of their current payoff with some “aspiration,” the value of which is drawn from a uniform distribution over $[a_L, a_H]$, where, in this particular framework, $a_L \leq 0$ and $a_H \geq 4$. When a proposer’s payoff exceeds her aspiration, for example, she retains her current strategy, but when it falls short, she is assumed to “change” it, where the likelihoods that strategies are ple ones. We do not allow these aspirations to evolve over time, for example, or consider peer influence. For an overview of recent developments, see Bendor, Mookherjee and Ray [2000].

⁶ As it turns out, 41 of the 377 (11%) of the fair offers we observed in our experiment were rejected. We should note, however, that 40 of these 41 came in our first session, and that three disenchanted responders with aspirations close to two were responsible. Dan Goldman, a student and participant in the experiment, later identified two possible reasons for the rejection of fair offers: “spite” on the part of those who would never realize their induced aspiration, and a preoccupation with relative outcomes on the part of those well above them.

adopted are equal to their current shares in the population. (This also assumes, of course, that the proposer either observes the composition of her own population or perhaps samples and imitates.) We use quotation marks because these changes are sometimes more nominal than real: when all of the proposers are fair, for example, even the disenchanted must remain so.

It follows, therefore, that that the shares of fair proposers will evolve as:

$$s_F^P(t + \Delta) = s_F^P(t) - \Delta p_F^P(t) s_F^P(t) + s_F^P(t) [\Delta p_F^P(t) s_F^P(t) + \Delta p_S^P(t) s_S^P(t)]$$

where $p_F^P(p_S^P)$ is the likelihood that a fair (selfish) proposer falls short of her aspiration. The second term on the right hand side is the number of fair proposers who become disenchanted in the current period, and the third is the product of the total number of unsatisfied proposers, fair and unfair, and the current share of fair proposers, or the number of “new” fair proposers. Since $p_F^P(t) = (a_H - \pi_F^P(t))/(a_H - a_L)$ and $p_S^P(t) = (a_H - \pi_S^P(t))/(a_H - a_L)$, where $\pi_F^P(t)$ and $\pi_S^P(t)$ are the current payoffs to fair and selfish proposers, it follows that:

$$\frac{s_F^P(t + \Delta) - s_F^P(t)}{\Delta} = \left(\frac{1}{a_H - a_L}\right) s_F^P(t) (\pi_F^P(t) - \bar{\pi}^P(t)) \quad (2.1)$$

where $\bar{\pi}^P = s_F^P(t) \pi_F^P(t) + s_S^P(t) \pi_S^P(t)$ is the mean for all proposers.⁷ Likewise, for responders, we have:

$$\frac{s_A^R(t + \Delta) - s_A^R(t)}{\Delta} = \left(\frac{1}{a_H - a_L}\right) s_A^R(t) (\pi_A^R(t) - \bar{\pi}^R(t)) \quad (2.2)$$

As $\Delta \rightarrow 0$, (2.1) and (2.2) comprise a scaled version of the continuous time RD:

$$\begin{aligned} \dot{s}_F^P(t) &= \left(\frac{1}{a_H - a_L}\right) s_F^P(t) (\pi_F^P(t) - \bar{\pi}^P(t)) \\ \dot{s}_A^R(t) &= \left(\frac{1}{a_H - a_L}\right) s_A^R(t) (\pi_A^R(t) - \bar{\pi}^R(t)) \end{aligned}$$

⁷ Since it is well known the vector field is invariant under RD, we do not consider the behavior of $s_S^P = 1 - s_F^P$.

The particular form of the RD in this case is:

$$\begin{aligned}\dot{s}_F^P(t) &= (1/4)s_F^P(t)(1 - s_F^P(t))(2 - 3s_A^R(t)) \\ \dot{s}_A^R(t) &= (1/4)s_A^R(t)(1 - s_A^R(t))(1 - s_F^P(t))\end{aligned}\tag{3}$$

for $a_L = 0$ and $a_H = 4$.

As alluded to above, there are two stable outcomes under (2.3): $(s_F^P(t) = 0, s_A^R(t) = 1)$ is locally asymptotically stable, and the connected set $(s_F^P(t) = 1, 0 \leq s_A^R(t) < 2/3 - \epsilon)$ is Liapunov stable.

