Learning in Tournaments with Inter-Generational Advice

Ananish Chaudhuri
University of Auckland

Andrew Schotter
New York University

Barry Sopher
Rutgers University

Abstract

We study learning in a simulated tournament using an inter-generational framework. Here a group of subjects are recruited into the lab and play the stage game for 10 rounds. After his participation is over, each player is replaced by another player, his laboratory descendant, who then plays the game for another 10 rounds as a member of a fresh group of subjects. A particular player in generation t+1 can (1) see the history of choices by his generation t predecessor and (2) receive advice from that predecessor via free-form messages that generation t players leave for their generation t+1 successors. We find that the presence of advice makes a difference in that the experimental groups who get advice perform better – their decisions are closer to the Nash equilibrium – compared to a control group of subjects that plays the game with no recourse to such advice.

We are indebted to the Rutgers University Research Council for providing the funds for this study. We also thank Paul Brown for his assistance with organizing the data. The usual disclaimer applies.
Citation: Chaudhuri, Ananish, Andrew Schotter, and Barry Sopher, (2006) "Learning in Tournaments with Inter-Generational Advice." *Economics Bulletin*, Vol. 3, No. 26 pp. 1-16
Submitted: July 17, 2006. **Accepted:** September 20, 2006.
URL: http://economicsbulletin.vanderbilt.edu/2006/volume3/EB-06C90004A.pdf

1. Introduction

We study whether word-of-mouth advice from a predecessor to a successor in an inter-generational paradigm facilitates learning in a simulated tournament. While a particular game may have a clear equilibrium prediction, it is not always easy for boundedly rational agents to hone in on that equilibrium via a process of introspection as is often assumed in theoretical models of strategic behaviour. Rather agents need to engage in a considerable amount of learning or experimentation in order to arrive at the best response. Camerer (2003, Ch. 6) provides an overview of this line of work There is evidence that advice coming from a person with prior experience in the game has an impact on decision making in a wide variety of contexts. See Schotter (2003) for a review. We describe the constituent game first and then explain our inter-generational framework.

Each subject takes part in a two-person tournament similar to those described in greater detail in Bull, Schotter and Weigelt (1987) and Merlo and Schotter (1999). In Bull, Schotter and Weigelt (1987), randomly paired subjects in each round choose an integer called the "decision number" (*n*) between 0 and 100. This decision number has a quadratic cost associated with it of the form $C(n) = n^2/k$ where "*k*" is a constant. After the decision number is chosen by each subject, a random number is generated for each subject from a uniform distribution over the closed interval $[-a, +a]$. These two numbers – each player's decision number and the random number chosen for him – are added together to get the "total number" for a subject. Payoffs are determined by comparing the total numbers for each player in any pair. The player with the larger total number is awarded a larger prize M while the player with the smaller total is awarded a smaller prize, m where $M > m$. The payoff to a subject in any round is the difference between the prize won for that round (M or m) and the cost associated with the decision number that the subject chose for that particular round, i.e. the per period payoff is equal to $M(m) - n^2/k$. There is a trade-off in the choice of the decision number in that choosing a higher number increases the probability of winning the big prize M but that win comes at the expense of the higher cost and thus choosing a high number does not necessarily translate into a high payoff. By setting $k = 500, M = 29, m = 17.2$ *and* $a = 40$, this two person tournament has a unique Nash equilibrium at $n = 37$.

In this paper, as in Merlo and Schotter (1999), we turn this two person tournament into an individual decision making game by replacing one of the players with a computerized automaton programmed to always choose $n = 37$. Subjects are informed that the computerized opponent will choose 37 in every round.

The inter-generational framework is implemented in the following way. A group of subjects are recruited into the lab and play the stage game for 10 rounds (we will explain the exact experimental design and parameters shortly). After his participation is over, each player is replaced by another player, his laboratory descendant, who then plays the game for another 10 rounds as a member of a fresh group of subjects. The generations are non-overlapping. A particular player in generation $t+1$ can (1) see the **history** of choices by his generation t predecessor and (2) receive **advice** from that predecessor via free-form messages that generation t players leave for their generation $t+1$ successors. While the generations are non-overlapping still there is an inter-generational flavor in that payoffs span generations. Payoff to a generation t player is equal to what he has earned during his lifetime plus 100% of what his laboratory descendant earns. This provides an incentive to the subjects to pass on meaningful advice to their successors.

We can draw a parallel between our approach and the observational learning studied in Merlo and Schotter (2003) where players are paired with one member of the pair being the "doer" and the other being the "observer". Initially the "doer" participates in a constituent game (very similar to ours) for a number of rounds with the "observer" watching passively and without any communication between the two. They find that the "observers", when called upon to perform the same task, perform significantly better than the original "doers'. In our study each player in generation *t+1* not only received free-form advice from the generation *t* player but can also see the decisions made by the latter. Thus the set-up is similar to Merlo and Schotter's observational learning except in our set-up the observation is "indirect" in that the generation *t+1* player is not watching "over the shoulders" of the generation *t* player but rather receives a piece of paper which shows (1) all the decisions made by the generation *t* player and the payoff outcomes of those decisions and (2) the free-form advice left by the same player.

We find that the presence of advice makes a difference. The experimental groups who get advice perform better – their decisions are closer to the Nash equilibrium of 37 – compared to a control group of subjects that plays the game with no recourse to such advice. This is particularly true of those subjects who receive reasonably accurate advice about the Nash equilibrium choice. We also find that the deviations from the Nash equilibrium decline over generations. We proceed as follows. In the next section we describe the experimental design. We present our results in Section 3. In Section 4 we provide some concluding remarks.

