

Do We Detect and Exploit Mixed Strategy Play by Opponents?

Jason Shachat
IBM Watson Research Center
jshachat@us.ibm.com

J. Todd Swarthout
University of Arizona
swarthout@econlab.arizona.edu

January 15, 2003

Abstract: We conducted an experiment in which each subject repeatedly played a game with a unique Nash equilibrium in mixed strategies against some computer-implemented mixed strategy. The results indicate subjects are successful at detecting and exploiting deviations from Nash equilibrium. However, there is heterogeneity in subject behavior and performance. We present a one variable model of dynamic random belief formation which rationalizes observed heterogeneity and other features of the data.

I. Introduction

A Nash equilibrium of a normal form game can be identified as the fixed point of the players' best response correspondences. The notion that each player anticipates his opponents' actions and best responds to this belief has been an effective approach in the analysis of strategic decision making. However, when opponents adopt mixed strategies, we should carefully consider the assumption that a player forms accurate beliefs and best responds. This is especially so in constant-sum games which do not have a pure strategy Nash equilibrium.¹ In such games, when opponents don't play their Nash equilibrium strategies, a player typically has a unique pure strategy best response that gives an expected payoff greater than the Nash equilibrium level. Whether a person detects such payoff increasing opportunities, however, is an open question. In this study we report an experiment in which subjects repeatedly play a zero-sum game against different constant mixed strategies, and examine to what extent subjects detect and exploit such opportunities.

In our experiment, a subject was assigned one of the two possible roles in an asymmetric matching pennies game. The subject then played 200 repetitions against a constant mixed strategy. The subject was told that he was playing against another decision maker, but was not informed that this decision maker was computerized nor the nature of the decision maker's strategy. We varied the mixed strategy faced by different subjects to cover a broad spectrum of possible mixed strategies. This enables us to ascertain the extent to which best response correspondences describe subject behavior.

There are two main results. First, subjects come surprisingly close to best responding to an unknown mixed strategy, even if the mixed strategy is no more than fifteen percent above or below the Nash equilibrium strategy. Second, subjects are quite successful at exploiting mixed strategies that deviate from Nash equilibrium and, as a consequence, increase their payoffs above Nash equilibrium levels. However, both results only hold on average as there is heterogeneity across subject behavior and earnings. In an attempt to characterize this heterogeneity, we introduce a single variable adaptive belief model that

¹ The minimax and Nash equilibrium solutions coincide in this setting, and we could proceed only referring to the minimax solution and strategies. However, we proceed using the Nash equilibrium framework because we wish to focus on the concept of best response.

we call the “Random Hierarchical Belief” model. The model is based on the notion that before each stage game a subject randomly draws a belief from a Beta distribution and then best responds to this belief. The two shape parameters of the Beta distribution are discounted counts of the opponent’s previous action choices. The discount rate of these counts is a subject specific variable that generates differential behavior. We estimate this discount variable for each of our subjects and then present simulations to demonstrate that the model generates the qualitative dynamics and heterogeneity reflected in the data.

Our study extends and clarifies the results of previous studies of play in 2×2 zero-sum games with a unique equilibrium in mixed strategies. Lieberman [1961] and Fox [1972] both studied subject play against a mixed strategy that deviated from Nash equilibrium. They discovered that subjects significantly adjust their play and increase their earnings. However, in these studies only a single non-Nash strategy was evaluated, and this strategy differed from the Nash strategy by a probability greater than twenty-five percent. In our study, we systematically vary the mixed strategies to obtain a more complete characterization of how humans play in these situations and we take advantage of more sophisticated software to provide each subject a complete history past play.

Next we describe the experimental design and procedures. In the third section we address how subjects’ choice frequencies adjust from the first half to the second half of the experiment, examine whether subjects increase their payoffs above Nash equilibrium levels, and present the Random Hierarchical Belief model to explain the dynamics and heterogeneity found in the data. In the final section, we offer some concluding remarks.