2. A Modified Aspiration Model

We introduce two modifications to the treatment of social evolution in BGS. First, those with unrealized aspirations are now required to adopt *new* strategies: the disenchanted cannot return or “switch back” to their initial choices, no matter how common these are. (This does not preclude switches and, if and when there is disappointment in future rounds, switchbacks.) With just two strategies available to the members of each population, the *transition function* is a simple one, and its information requirements minimal: fair proposers who fall short of their aspirations must become selfish ones, for example, and do not need to know the composition of either population to do so. In discrete time, the proportions of fair and selfish proposers will therefore evolve as:

$$\begin{aligned}s_F^P(t + \Delta) &= (1 - \Delta p_F^P(t))s_F^P(t) + \Delta p_S^P(t)s_S^P(t) \\ s_S^P(t + \Delta) &= (1 - \Delta p_S^P(t))s_S^P(t) + \Delta p_F^P(t)s_F^P(t)\end{aligned}$$

It follows that $\sum_j s_j^P(t + \Delta) = \sum_j s_j^P(t)$, so that $\sum_j s_j^P(0) = 1 \rightarrow \sum_j s_j^P(t) = 1$ for each t - that is, population shares will never “wander off the simplex” - so that we can substitute $1 - s_F^P(t)$ for $s_S^P(t)$ and limit attention to the first of these laws of motion:

$$s_F^P(t + \Delta) - s_F^P(t) = -\Delta p_F^P(t)s_F^P(t) + \Delta p_S^P(t)(1 - s_F^P(t))$$

Likewise, for responders, we have:

$$s_A^R(t + \Delta) - s_A^R(t) = -\Delta p_A^R(t) s_A^R(t) + \Delta p_R^R(t) (1 - s_A^R(t))$$

Combining these and letting $\Delta \rightarrow 0$ produces:

$$\begin{aligned} \dot{s}_F^P(t) &= -p_F^P(t) s_F^P(t) + p_S^P(t) (1 - s_F^P(t)) \\ \dot{s}_A^R(t) &= -p_A^R(t) s_A^R(t) + p_R^R(t) (1 - s_A^R(t)) \end{aligned} \tag{3.1}$$

These constitute the “no switchback dynamics” or NSD for MUG.

The connections between standard notions of evolutionary equilibrium and the stable rest points of evolutionary dynamics, a characteristic feature of the RD, vanish under the NSD. For example, if the proposers who make selfish offers and the responders who turn down these offers are *ever* dissatisfied, the shares that correspond to the perfect equilibrium of MUG will not even be a rest point under NSD, let alone a stable one. Furthermore, this condition will (almost) never be satisfied: if more than a “small” subset of the responder population aspires to more than one, for example, the proportion of those who reject selfish offers must soon rise. For similar reasons, the set of locally stable states in which no proposer is selfish and two thirds or fewer of responders would agree to an unequal split, a subset of the Nash equilibria of MUG, will not be an attractor either. To the extent that the experimental evidence is consistent with a limit point well inside the interior of the phase space, however, a point somewhere in the northeast section of the state space, this is a strength, not a weakness.

We are not the first, of course, to suggest that non-Nash outcomes can be stable. Drawing on the work of McKelvey and Palfrey [1995], for example, Chen, Friedman and Thisse [1997, 37] define a variant of the quantal response equilibrium, the “boundedly rational Nash equilibrium” or BRNE, in “which the strategy of each player is a vector of discrete choice probabilities which is a random choice [modified multinomial logit] best response to the choice probabilities of the remaining

players.”⁸ Chen *et al* show that all finite games have BRNEs and that under broad conditions, fictitious play will converge to a unique BRNE. As shown below, the stable rest point of the NSD corresponds to a BRNE of MUG in which proposers and responders are both “more rational” than consistent with, for example, Luce’s [1959] notion of “probabilistic choice.”