2. Experimental Design

The experiments are computerized. A session starts with each subject seated in front of a computer. Subjects are seated far apart from one another ensuring that no one can look at any one else's screen. They are handed the instructions (see the Appendix) which are printed on paper. Subjects are given time to read over the instructions on their own. Then the experimenter reads the instructions out loud. Questions are answered prior to the beginning of the game. Each subject then begins round 1 of the game by picking a number and entering this number on the user interface using his keyboard. The number "*n*" that the subject picks ranges from 0 to 100. There is a quadratic cost function associated with the number picked which takes the form $n^2/500$. Given that *n* varies from 0 to 100, the associated cost ranges from 0 to 20. See the instructions included in the Appendix for the decision cost table. The cost information is part of the instructions given to the subjects and the computer shows the particular cost associated with the number chosen once the subject has entered his chosen number. The computer then generates a random number for the subject and one for itself. The subject is told what his random number is and what the subject's total number (decision number plus the random number) is. Then the subject learns whether he received the large fixed payment of 29 or the smaller one of 17.2. The subject's earning for the round is the prize (29 or 17.2) minus the cost associated with the number chosen. For instance suppose in a particular round a subject picks 65. The cost associated with this effort level is 8.45. Suppose the random number picked by the computer for this subject is 20 while the random number picked by the computer for itself is 25. Then the subject's total number is 85 (65+20) while the computer's total is 62 (37+25) since the computer always chooses 37. In this case the subject has a larger total number and wins the larger prize of 29. Given the cost of 8.45 associated with picking 65, his actual earning then would be 29 – 8.45 = 20.55. On the other hand if the subject's total number was less than the computer's total then the subject would get 17.2 minus the cost. The round then comes to an end and the next one begins. Each subsequent round proceeds in the same manner.

We institute two different treatments. In our experimental treatment with advice and history subjects play the stage game for 10 rounds. Subjects in generation *t+1*, after they have been given the instructions and before they start round 1, receive a piece of paper which has the free-form advice left by the generation *t* predecessor and also shows the history of the 10 decisions made by the predecessor and the payoff outcomes of those decisions. See Table 1 for a representative example of the history that each player receives. Thus each player gets to see the choices made (and the results of those choices) by the previous player and then makes 10 choices of his own. This is true for players in all generations except the first who do not get to see any history of choices. Each player receives two payoffs except it works differently for players in the last generation and everybody else. Players in all but the last generation, get a first payment equal to what they earn during their lifetime of 10 rounds plus a second payment which is equal to the earnings of their generation *t+1* successors over their 10 rounds. This second payment is given to them at a later date. The very last generation is paid differently. After the players in the last generation completed their stipulated 10 rounds of the game, they are asked to participate in a "surprise quiz". They are told that they will play the game one more time. Except the payoff from this additional decision will be 10 times the per-round payoff of the previous 10 rounds and they will be paid the total of their earnings during the regular 10 rounds plus a second payment which is equal to the earnings from the surprise quiz.

Behavior of this experimental group which gets to see the decisions made by the immediate predecessor and also receives advice from the same is compared to the behavior of a control group. Each subject in the control group plays the game for 20 rounds without any advice or history. At the end of those 20 rounds each subject in the control group takes part in a "surprise quiz" where they make their decision about the number one more time except the payoff associated with this decision is 10 times the payoff associated with each decision during the 20 regular rounds.

The experiment was carried out at Rutgers University.[1] There are 60 subjects in the experiment. There are 10 players in the control group who play the game for 20 rounds without any advice or history and then take part in the surprise quiz. This gives us 200 observations for the 20 rounds plus the 10 surprise quiz choices. Among the inter-generational groups that play with history and advice, we have 10 families – Family 1 through Family 10. Each family consists of five generations with players in generation *t+1* receiving advice and history from generation *t* before participating in 10 rounds of play. So we have 50 subjects in the advice treatment who make 10 decisions each for a total of 500 observations. Moreover the player in the last generation of each of the 10 families takes part in the surprise quiz giving us 10 surprise quiz decisions for these advice groups. Thus we have 720 observations in all. Table 2 shows the exact experimental design.

One point needs to be clarified here. Within a particular family we have five generations with each player in generation *t+1* getting to see advice and history from generation *t*, except generation 1. Thus within a family the observations for the five players are not independent. However each family does constitute an independent observation. Thus the 10 families of subjects in the advice treatment all constitute independent observations.[2]

The payoffs are denoted in francs and experimental earnings are converted into US dollars at the rate of 0.0375 dollars per franc. A typical ten round session lasts about 45 minutes while the 20 round sessions for the control groups lasted about an hour and a quarter. There is quite a bit of variation in earnings but on average subjects earned US $15.00.

[1] All the sessions were run within 5 days with multiple sessions on most days. We believe that this short time-span, coupled with the fact that subjects were recruited from a large pool, ensured that there was no communication between earlier and later generations.

[2] In running these experiments we face a trade-off between running more generations and running more families. In order for reliable patterns of behaviour to emerge it is important to run a substantial number of generations because it is unlikely that any interesting patterns would emerge within two or three generations. But the problem with running more generations is that these observations within a family are not independent. Thus one needs to run multiple families in order to generate independent observations. We ended up with 10 families, each of which represents an independent observation.

3. Results

3.1 The Nature of the Advice Left

We are interested in understanding if the presence of advice helps subjects make better decisions. In order to understand the impact of advice on subsequent decisions or lack thereof we need to analyze the nature of the advice left and quantify that advice.