II. Experimental Design and Protocols

We employ a zero-sum asymmetric matching pennies game (introduced by Rosenthal, Shachat, and Walker [2002]). In the game each player can move either Left or Right. The normal form representation of the game is given in the table below. The game has a unique Nash equilibrium in which each player chooses Left with probability two-thirds. When Column doesn’t adopt the equilibrium strategy, Row’s best response is to play Left if Column chooses Left with a probability greater than two-thirds, and to play Right otherwise. Likewise when Row doesn’t adopt the equilibrium strategy, Column’s best

response is to play Right if Row plays Left with a probability greater than two-thirds, and to play Left otherwise. In equilibrium, Row's expected payoff is $2/3$ and Column's expected payoff is $-2/3$.

		Column Player	
		Left	Right
Row Player	Left	1,-1	0,0
	Right	0,0	2,-2

We conducted all experimental sessions in the Economic Science Laboratory at the University of Arizona during the Fall of 2002. We report results from seven sessions, using a total of 102 undergraduate students. Each session contained between 8 and 22 subjects. Half of the subjects were assigned as Row players, and the other half were assigned as Column players.

Each subject was seated at a computer workstation such that no subject could observe another subject's screen. Subjects first read computerized instructions that detailed both how to enter decisions and how earnings were determined. Then, 200 repetitions of the game were played. Column subjects were initially endowed with a balance of 250 tokens, while Row players began with no tokens: each token was valued at 10¢. Each subject's total earnings consisted of a \$5 show-up payment plus his token balance after the 200th repetition. No Column subjects went bankrupt.²

At the beginning of each repetition, a subject saw a graphical representation of the game on the screen. Each Column subject's game displays transformed so that he appeared to be a Row player. Thus, each subject selected an action by clicking on a row,

² Notice that the possible Row subjects' stage game payoffs were all non-negative (gains) and the possible Column subjects' payoffs were all non-positive (losses.) In studies of individual decision making under uncertainty, for example Kahneman and Tversky (1979), it is often concluded that individuals are risk averse when all possible outcomes are gains and risk loving when all possible outcomes are losses. In our experiment such preferences shift the opponent's mixed strategy that renders a Row or Column subject indifferent between Left and Right to something less than two-thirds. As we will see shortly, there isn't strong evidence of such an effect in the data.

and then confirmed his selection. After the repetition was complete, each subject saw the outcome highlighted on the game display, as well as a text message stating both players' actions and his own earnings for that repetition. Finally, at all times a subject's current token balance and a history of past play were displayed. The history consisted of an ordered list with each row displaying the repetition number, the actions selected by both players, and the subject's payoffs from the specific repetition.

Each subject played against a fixed computerized mixed strategy. The various mixed strategies adopted and the number of subjects who played against them are presented below. Each subject was informed that he was going to play against the same decision maker for all repetitions: he was not informed that the decision maker was a computer or the nature of the decision maker's strategy. Although human subjects never played against each other, each Row subject was matched with a Column subject: this was done to reduce the chance a subject believed he was playing against a computer. Specifically, while the computer generated instantaneous action choices, the software did not reveal the computer's action until both paired human subjects had made action selections. This process allowed the pair to progress at a more natural rate determined by the response speed of the two subjects.

Percentage Left by Mixed Strategy	Number of Subject Pairs
19%	4
27%	4
35%	4
43%	4
51%	7
59%	4
67%	3
75%	7
83%	8
91%	6

III. Data Analysis

We start the data analysis by addressing to what extent subjects best respond to different mixed strategies. We find that a subject's play is likely to move substantially

towards his best response when his opponent's choice frequencies are more than fifteen percent above or below the Nash equilibrium frequencies. Correspondingly, we find that subjects achieve a statistically significant increase in payoffs above Nash equilibrium levels when facing mixed strategies that deviate from the Nash equilibrium by more than fifteen percent. However, there is heterogeneity across subjects to the degree they best respond and maximize potential payoffs. We present a single parameter random belief adjustment model that rationalizes this heterogeneity.

III.1 Best Response and Payoff Gains

A natural starting point is to inspect how often each subject best responds when his opponent's choice frequencies deviate from the Nash prediction. We present this view of the data for the Row subjects in Figure 1 and for the Column subjects in Figure 2. In each of these figures, the solid line represents the subjects' best response correspondence. Also, each arrow is a summary of play for a single human/computer pair. The origin of the arrow is located at the joint frequency of Left play in the first 100 stage games, and the tip of the arrowhead is located at the joint frequency of Left play in the second 100 stage games. These arrows show the adjustments subjects make from the first-half to the second-half experiment regarding how often subjects best respond.