As these observations hint, the distribution of aspiration levels matters more under NSD. Under the alternative RD, for example, as a_H rises - that is, as the numbers of proposers and responders who fall short of their respective aspirations increases - the pace of evolution is affected, but its character is not. That is, the solution orbits are the same, but velocities on these orbits are not. Under the NSD (1), on the other hand, this increase would push the interior limit point(s) to $(1/2, 1/2)$, for intuitive reasons: in discrete time, $\Delta s_F^P(t)$ fair proposers, *all* of those who evaluate their performance in a particular period, will become selfish, while all Δs_S^P of the selfish ones who self-evaluate will become fair, and these flows will not offset one another unless $s_F^P = s_S^P = 1/2$.

This leads to our second modification. BGS [87] mention differences in the distribution of aspiration levels as a natural extension of their model, but also note, in effect, that with switchback, it is the basins of attraction, not the attractors themselves, that are affected. We shall allow for differences in the (still uniform, however) distribution, too, but because the limit points of the NSD are sensitive to these, a selection criterion is called for. The levels induced in our subjects, for example, were consistent with the requirement that *no one is bound to be satisfied or dissatisfied* in all possible states of the world. In more practical terms, we suppose that proposers draw, or have drawn for them, from $U[0, 3]$, and responders from $U[0, 2]$.

It follows that under these conditions, $p_F^P(t) = 1/3$, $p_S^P(t) = 1 - s_A^R(t)$, $p_A^R(t) =$

⁸ We thank Larry Samuelson for bringing this connection to our attention.

$(1/2)(1-s_F^P(t))$ and $p_R^R(t) = 1-s_F^P(t)$. One third of the fair proposers who reconsider their situation in a particular period, for example, will become selfish, no matter what the characteristics of the responder population. This is the expected result: fair proposers receive 2 for certain, and with a uniform distribution of aspirations between 0 and 3, one third will not be satisfied with this. For similar reasons, the observation that while responders' "likelihood of disappointment" varies with the number of fair proposers, the likelihood that those who turn down unequal splits is twice that of those who do not is also more or less intuitive.

Substitution for the p_j^i 's and π_j^i 's in (3.1) leads, after further simplification, to the precise form of the NSD:

$$\begin{aligned} \dot{s}_F^P(t) &= -\frac{1}{3}s_F^P(t) + (1-s_F^P(t))(1-s_A^R(t)) \\ \dot{s}_A^R(t) &= (1-s_F^P(t))(1-\frac{3}{2}s_A^R(t)) \end{aligned} \quad (3.2)$$

The associated phase diagram is depicted in Figure 1. There is a single, asymptotically stable, equilibrium, $(s_F^P = 1/2, s_A^R = 2/3)$, in which half of the offers are fair, and two thirds of all unfair offers are accepted.⁹ This prediction is sharper than that obtained under the RD, and more consistent, or at least no less consistent, with the experimental evidence (Roth 1995). It is also a more "turbulent" equilibrium, another characteristic of the experimental data: one third of *all* proposers, fair and selfish, switch each period, as do half of the responders who reject unfair offers and one quarter of the responders who do not.¹⁰ We observe, too, that this equilibrium

⁹ The trace of the relevant Jacobian, evaluated at this point, is equal to $-17/12 < 0$, the determinant is $1/2 > 0$, and since $(17/12)^2 > 4(1/2)$, the eigenvalues are negative and unequal, so that the rest point is locally asymptotically stable.

¹⁰ In the sequential bargaining experiment elaborated on in Carpenter [2000], sixty-six percent of first movers change their offers from period to period. This fraction seems even larger given the central tendency of offers was not significantly

is invariant with respect to common affine transformations, so that the conversion of experimental monetary units into dollars, or the use of rewards for participation, have no effect, provided the endpoints of the distributions of aspirations are also transformed.

If these proportions are instead (re)interpreted as mixed strategy profiles for a one shot version of MUG, this equilibrium corresponds to a BRNE in which responders' "degree of rationality" μ_R is $\ln 2/\ln 1.5$, but proposers' μ_P is indeterminate.¹¹ On the continuum of possible μ -values, 0 is associated with equal choice different from period to period. It should be noted, however, that the turbulence can be "tuned down" in our model if it assumed that proposers and responders evaluate their situation less frequently.