As mentioned above subjects are allowed to leave free-form advice to their successors. A number of subjects made explicit recommendations to their successors regarding what number they should choose. For instance a subject may write "*You have NO CONTROL over the random number. So, make sure you win the decision round. To do that you need to select a number greater than 37. If you select too high a number, you incur high costs (costs seem to increase exponentially). I would recommend 38.*" In such cases the advice coding is easy and it is coded simply using the number that is being recommended. However sometimes the advice specified a range and in such cases we simply calculate the average of that range. A number of subjects obviously realized that there was a clue in the fact that the computerized opponent always chose 37 and left advice to choose close to 37 but "slightly above it" or variants of that such as *"You should play a number higher than 37. But it should not be too much higher than 37 as your decision cost will also be higher."* In such cases where a subject did not specify a particular number but stressed that it should be "*close to 37 but slightly above it*", we assign the number "39" to this advice. A number of subjects advised their successors to choose "small numbers" such as *"I chose to be cost-averse. I started the session choosing low numbers: 10, 30, 60, 90 . . Each time, I found that I was (sic) my total number was below that of the computer's. However, the increasing cost of the larger numbers subtracted a significant portion of my 17.2 fixed payment. In order to maintain the largest profit, I minimized my losses by returning to small numbers with low associated costs."* All such advice suggesting "small numbers" was coded as "25". Finally, possibly owing to the rather complex decision problem involved, there were a lot of subjects who left advice that had no substantive content at all such as "*the more random and spread out you stay with picking your decision numbers the better it is.*" or *"i found that the combination of odd and even number helped me increase my score. but towards the end my total earnings went down. i think that the computer understands the pattern we put our numbers in. if according to our pattern our numbers exceed 100, we will earn less total earnings.*" These subjects do not make any specific recommendations to their successors at all and all such advice was coded as "101". Figure 1 presents the distribution of advice left by the subjects. 17 out of 50 (34%) subjects left advice that did not make a specific recommendation regarding the choice of a number. But 21 subjects (42%) left advice which asked their successors to choose a number between 35 and 41 and 17 subjects (34%) advised choosing a number between 37 and 40.

3.2 Groups with Advice Make Decisions Closer to the Nash Equilibrium

In what follows we focus on the **absolute deviations from the Nash equilibrium prediction of 37** rather than the actual numbers chosen by the subjects in the two groups. Figure 2 shows the absolute deviations from 37 for the control and the advice groups. As is clear from Figure 2 the deviations from 37 for the groups with advice are generally smaller than the deviations for the control subjects. However it is also clear that while the groups with advice are closer to the Nash equilibrium there is no clear convergence in either group over time.

The average absolute deviation for the control subjects is 19.01, while that of the advice subjects is 16.47. The difference is marginally significant using a t-test (t = 1.79, p = 0.07) but not significant using the non-parametric Wilcoxon ranksum test (z = 1.27, p = 0.20). However a more useful comparison is between the control treatment and those who receive a specific recommendation regarding the number to choose. We have 27 subjects (i.e. 270 observations) who receive such advice. There are 200 observations for the control subjects. The mean absolute deviation for the control groups is 19.01 while that of the subjects with meaningful advice is 14.41 a difference that is significant both on a t-test (t = 2.91, p = 0.00) as well as a Wilcoxon ranksum test (z = 2.63, p = 0.00). Figure 3 shows the distribution of absolute deviations for the two groups. It is clear that the choices of the subjects who receive a specific recommendation (right panel) are closer to the Nash prediction of 37 (i.e. their absolute deviations are closer to zero) compared to those of the control subjects (left panel).

If we look at the choices made by those who did not receive any specific advice versus those who did - here of course we are focusing only on the subjects who do receive advice – then we find that the average absolute deviation from 37 for those who receive meaningless advice (i.e. advice that made no specific recommendation) is 19.82 while for the subjects who received a meaningful recommendation the average absolute deviation is 14.41. The difference is highly significant using either a t-test (t = 3.30, p = 0.00) or a Wilcoxon ranksum test (z = 3.87, p = 0.00). This is borne out by the histograms showing the distribution of absolute deviations by quality of advice received in Figure 4. In Figure 4 the left panel shows the decisions of those subjects who receive meaningless advice (advice quality = 0) while the right panel shows the decisions of those subjects who get meaningful advice (advice quality = 1). Looking at the right panel, it is clear that the mass point of absolute deviations moves to the left for those who receive meaningful advice. Thus it is obvious that the presence of advice that makes a specific recommendation has a strong impact in driving decisions towards the Nash equilibrium. Figure 5 highlights this point by showing the absolute deviations for **round 1 only** for those who received meaningless advice in the left panel and those who received meaningful advice on the right panel. It is clear that there is a shift of the mass point of deviations towards zero in the right panel showing that in the first round of any generation, subjects who received meaningful advice made choices which are equal to or close to the Nash equilibrium of 37. We find that for subjects who received meaningful advice the correlation coefficient between the number chosen and the advice received is 0.31 (p = 0.00, n = 270). If we focus on the correlation between number chosen in the first period (i.e. immediately after receiving the advice) and the advice received then this correlation increases to 0.50 (p = 0.00, n = 27).

One question that we can ask at this juncture is the following: how much would history help if advice was not available? If a subject did not receive any specific recommendation then it would seem to make sense that this subject should look at the choices made by the predecessor. In order to look at the correlation between a subject's choice and the numbers chosen by his predecessor, we compute the average of the 10 choices made by the predecessor and look at the correlation between a subject's choice and the average choice made by the predecessor. This of course is relevant for the subjects in generation 2 onwards in the advice treatment. We find that there is strong correlation between the predecessor's average choice and a subject's own choices – a correlation coefficient of 0.11 (p = 0.02, n = 400).

However if we look separately at those who received meaningful advice versus those who received meaningless advice then we find the following. The correlation coefficient between a subject's own choices and the average choice of the predecessor is -0.12 (p = 0.16, n = 130) for those subjects who receive meaningless advice. While for those who receive meaningful advice, the correlation coefficient between a subject's own choices and the average choice of the predecessor is 0.20 (p = 0.00, n = 270).

If we look at the correlation between the actual number chosen by an individual subject and the advice left by that subject then we find that they are correlated. (Correlation coefficient = 0.16, p = 0.00, n = 500). This suggests the following: those who make better decisions (i.e. choose numbers close to 37) leave better advice, which in turn leads to better decisions by their successors. However if a player made "good" decisions but left "bad" advice then it is highly likely that his successor will make "bad" decisions since in choosing a number it is the advice, rather than the predecessor's actual choices, that plays the crucial role.