We can make several observations from these figures. First, the further his opponent deviates from Nash equilibrium frequencies of Left play the more likely a subject is to best respond. However, this statement needs two qualifications. First, the opponents' deviations must be sufficiently far from equilibrium to see all subjects' move close to the best response. Also, it is clear the subjects' frequencies of Left play differ in the magnitudes of adjustment from the first half to second half of the experiment. However, we can't identify if these movements towards best response result from learning about the opponent or learning about the functioning of the experiment. Finally, when his opponent's play is near the Nash equilibrium strategy, the human's proportions are biased towards levels below the Nash equilibrium proportion.

We provide a statistical evaluation of whether a subject's play is significantly towards his best response. First, we establish a baseline for when play is decidedly in the

direction of best responding. When the subject's best response is to play Left, we say that his play is "better responding" if his probability of Left play exceeds two-thirds. Similarly, when the subject's best response is to play Right, we say he is better responding if his probability of playing Left is less than one-half.³ Utilizing these baselines, we construct two hypothesis tests for play in the last one hundred stage games. The first is a binomial test for which the null hypothesis is that the subject's probability of Left play equals two-thirds and the alternative hypothesis is that this probability exceeds two thirds. At the five percent level of significance, we reject the null in favor of the alternative whenever a subject plays Left more than seventy-five times. We depict the critical region of this test on Figures 1 and 2 with a dashed line at the subject proportion of .75 within the area for which Left is a best response. The second hypothesis test is another binomial test with the null hypothesis that the subject's probability of Left play is fifty percent and the alternative is that the probability is less than fifty percent. At the five percent level of significance, we reject the null hypothesis whenever a subject plays Left fewer than forty-one times. We depict the critical region of this test on Figures 1 and 2 with a dashed line at the subject proportion of .41 within the area for which Right is a best response.

We note that frequencies of Left play fall out of the two critical regions of better responding for only 16 of 51 Row subjects and 17 of 51 Column subjects. For Row subjects, Left frequencies are all within the critical region for better responding towards Left when the computer's frequency exceeds 80 percent and also within the critical region for better responding towards Right when the computer's frequency of left is less than 50 percent. Likewise, all Column subjects are in the critical region for better responding towards Right when the computer's Left frequency exceeds 80 percent. However, the uniform movement towards Column's critical region for Left doesn't occur until the opponent's frequency of Left falls below 35%. Figure 2 demonstrates marked heterogeneity in the Column subjects' tendencies to move towards the best response when the opponent's frequency of Left play is below the Nash equilibrium levels. We will see that this results in differential earnings for the Column subjects.

³ We are choosing the benchmark of fifty percent because we have already noted that the proportion of Left human play is biased below two-thirds when the facing Nash equilibrium proportion. We feel that in this instance setting the Null at two-thirds would bias our conclusions towards subjects better responding.

The next metric we consider is the subjects' average stage game earnings. We ascertain whether subjects successfully exploit non-Nash equilibrium mixed strategies and how close they come to maximizing potential payoffs. In Figures 3 and 4, for the last 100 stage games, we plot each subject's average stage game payoff versus his opponent's frequency of Left. An open circle indicates a subject's earnings that we can't reject are the same as the Nash equilibrium payoffs, and the solid triangle indicates a subject's earnings that we conclude exceed the Nash equilibrium level. These conclusions are reached via a hypothesis test performed at a five percent level of significance. The solid lines found on Figures 3 and 4 represent the expected payoff from playing the pure strategy best response. As is commonly known, in these games a player's payoff function is relatively flat around his opponent's Nash equilibrium strategy. This is evident as we see mostly open circles in the frequency range of fifty to eighty percent. However, when a computer decision maker deviates from the Nash proportion by more than fifteen percent the subjects successfully increase their payoffs. This is not true in the case where Column subjects face mixed strategies less than two-thirds. Here we observe that some subjects fail to exploit mixed strategies as low as thirty percent while other subjects' earnings are close to the maximum expected payoff.

While subjects are surprisingly adept, on average, at detecting and exploiting non optimal mixed strategies, the heterogeneity of these abilities across subject can't be ignored. We provide a possible explanation of this heterogeneity by formulating a simple adaptive model of play.