¹¹ Letting the mixed strategies be $(s_F^P, 1-s_F^P)$ and $(s_A^R, 1-s_A^R)$, the two conditions for a BRNE are:

$$s_F^P = \frac{2^{\mu_P}}{2^{\mu_P} + (3s_A^R)^{\mu_P}}$$

and

$$s_A^R = \frac{(1 + s_F^P)^{\mu_R}}{(1 + s_F^P)^{\mu_R} + (2s_F^P)^{\mu_R}}$$

where μ_P and μ_R are the aforementioned "degrees of rationality." For $(s_F^P = 1/2, s_A^R = 2/3)$, these will be satisfied for $\mu_R = \ln 2/\ln 1.5$ and all μ_P . The value of μ_P is indeterminate because when $s_A^R = 2/3$, the expected values of fair and selfish offers are equal and there is no premium for "more rational" behavior. Suppose, however, that responders sometimes tremble when confronted with a fair offer, and let the expected outcome under *(fair, reject)* be $(2 - \delta, 2 - \delta)$. It is then not difficult to show that as $\delta \rightarrow 0$, $s_F^P \rightarrow 1/2$, $s_A^R \rightarrow 2/3$, $\mu_R \rightarrow \ln 2/\ln 1.5$, but $\mu_P \rightarrow 3$. It is for this reason that we conclude both proposer and responder are more rational than Luce's [1959] probabilistic choosers.

probabilities, 1, with Luce’s [1959] notion of probabilistic choice, and ∞ , with “full rationality,” from which we conclude that responders and, for reasons outlined in the footnote, proposers are more rational than, for example, probabilistic choosers would be. It is tempting, therefore, to view the NSD as a selection mechanism for BRNEs.

Last, and in anticipation of some of our experimental results, observe that initial states “close” to the northeast corner of state space ($s_F^P = 1, s_A^R = 1$) are not “pulled across the top,” to the point corresponding to the subgame perfect equilibrium, as in BGS, but rather into the interior of the space, consistent with the behavior we observed.

Intuition suggests that the introduction of some “decision noise” should not have much effect on our already turbulent equilibrium. To verify this, suppose that a fraction θ^P of proposers, and θ^R of responders, commit self-evaluation errors - that is, a share θ^P of proposers, both fair and unfair, who should be satisfied conclude otherwise, and then switch, and that the same share who should be dissatisfied fail to do so, and likewise for responders. In general terms, the modified NSD are:

$$\begin{aligned}
 \dot{s}_F^P(t) &= -((1 - \theta^P)p_F^P + \theta^P(1 - p_F^P))s_F^P(t) \\
 &\quad + ((1 - \theta^P)p_S^P + \theta^P(1 - p_S^P))(1 - s_F^P(t)) \\
 \dot{s}_A^R(t) &= -((1 - \theta^R)p_A^R + \theta^R(1 - p_A^R))s_A^R(t) \\
 &\quad + ((1 - \theta^R)p_R^R + \theta^R(1 - p_R^R))(1 - s_A^R(t))
 \end{aligned} \tag{3}$$

The effects of such noise on the equilibrium shares s_F^P and s_A^R are recorded in Table 1. The introduction of minimal noise ($\theta^P = 0.01, \theta^R = 0.01$) has almost no effect on the (still stable) equilibrium: the share of fair proposers rises, from 50 percent to 50.3, and that of responders who reject unfair offers falls, from 66.7 percent to 66.2. Since the rest point is hyperbolic,¹² such “persistence” is more or less expected. The surprise, perhaps, is that as the level of noise in both populations increases a

¹² That is, the relevant Jacobian has no zero or purely imaginary eigenvalues. For

substantial amount, to, say, 10 percent, the share of fair proposers rises just a little more, to 52.2 percent, while the proportion of responders who reject unfair offers falls, also a little bit, to 62.4 percent. In more general terms, the equilibrium share s_F^P (s_A^R) is a decreasing (increasing) function of θ^P , and an increasing (decreasing) function of θ^R with, in a loose sense, responder noise the more decisive influence. There is perhaps a loose parallel here to BGS, who find that responders must be “noisier” than proposers for the perfect equilibrium not to become the unique limit point.

Last, we conjecture that a more elaborate NSD that accounted for the small size of our experimental populations - in the third session, for example, there are four proposers and four responders, so that each “disappointment” has a substantial effect on population composition - would produce final states even further in the northeast quadrant.