3.3 Deviations from Nash Equilibrium Decrease over Generations

We find that among the subjects in the advice treatment the absolute deviation from 37 decreases over generations. See Table 3. Here we are taking the average of the absolute deviation from 37 over all 10 rounds for all players in a particular generation. The average absolute deviation over all 10 rounds for the 10 players in the first generation is 17.68. This generation does not receive advice but leaves advice. The average absolute deviation goes down to 15.11 in generation 2 (the first generation to receive advice) but increases in generation 3 to 19.42 before decreasing to 15.67 and 14.48 in generations 4 and 5 respectively.

Figure 6 shows the distribution of the absolute deviations for the five generations in the advice treatment. For purposes of comparison we have provided the histogram for the control group (generation = 0) as well. As is clear from Figure 6 the mass point of the distribution moves towards the left (i.e. towards smaller deviations from the Nash equilibrium of 37) over generations.

In Table 4 we present some additional parametric results showing that the absolute deviation from the Nash equilibrium decreases over generations. Since we have data for a cross-section of subjects who make ten decisions each, the data has a panel structure. Moreover absolute deviations are bound by zero from below (when the actual decision is equal to 37) and by 63 from above (for decisions = 100). Thus the data is doubly censored and we model this using a double censored random effects Tobit regression. Furthermore, for this part of the analysis we will restrict our attention to only generations 2 through 5 since it is only these four generations which receive advice from their predecessors. The dependent variables include (1) "Round" which ranges from 0 to 10; (2) "Advice Received", and (3) "Generation" which takes the values of 2, 3, 4 or 5. As one can see from Table 4 the coefficient for generation is negative and significant showing that the deviations from the Nash equilibrium do decrease over the four generations which receive advice. The coefficient for "round" is positive and significant which attests to the fact that there is little learning going on for individual subjects over the 10 rounds that they play.

3.5 Distribution of Surprise Quiz Choices

Finally we look at the surprise quiz choices made by the two groups. We have 10 players in the control groups who each made a surprise quiz choice. We also have 10 families in the advice treatment with the player in the last generation of each family making a surprise quiz choice. The mean and median surprise quiz choices for the control groups are 44.1 and 44 respectively while for the advice treatment the mean and the median are 55.8 and 45.5 respectively. However if we carry out a t-test or a Wilcoxon ranksum test to test for the equality of the two distributions of absolute deviations then we do not find a significant difference. Thus we cannot reject the null hypothesis that these observations are drawn from the same population.

4. Concluding Remarks

In this paper we have adduced evidence that in a reasonably complex decision making situation involving a tournament set-up, word-of-mouth advice left by players who have some experience in the game has an impact and helps facilitate convergence to the Nash equilibrium. Our results then provide corroboration to the results presented in Merlo and Schotter's (2003) findings about observational learning.

However having made that claim we hasten to add the following two caveats. First, while our inter-generational set-up does bear testimony in favor of advice facilitating convergence to the Nash equilibrium, such convergence is rather slow. In that sense it would be instructive to have more than five generations in each family. More generations might help establish a clearer pattern. Of course one must bear in mind the trade-off between more generations and more families as pointed out in footnote 2 above.

Second, the nature of the advice left is crucial. As we have shown above, succeeding generations make decisions that are strongly correlated with the advice they get. However this has both a positive and a negative impact. When the advice received is "good", the subsequent decisions are close to the Nash equilibrium. But if the advice is "bad" then that moves successors away from the equilibrium even if the predecessor made decisions close to the equilibrium. Thus the knowledge transmission over generations relies heavily on (1) an individual player's ability to figure out the equilibrium best-response, (2) his ability to communicate that to the successor via written advice and (3) the successor's willingness to abide by that advice.

References:

Bull, C., A. Schotter and K. Weigelt (1987), "Tournaments and Piece Rates: An Experimental Study", *Journal of Political Economy*, 95 (1), p. 1 – 33.

Camerer, C. (2003), Behavioral Game Theory, Princeton University Press, Princeton, NJ.

Merlo. A. and A. Schotter (2003), "Learning by Not Doing: An Experimental Investigation of Observational Learning", *Games and Economics Behavior*, 42, p.116-136

Merlo. A. and A. Schotter (1999), "A Surprise Quiz View of Learning in Economic Experiments", *Games and Economic Behavior*, 28, p. 25-54.

Schotter, A. (2003), "Decision Making with Naïve Advice", *American Economic Review*, 93 (2), p. 196-201.

Table 1: Representative History of Past Plays Shown to Players in Generation 2 onwards in the Advice Treatment

Period	Decision Number	Random Number	Cost	Total Number	Fixed Payment	Earnings
1	50	14	5	64	29	24
2	50	-11	5	39	29	24
3	100	-28	20	72	17.2	-2.8
4	50	23	5	73	29	24
5	100	10	20	110	29	9
6	50	26	5	76	29	24
7	75	-15	11.25	60	29	17.75
8	50	32	5	82	29	24
9	25	32	1.25	57	17.2	15.95
10	50	7	5	57	29	24

Table 2: Design of the Experiment

<table>
<tr><td colspan="10" align="center"><u>Control Group</u></td></tr>
<tr><td colspan="10">Players: #1 #2 #3 #4 #5 #6 #7 #8 #9 #10
Each player plays for 20 rounds and then participates in “Surprise Quiz” at the end</td></tr>
<tr><td colspan="10" align="center"><u>Advice Groups</u></td></tr>
<tr><td>Family 1</td><td>Family 2</td><td>Family 3</td><td>Family 4</td><td>Family 5</td><td>Family 6</td><td>Family 7</td><td>Family 8</td><td>Family 9</td><td>Family 10</td></tr>
<tr><td>Players
#11
#12
#13
#14
#15</td><td>Players
#21
#22
#23
#24
#25</td><td>Players
#31
#32
#33
#34
#35</td><td>Players
#41
#42
#43
#44
#45</td><td>Players
#51
#52
#53
#54
#55</td><td>Players
#61
#62
#63
#64
#65</td><td>Players
#71
#72
#73
#74
#75</td><td>Players
#81
#82
#83
#84
#85</td><td>Players
#91
#92
#93
#94
#95</td><td>Players
#101
#102
#103
#104
#105</td></tr>
<tr><td colspan="10">Last player of each family such as players #15, #25, #35 etc. take part in surprise quiz</td></tr>
</table>