III.2 Random Hierarchal Beliefs Model of Heterogeneity and Adjustment

In this subsection we present a simple one-variable model of random belief formation. We then estimate the single variable for each subject. Lastly we present a simulation, which demonstrates the ability of the model to rationalize the heterogeneity observed across subjects.

Recall the two players are Row and Column, which we will denote by r and c . Stage games are indexed by n . Each player i 's set of actions is $A_i = \{L, R\}$ and the action player i selects in stage game n is a_{in} . Player i 's set of mixed strategies is $\Sigma_i = [0,1]$. A mixed

strategy, $\sigma_{in} \in \Sigma_i$, is the probability that player i selects L in stage game n . Finally, let b_{in} be player i 's belief of what player j 's mixed strategy will be in stage game n .

We now propose a one-variable model to describe the dynamics of how subjects played the game. We assume that before each stage game a player's belief is determined by a draw from a distribution over his set of possible beliefs, and that a player selects the best response to this belief. We call this model the Random Hierarchical Belief (RHB) model because subjects' beliefs are determined by a hierarchical probability structure.

The belief b_{in} is a random variable which has a Beta distribution function, $\beta(\cdot)$, with the player specific parameters s_{iLn} and s_{iRn} . The support of a Beta distribution is the unit interval with a mean of $s_{iLn}/(s_{iLn} + s_{iRn})$ and a mode, if both s_{iLn} and s_{iRn} are greater than one, equal to $(s_{iLn} - 1)/(s_{iLn} + s_{iRn} - 2)$. When both parameters are one, the Beta distribution is simply the Uniform distribution. The parameters of the Beta distribution have an important interpretation in Bayesian statistics. If one is modeling a binomial process and starts with a Beta prior density, the posterior density is also Beta for which the first parameter s_{iLn} is incremented by the number of successes and the second parameter, s_{iRn} , is increment by the number of failures. These parameters are often called the prior sample sizes.

Our model follows the spirit of this Bayesian interpretation; the parameters s_{iLn} and s_{iRn} are determined by the observed history of play according to the following rules:

$$s_{iLn} = \delta * s_{iLn-1} + I\{a_{jn-1}=L\} \quad \text{and} \quad s_{iRn} = \delta * s_{iRn-1} + I\{a_{jn-1}=R\} \quad \text{for } n > 1, \text{ and} \\ s_{iLn} = s_{iRn} = 1 \quad \text{for } n = 1.$$

The unobservable variable δ is a discount rate for the two parameters and $I\{a_{jn-1}=a_j\}$ is an indicator function for the event that player j chose action a_j in the stage game $n-1$.

Furthermore, by setting the initial values $s_{iL1} = s_{iR1} = 1$, the player's belief in the first stage game is drawn from a uniform distribution.

Consider an example. Suppose Row has a δ of one-half. In the first period his belief about Column's mixed strategy is drawn from the uniform distribution on the unit interval. The probability he draws a belief for which L is a best response is $1/3$, i.e. $1 - \beta(2/3, 1, 1) = 1/3$. Therefore the model predicts $\Pr(a_{r1}=L) = 1/3$. Now suppose that his opponent chooses R in the first period. With this outcome $s_{rL2} = .5$ and $s_{rR2} = 1.5$, and in

the second stage game the probability that Row draws a belief for which L is the best response, is $1 - \beta(2/3, .5, 1.5) = .09$. If Column chose R in the second stage game then $s_{rL3} = .25$ and $s_{rR3} = 1.75$, and the probability Row chooses L in stage game three is $1 - \beta(2/3, .25, 1.75) = .03$. Consider a last iteration in which the Column player selects L in stage game three. In this case $s_{rL4} = 1.125$ and $s_{rR4} = .875$, and the Row player selects L in the fourth stage game with the probability $1 - \beta(2/3, 1.125, 1.75) = .41$.

To better appreciate the flexibility of the RHB model, consider two special cases. As the discount rate δ approaches zero, behavior approaches a simple best response dynamic. Also, when the discount rate is one, the mode of the belief distribution follows a fictitious play process and the belief is drawn from a Bayesian posterior distribution on the opponent's mixed strategy.