3. Experimental Evidence

To examine whether the standard model of aspiration-based social learning developed in BGS, Weibull [1995], and Vega-Redondo [1996] or the current model based on the no switchback principle best describes behavior in MUG, we ran four computerized experimental sessions. Fifty students, representing various majors, were recruited from the undergraduate population at Middlebury College. The experiment was computerized and based on the ultimatum minigame presented in BGS (see above) with payoffs stated in terms of experimental monetary units, EMUs, that were translated into cash at the end of the experiment. Proposers were asked to choose between a selfish proposal, 3EMUs for the proposer and 1EMU for the responder, and a fair proposal 2EMUs for each player. Responders were then given the opportunity to accept or reject the proposal.

details, see, for example, Glendenning [1994].

Because we are interested in the ending state of a social learning process, we were careful to take precautions to prevent any possible endgame effects. We hypothesized that subjects may tend to disregard the history of play when near the end of a session, especially if they have no chance of meeting their aspiration, and therefore the instructions were worded vaguely stating that the experiment would proceed for as many rounds as time permitted. An hour and a half was allocated for each session, but after piloting the procedures in an informal setting, we discovered by debriefing participants that many lost interest after round 25. With this in mind, each session ran for 20 rounds, which took about an hour. Further, participants maintained the same role for the entire experiment, but were randomly reassigned a new partner after each round.

There is one other noteworthy design feature of the experiment. To be as fair as possible to the aspiration-based model, we decided to induce aspiration levels in our participants. We accomplished this by modifying the procedures used in Siegel and Fouraker [1960]. At the beginning of each session, participants were randomly assigned an aspiration level from an interval that depended on the participants role in the experiment (recall the above discussion of asymmetric aspiration intervals). Proposer aspiration levels were drawn from the interval $(0, 3)$ and responder aspiration levels were drawn from $(0, 2)$. This asymmetry is appropriate given responders could never earn more than 2EMUs in a round. To make the aspiration level salient, participants were told that if their average earnings at the end of the experiment met or exceeded their aspiration level, they would be given the chance to double their earnings. When paying the participants at the end of the experiment, anyone whose average earnings exceeded their aspiration level was given a die to roll. If the die landed with either a 1 or a 2 up, the participants earnings were doubled.

Table 2 summarizes the starting and ending states for each session. Three of the four sessions start in the interior of the strategy space and, taken together, the four

sessions provide very different initial conditions for the experiment.¹³ Just as our phase diagram, Figure 1, sweeps the entire strategy space when examining potential paths to equilibrium, the differences in starting states allow us to be confident that our experimental analysis is not limited to local behavior in one limited region of the simplex. One can also see that the final states vary by session, but tend to stay in the northeast quadrant of the simplex as predicted by the no switch-back model of learning.

Table 2 also lists the average behavior over all twenty rounds, the average aspiration level drawn by proposers and responders, the fraction of players in each session who reach their aspiration, and the number of participants per session. Interestingly, aspirations and meeting aspirations appear to correlate with average play in the experiment which is evidence that our aspiration-inducement procedure was successful. More specifically, in accordance with subgame perfect play, higher proposer aspiration levels tend to reduce the number of fair offers and high responder aspirations appear to yield more acceptances. Further, participants seem to also respond to the size of the session.¹⁴ Large sessions tend to stay closer to the center of the simplex while our smallest session, 3, starts, ends, and remains close to the all fair, all accept vertex. We analyze these observations in more detail below.

Figure 2 presents the evolution of play. For each session we map the path

¹³ It is possible that differences in initial conditions reflect differences in mean aspirations across sessions, but the relationship is not obvious *a priori*, and difficult to test with just four observations/sessions. We suspect, however, that as the mean aspiration of proposers rises, there should be fewer fair offers in the first round, and that of responders rises, there will be more “spite.”

¹⁴ Friedman [1996] also mentions group size effects on the convergence to behavioral equilibria.

taken on the strategy simplex. Numbers indicate the transitions in the evolution of play in chronological order. Clearly, play never starts, ends, or even approaches the subgame perfect equilibrium of the ultimatum minigame. However, we are more interested in whether play proceeds in the direction of the perturbation-induced equilibrium calculated in BGS, or if play remains in the interior of the simplex as predicted by the no switch-back model.