Table 3: Absolute Deviations over Generations

Group	Absolute Deviation from Nash Prediction of 37
Control Group	19.01
Advice Group Generation = 1	17.68
Advice Group Generation = 2	15.11
Advice Group Generation = 3	19.42
Advice Group Generation = 4	15.67
Advice Group Generation = 5	14.48

Table 4: Random Effects Tobit Regression for Absolute Deviations from the Nash Equilibrium of 37

Variable	Coefficient	Standard Error	Z-statistic	p-value
Round	1.018	0.215	4.74	0.000
Advice Received	0.210	0.054	3.92	0.000
Generation	-2.963	1.323	2.24	0.025
Constant	8.049	5.627	1.43	0.153
Sigma_u	10.453			
Sigma_e	12.146			
Number of observations = 400		Number of groups = 80		
Wald χ^2 = 39.61		370 uncensored observations		
Prob > χ^2 = 0.000		27 left-censored observations		
Log likelihood = -1514.361		3 right-censored observations		

Figure 1: Histogram for Distribution of Advice

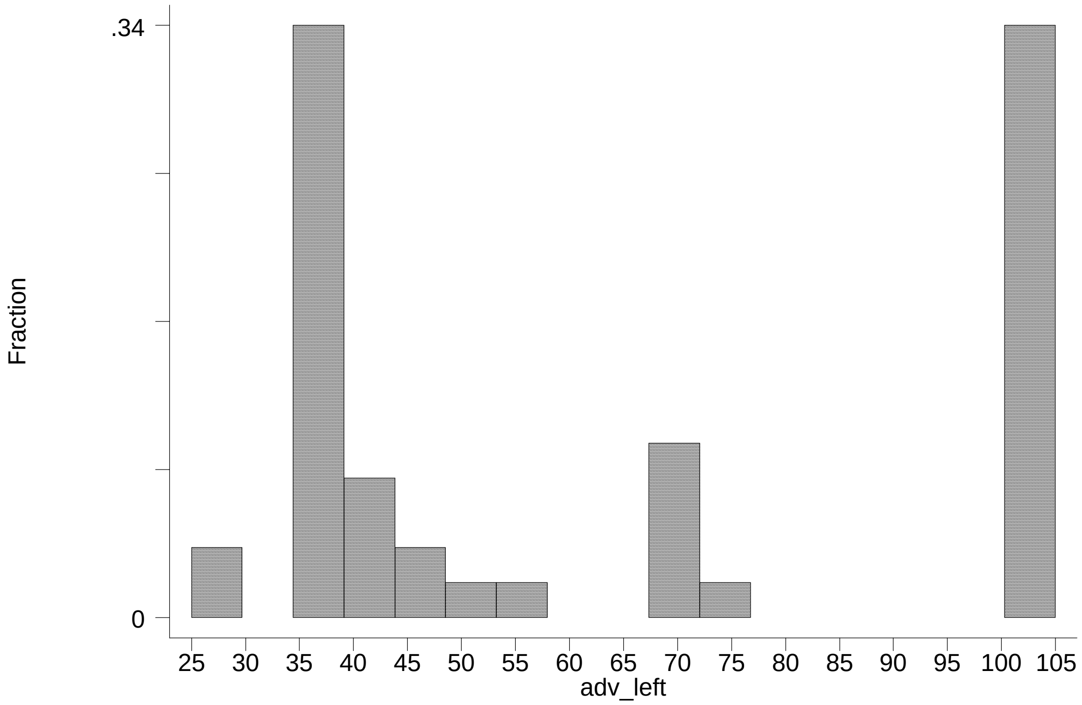

Figure 2: Deviation from the Nash Equilibrium Value of 37 for the Two Groups across Rounds

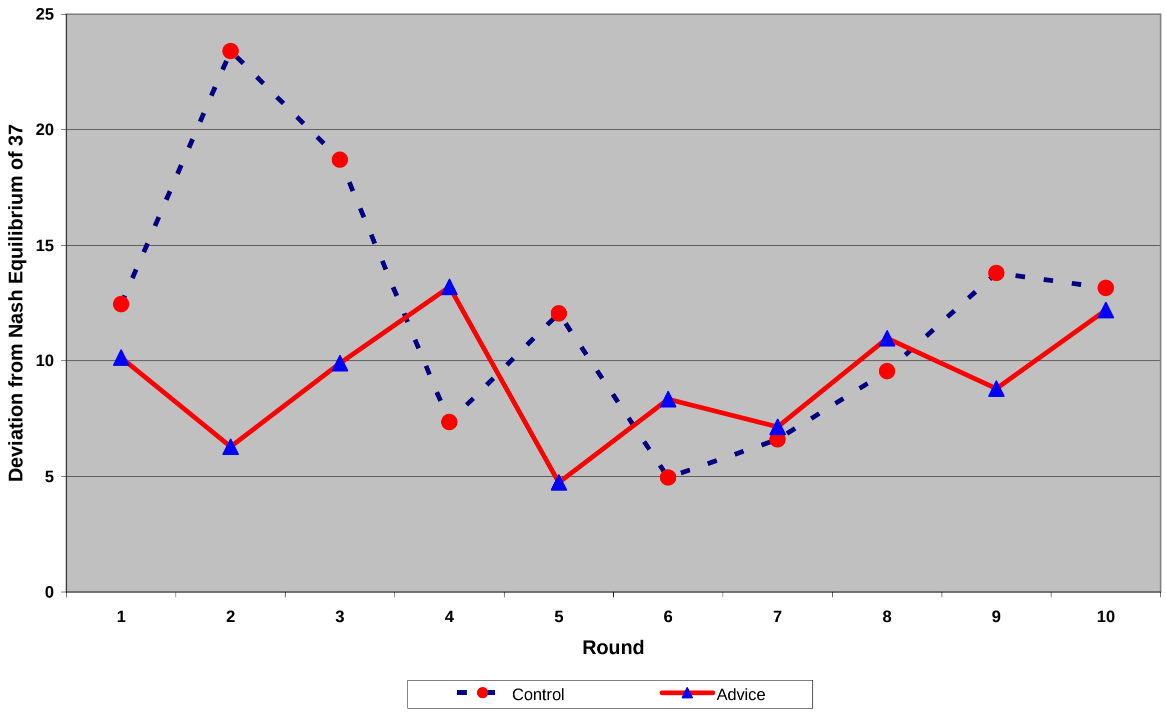

Figure 3: Histogram of Absolute Deviations from the Nash Prediction of 37 for the Control and Advice treatment subjects