With the exception of two transitory states in session one (rounds 15 and 16) which approach the BGS equilibrium ($\approx 1, 2/3$), play either remains in the interior of the simplex or moves to a state on the border where everyone offers an equal split and all offers are accepted (sessions two and three). We conclude that rational, error-prone behavior does not describe play in this experiment. Note however, the majority of play cycles in the northeast quadrant of the strategy space as predicted by the no switch-back model of social learning.

As mentioned above, the aspiration levels assigned to our participants seemed to play a significant role in the experiment. We examine this in more detail by regressing players choices on their aspiration levels, deviations between aspiration levels and average payoffs, and the size of the session. Because we wish to isolate the role of aspiration in players decisions, we control for cross-sectional differences and learning by employing random effects. Table 3 summarizes the results of this analysis.

One might expect that even though the instructions clearly stated that individual choices would never be revealed, players may feel more anonymous in big groups. If anonymity causes more self-interested play, we expect more greedy proposals and more acceptances in our larger sessions. At the same time, if the aspiration levels we induced were salient, we also expect (as shown by Siegel and Fouraker [1960]) that player aspirations will tend to crowd out other-regarding feelings and therefore retard the evolution of play towards the all fair, all accept vertex. If our hypotheses are correct, then our large sessions with high aspirations provide the aspiration-

based model with its best chance of success.

Starting with proposer choices, we see from Table 3 that the sign on the session size coefficient is in the predicted direction, larger groups yield fewer fair offers, but the effect is insignificant. However, proposers react strongly to their aspiration level. Higher aspiration levels significantly reduce the likelihood of fair offers, even controlling for the deviation between a proposer's current average payoff and their aspiration level. We conclude that proposers are driven by the absolute level of their aspirations, but not necessarily the payoff implications of these aspirations (i.e. the deviation between aspirations and average payoffs does not affect behavior).

The anonymity of a session does affect the choices of responders. Contrary to our predictions about increased self-interest in large groups, responders are significantly more likely to reject an offer of given size in such groups. This suggests that anonymity triggers more, not less, spite, a result similar to Bolton and Zwick [1995]. Further, responders are more likely to accept each offer when they draw high aspiration levels. Similar to proposers, the deviation of a responder's current average payoff and the aspiration level works in the hypothesized direction (higher deviations make responders more likely to accept), but is not a significant influence.

We end our discussion of the experiment by noting that aspiration-based models of social evolution make specific predictions about behavior that we can test using in our data. We would expect players to be more likely to change strategies when their average payoff falls below their aspiration level. The results in Table 4 assess this prediction. The variable Aspiration Deviation is the difference between a player's current average payoff and his or her aspiration level. Equation one confirms that aspirations cause players to switch strategies. More specifically, unhappy players (i.e. Aspiration Deviation < 0) are more likely to switch than players who have met or surpassed their aspiration level. Equation two demonstrates that the effect of deviations from aspirations is attenuated by how long the game has been played. It appears that players are less likely to switch as the game proceeds indicating sessions

tend to settle onto a behavioral equilibrium or norm. Moreover, there doesn't appear to be any interaction between how much a player falls short (or is above) his or her aspiration level and how long the game has been played. That is, players neither seem to panic and switch more or relax and switch less as the game progresses suggesting the fact that the end of the experiment was unknown prevents any unmodeled endgame behavior. Finally, controlling for aspiration deviations and how long the game has been played, proposers are significantly more likely to switch strategies than responders are indicating spite dominates greed in our experiment.

4. Conclusion

Our purpose was twofold in this paper. First, we were interested in developing a model of the evolution of play in the ultimatum game that (1) was based on the assumption that dissatisfied players switched strategies for certain, and (2) required that players draw aspirations from the set of available game payoffs. Our hope was that such a model would predict outcomes better than the standard aspiration-based replicator dynamic. Second, to assess the success or failure of our modifications to the standard evolutionary dynamic, we were also interested in running an experiment designed to replicate the conditions necessary for an aspiration-based model to predict; namely, we decided to run a binary choice version of the game and induce aspirations in our participants.