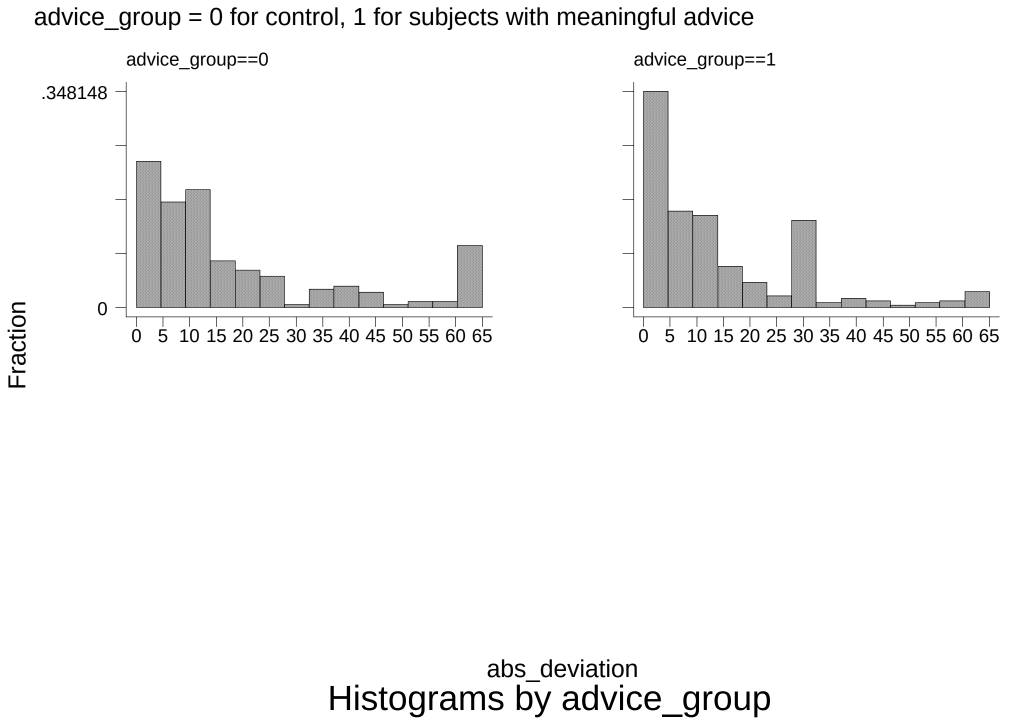

Figure 4: Histogram of Absolute Deviations by Advice Quality

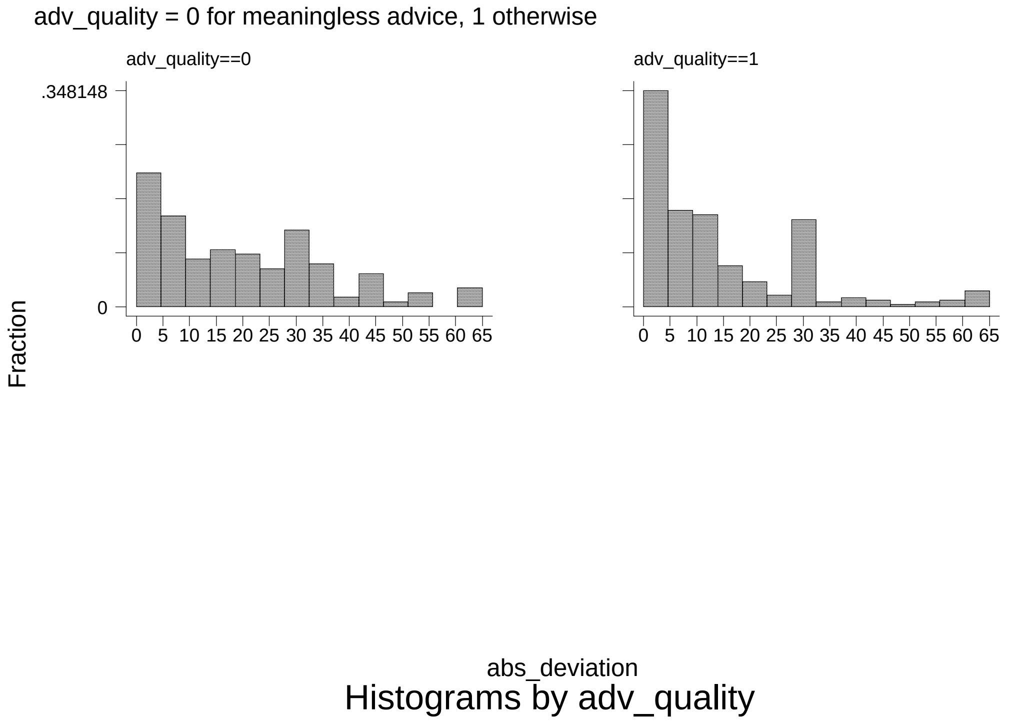

Figure 5: Histogram of Absolute Deviations by Advice Quality for Round 1 only

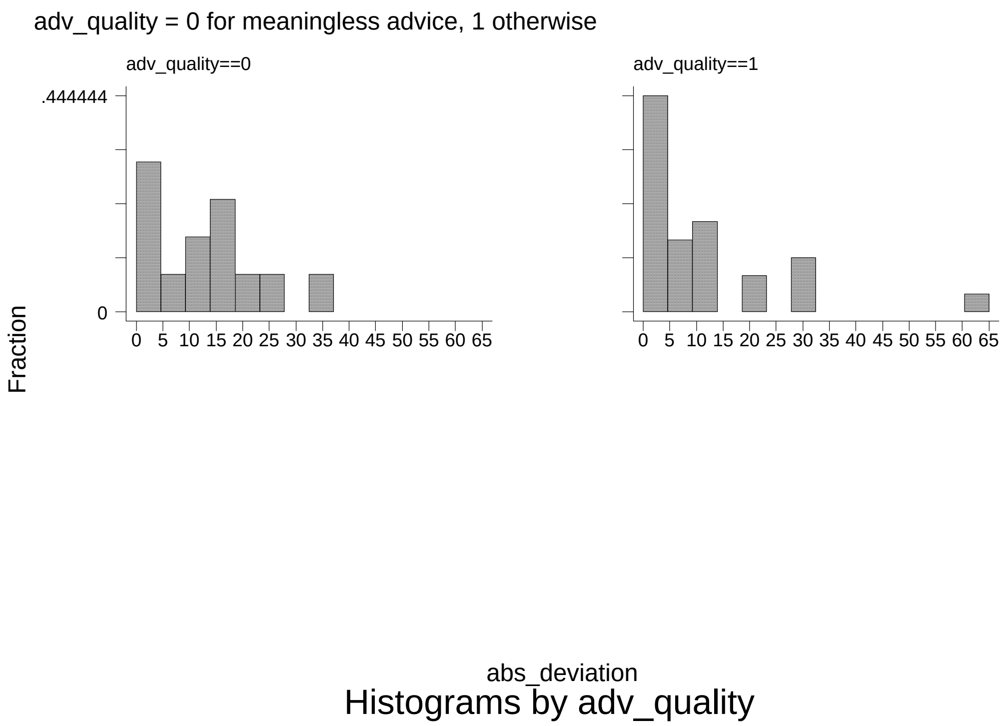

Figure 6: Histogram of Absolute Deviations from the Nash Prediction of 37 over Generations

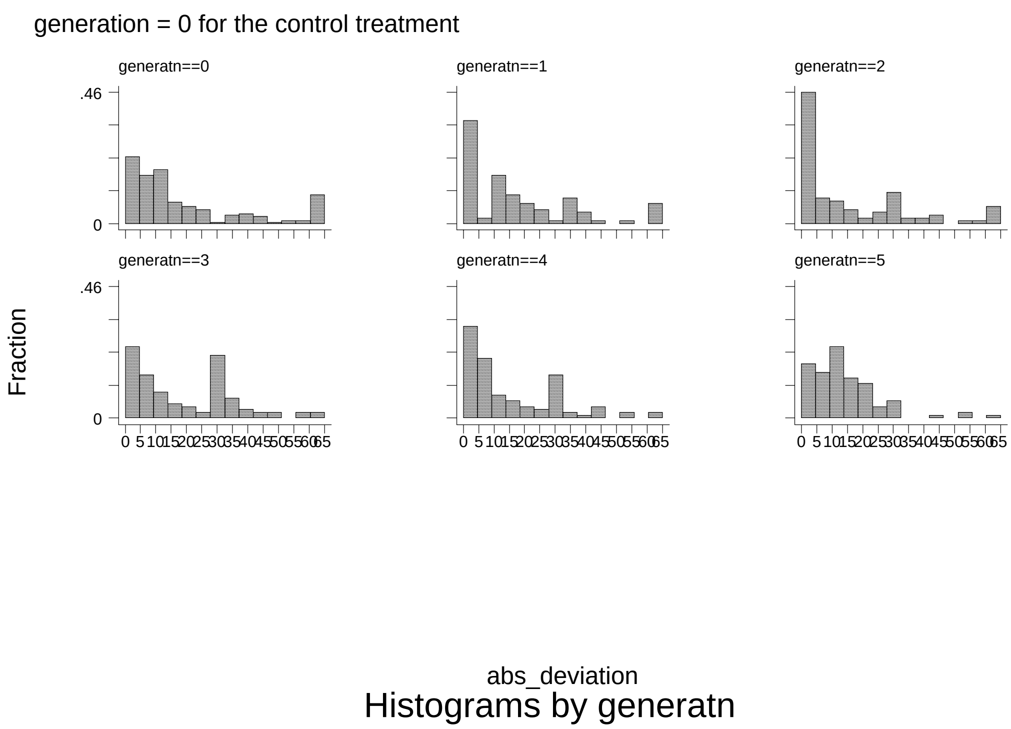

Appendix

Player ID # ____________

Instructions

Introduction

This is an experiment about decision making. The instructions are simple, and if you follow them carefully and make good decisions, you could earn a considerable amount of money, which will be paid to you in cash.

All monetary amounts in the experiment are denominated in an experimental currency called francs. At the end of the experiment your earnings in francs will be converted into dollars at the rate of .0375 dollars per franc. You will also be paid an additional $5.00 simply for showing up and completing the experiment.

Specific Instructions

As you read these instructions you will be in a room with a number of other subjects. Each subject has been randomly assigned an ID number and a computer terminal. The experiment consists of 10 decision rounds. In each decision round you will be paired with a computerized subject which has been programmed to make the same decision in every round. The computerized subject randomly matched with you will be called your *computerized pair member*. Your computerized pair member will remain the same throughout the entire experiment.