Concerning our first objective, we find that a model of social evolution wherein agents abandon strategies that produce payoffs falling short of their aspirations for sure results in a unique asymptotically stable attractor much closer to the center of the strategy space than equilibria under the standard (noisy) dynamic. This result is noticeably more consistent with existing experimental results. That is, in most repeated versions of the ultimatum game, each period generates both fair and selfish offers and selfish offers are rejected with non-vanishing probability (e.g. Prasnikar and Roth [1992]). Further, if we allow for asymmetries in the distribution of aspirations that are role-dependent, our equilibrium moves even closer to actual

play.

We summarize the results of our experiment as follows. Regression analysis (Table 3) suggests that our aspiration manipulation was successful. In our experiment induced aspirations have the predicted effect of pushing play in the direction of the subgame perfect equilibrium (i.e. fewer fair offers and more acceptances), but these forces are not strong enough so that the subgame perfect equilibrium was realized in any session. Instead, group size tends to attenuate the effect of aspiration on responders (i.e. responders are emboldened to reject in larger, more anonymous settings). The end result is best viewed in Figure 2 - controlling for aspiration levels and group size, the no switchback dynamic is a better predictor of the evolution of play than either the subgame perfect equilibrium or the connected set of equilibria in which all offers are fair. Lastly, our experiment indicates that aspiration-based models are a sensible way to think about social evolution: our second set of regressions (Table 4) demonstrates that players make strategic choices based on deviations from induced aspirations.

These results suggest two future directions for research in this area. First, from an experimental point of view, we were surprised by the magnitude of the effect of induced aspirations on the experimental outcomes. We speculate that inducing aspirations in other well understood game environments (e.g. public goods, or common pool resources) will also yield interesting results tractable by evolutionary models. Second, we are encouraged by our theoretical results which indicate that tailoring the standard story of social evolution to better fit a given situation yields results more consistent with observed behavior. Other manipulations are obvious, but we will mention one we feel is particularly interesting. We suspect that an even better way to think about aspirations is that they evolve with the history of play, as in Karandikar *et al* [1998]. In future work, we plan to explore the implications of endogenous aspirations without switchbacks, and hope to report our results in the near future.

References

- Bendor, J., D. Mookherjee and D. Ray [2000], Aspiration-based Reinforcement Learning in Repeated Games: An Overview, mimeo.
- Binmore, K. G., J. Gale and L. Samuelson [1995], Learning to be Imperfect: The Ultimatum Game, *Games and Economic Behavior*, 8: 56-90.
- Bolton, G. and R. Zwick [1995], Anonymity versus Punishment in Ultimatum Bargaining, *Games and Economic Behavior*, 10: 95-121.
- Carpenter, J. [2000], Bargaining Outcomes as the Result of Co-ordinated Expectations, mimeo.
- Chen, H-C., J. W. Friedman and J.-F. Thisse [1997], Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach, *Games and Economic Behavior*, 18: 32-54.
- Costa-Gomes, M. and G. Weizsacker [2001], Another Look at Decision Rules in Games: An Experimental Study, mimeo.
- Friedman, D. [1996] Equilibrium in Evolutionary Games: Some Experimental Results, *Economic Journal*, 106: 1-25.
- Glendenning, P. [1994] *Stability, Instability and Chaos: An Introduction to the Theory of Nonlinear Differential Equations*, New York: Cambridge University Press.
- Güth, W., R. Schmittberger and B. Schwarz [1982], An Experimental Analysis of Ultimatum Bargaining, *Journal of Economic Behavior and Organization*, 3: 367-388.
- Karandikar, R., D. Mookherjee, D. Ray and F. Vega-Redondo [1998], Evolving Aspirations and Cooperation, *Journal of Economic Theory*, 80: 292-331.
- Luce, R. D. [1959], *Individual Choice Behavior: A Theoretical Analysis*, Wiley: New York.
- McKelvey, R. D. and T. R. Palfrey [1995], Quantal Response Equilibria for Normal

- Form Games, *Games and Economic Behavior*, 10: 6-38.
- Prasnikar, V., and A. Roth [1992], Considerations of Fairness and Strategy: Experimental Data From Sequential Games, *Quarterly Journal of Economics*, 102: 865-888.
- Roth, A. [1995] Bargaining Experiments, in J. Kagel and A. Roth, eds, *The Handbook of Experimental Economics*. Princeton University Press: Press, 363-391.
- Samuelson, L. and J. Zhang [1992], Evolutionary Stability in Asymmetric Games, *Journal of Economic Theory*, 57: 363-391.
- Siegel, S., and L. Fouraker [1960], *Bargaining and Group Decision Making: Experiments in Bilateral Monopoly*, McGraw-Hill: New York.
- Stahl, D. O. [2001], Action Reinforcement Learning versus Rule Learning, mimeo.
- Van Huyck, J. J. Cook and R. Battalio [1994], Selection Dynamics, Asymptotic Stability, and Adaptive Behavior, *Journal of Political Economy*, 102: 975-1005.
- Van Huyck, J., R. Battalio, S. Mathur and P. Van Huyck [1995], On the Origin of Convention: Evidence from Symmetric Bargaining Games, *International Journal of Game Theory*, 24: 187-212.
- Vega-Redondo, F. [1996], *Evolution, Games, and Economic Behavior*, Oxford University Press: Oxford.
- Weibull, J. [1995], *Evolutionary Game Theory*, MIT Press: Cambridge.

Table 1: The Effect of Decision Noise on Equilibrium

		θ^R			
		0	0.01	0.10	0.25
θ^P	0	0.500,0.667	0.503,0.662	0.531,0.623	0.566,0.565
	0.01	0.500,0.667	0.503,0.662	0.530,0.623	0.564,0.565
	0.10	0.500,0.667	0.503,0.662	0.522,0.624	0.549,0.567
	0.25	0.500,0.667	0.501,0.662	0.512,0.624	0.528,0.569

Table 2: Summary of Play and Aspiration

	Session 1	Session 2	Session 3	Session 4
Start State	(0.55,0.55)	(0.83,0.83)	(1,1)	(0.33,0.50)
End State	(0.78,0.67)	(0.83,1)	(1,1)	(0.67,0.67)
Mean State	(0.77,0.62)	(0.76,0.83)	(0.93,0.95)	(0.62,0.75)
Mean Proposer				
Aspiration	1.22	1.54	0.60	2.41
Mean Responder				
Aspiration	1.56	0.78	1.40	1.77
Fraction Who Reach				
Aspiration	0.50	0.75	0.88	0.17
Participants	18	12	8	12

Table 3: The Effects of Aspiration and Session Size

	Proposers	Responders
Constant	2.12*** (0.78)	2.80*** (0.99)
Aspiration Level	-0.68*** (0.23)	0.75** (0.36)
Aspiration Deviation	-0.20 (0.15)	0.23 (0.30)
Session Size	-0.02 (0.05)	-0.37*** (0.07)
Proposal		4.07*** (0.43)

Notes: The dependent variables are 1=fair(0=unfair) for proposers and 1=accept (0=reject) for responders. Both regressions are random effect probits, where *, ** and *** denote significance at the 10, 5 and 1 percent levels.

Table 4: The Determinants of Switching

	(1)	(2)	(3)	(4)
Constant	-1.11*** (0.10)	-0.76*** (0.15)	-0.77*** (0.15)	-1.66*** (0.21)
Aspiration Deviation	-0.19*** (0.08)	-0.17** (0.08)	-0.26** (0.12)	-0.25** (0.12)
Round		-0.03*** (0.01)	-0.03*** (0.01)	-0.03*** (0.01)
Dev×Round			0.01 (0.01)	0.01 (0.01)
Proposer				1.15*** (0.16)

Notes: The dependent variables is =1 if (a) proposers switch strategies between rounds $t - 1$ and t , and (b) responders switch, given responder is considering the same offer as last period. All of the regressions are random effect probits, where *, ** and *** denote significance at the 10, 5 and 1 percent levels.




Figure 1: Direction Field for MUG Under NSD

Session One




Figure 2a: The Evolution of Play: Session One

Session Two




Figure 2b: The Evolution of Play: Session Two

Session Three




Figure 2c: The Evolution of Play: Session Three

Session Four




Figure 2d: The Evolution of Play: Session Four