Unless you are in the first group to participate in this experiment, when you start the experiment you will receive advice on how to make your decisions from a subject who participated in the experiment immediately prior to you. This subject will earn an additional payment equal to your earnings from the 10 decision rounds that you complete in the experiment. Besides this advice, you will also be able to see the history of the previous 10 decision rounds of the subject who is giving you advice. After your 10 decision rounds you will leave advice to a new subject on how to make decisions, and you will receive an additional payment equal to the earnings of the subject you give advice to in his or her 10 decision rounds. Details on how you leave advice will be explained below.

Experimental Procedures

In the experiment you will perform a simple task. Attached to these instructions is a sheet called the *Decision Cost Table*. This sheet shows 101 numbers from 0 to 100 in column 1. These are your decision numbers. Associated with each decision number is a decision cost, which is listed in column 2. Note that the higher the decision number chosen, the greater is the associated decision cost. Your computer screen should look as follows as you enter the lab:

Enter a number between 0 and 100: ________

Period	Decision Number	Cost	Random Number	Total Number	Fixed Payment	Earnings	Total Earnings

In each decision round the computer will ask you to choose a decision number. Your computerized pair member will also choose a decision number. Remember that it will always choose the same decision number, which will be 37 in every decision round. You, of course, are free to choose any number you wish among those listed in column 1 of your Decision Cost Table. Therefore, in each round of the experiment, you and your computerized pair member will each select a decision number **separately** (and you know that it will always choose 37). Using the number keys, you will enter your selected number, which will appear in the space following "Enter a number between 0 and 100," and then hit the Enter key. At this point the computer will display the round number, the decision number you selected for the round, and the decision cost associated with your decision number. The computer will then prompt you to press the Enter key again, at which time the computer will generate your Random Number. Your random number will be one of the 81 whole numbers that lie between -40 and +40 (including zero), and each of these numbers is equally likely to be selected in each round. Your computerized pair member's random number is also equally likely to be any one of the 81 whole numbers lying between -40 and +40. The process that generates your random number is independent of the process that generates your computerized pair member's random number. That is, there is no relationship between the random number generated for you and the random number generated for your computerized pair member. Your Random Number and your Total Number (Decision Number + Random Number) will both be displayed. At this time the computer will have also generated your computerized pair member's Random Number and will compute It's Total Number as well.

Computation of Payoffs

Your payment in each decision round will be computed as follows. After your random number has been selected and your computerized pair member's random number has been selected, the computer will then prompt you to press the Enter key again, at which time the computer will compare your Total Number to the Total Number of your computerized pair member. Your fixed payment will be 29 if your Total Number is bigger than or equal to your computerized pair member's Total Number. Your fixed payment will be 17.2 if your computerized pair member's Total Number is bigger than your Total Number. After determining which fixed payment you receive, the computer will subtract your associated decision cost from your fixed payment to determine your Earnings for the round. The computer will also calculate your Total Earnings over the course of the experiment, ie, the earnings you have accumulated in each round. Your fixed payment, earnings, and total earnings will then be displayed on the screen.

Continuing Rounds

After round 1 is over, you will perform the same procedure for round 2, and so on for 10 rounds. In each round you will choose a decision number and generate a random number by pressing the Enter key at the computer's prompt. Your Total Number will be compared to the Total Number of your computerized pair member, and the computer will calculate your fixed payment and earnings for the round. When round 10 is completed, the computer will convert your total earnings in francs into dollars and display this amount. We will pay you this amount plus the $5.00 show-up fee.

In rounds after round 1, the results from earlier rounds will be displayed on the screen, and you may refer to these results at any time. The results displayed are: the period number, the decision number in that period, the associated decision cost, your random number and total number, your earnings, and your total earnings (both in francs). You will also receive the same information from the 10 decision rounds completed by the subject who left you advice. You may refer to this information at any time as well.

At the end of your 10 decision rounds, you will be able to leave advice to a new subject in the experiment. You will type this advice using a word processor on your computer. This advice will be printed and given to a new subject, along with a printout of the history of your decisions. You will receive a second payment, equal to the earnings of the subject who receives your advice in his or her 10 decision rounds. You will be notified by E-mail or telephone when your second payment is ready.

If you have any questions please raise your hand and the monitor will come and answer it.

Let us begin.

Decision Cost Table

NUMBER	COST	NUMBER	COST	NUMBER	COST	NUMBER	COST
0	0.000						
1	0.002	26	1.352	51	5.202	76	11.552
2	0.008	27	1.458	52	5.408	77	11.858
3	0.018	28	1.568	53	5.618	78	12.168
4	0.032	29	1.682	54	5.832	79	12.482
5	0.05	30	1.8	55	6.05	80	12.8
6	0.072	31	1.922	56	6.272	81	13.122
7	0.098	32	2.048	57	6.498	82	13.448
8	0.128	33	2.178	58	6.728	83	13.778
9	0.162	34	2.312	59	6.962	84	14.112
10	0.2	35	2.45	60	7.2	85	14.45
11	0.242	36	2.592	61	7.442	86	14.792
12	0.288	37	2.738	62	7.688	87	15.138
13	0.338	38	2.888	63	7.938	88	15.488
14	0.392	39	3.042	64	8.192	89	15.842
15	0.45	40	3.2	65	8.45	90	16.2
16	0.512	41	3.362	66	8.712	91	16.562
17	0.578	42	3.528	67	8.978	92	16.928
18	0.648	43	3.698	68	9.248	93	17.298
19	0.722	44	3.872	69	9.522	94	17.672
20	0.8	45	4.05	70	9.8	95	18.05
21	0.882	46	4.232	71	10.082	96	18.432
22	0.968	47	4.418	72	10.368	97	18.818
23	1.058	48	4.608	73	10.658	98	19.208
24	1.152	49	4.802	74	10.952	99	19.602
25	1.25	50	5	75	11.25	100	20