Again each of our subjects played against some computer implemented fixed mixed strategy. For each of our subjects, we estimate δ by a maximum likelihood procedure. For a Row subject the probability he chose the action L in stage game n is

$$\Pr(a_{rn} = L) = 1 - \beta(2/3, s_{rL1}(\delta_r), s_{rR1}(\delta_r)).$$

The resulting log-likelihood function for each of our Row subjects is

$$\ln L = \sum_{n=1}^{200} \ln \left(I_{\{a_{rn}=L\}} \left(1 - \beta\left(\frac{2}{3}, s_{rLn}(\delta_r), s_{rRn}(\delta_r)\right) \right) + I_{\{a_{rn}=R\}} \beta\left(\frac{2}{3}, s_{rLn}(\delta_r), s_{rRn}(\delta_r)\right) \right).$$

Similarly the log-likelihood function for each of our Column subjects is

$$\ln L = \sum_{n=1}^{200} \ln \left(I_{\{a_{cn}=L\}} \beta\left(\frac{2}{3}, s_{cLn}(\delta_c), s_{cRn}(\delta_c)\right) + I_{\{a_{cn}=R\}} \left(1 - \beta\left(\frac{2}{3}, s_{cLn}(\delta_c), s_{cRn}(\delta_c)\right) \right) \right).$$

In Tables 1 and 2 we report for each subject the maximum likelihood estimate of δ and the result of a forecasting exercise. For each subject role estimates are listed in increasing order. The lowest estimate of the variable is .457 and the highest is 1.01.⁴ The third column reports the percentage of Left play by the opponent for all 200 stage games. The fourth through eighth columns report the data and results for a within sample forecasting exercise.

We asked how well the RHB model predicts play in the last 100 stage games. For each subject, we generate a sequence of choice probabilities of Left using his actual

⁴ $\Omega\epsilon$ *τρυνχάτεδ* *τηε* *εστιματέδ* *παλβε* *οφ* δ at 1.01, as the behavior of the likelihood function quickly deteriorates as δ exceeds one and estimates are difficult to obtain.

opponent's choices and his estimated value of δ . We report the average of this sequence in the fourth column. In addition, we calculate the square of the difference between the subjects' predicted choice probability and his actual choice – where a choice of Left is set to one and a choice of Right is set to zero – in each of the last 100 stages games and we sum these squared differences. In column six, we report the total sum of squared errors of the subject's estimated choice probabilities. We also report the subject's total proportion of left play for all 200 stage games in column five, and the total sum of squared errors of this proportion for the last 100 stage games in column seven. We report the difference of the two of sum of squared errors statistics in column eight. The RHB generally has a higher sum squared error in two circumstances; either a subject's frequency Left play is on the opposite side of fifty percent from the best response or a subject's frequency of best response is nearly one.

Finally we provide a simulation to show how the RHB characterizes diverse subject behavior. For each player type, we start by selecting the lowest and highest estimate obtained of δ . Then we simulate the two RHB models playing 200 stage games of the game against a mixed strategy. For both of the values of delta, we record the proportion of Left in the last 100 stage games by the RHB model. We do this exercise one hundred times for a particular mixed strategy. We then calculate the average proportions of Left play in the last 100 stage games for the two values of δ across the one hundred exercises. We report these averages as the RHB response to the mixed strategy. We do this simulation for the mixed strategies in the interval [.05, .99] using a step size of .01. The simulation generates a pair of response surfaces for the RHB model, one for the low estimate of δ and one for the high estimate of δ . These response surfaces are presented in Figures 5 and 6.

Figure 5 presents the response surfaces for the Row player and Figure 6 presents the response surfaces for the Column player. The low value of δ (.457 for the Row player and .619 for the Column player) produces a response surface that is almost a line segment that connects the two ends of the best response correspondence. On the other hand, the high value of δ (1.01 for both player types) produces a surface that indicates more frequent best responses. The two surfaces show how the RHB model characterizes the data. We display the scatter plot of human/computer joint left frequencies on the figure

and the scatter plot is quite similar to the response surfaces. For example, the subjects and the RHB model play Left substantially less often than the Nash equilibrium frequency of two-thirds when playing against mixed strategies that are close to the Nash equilibrium. We are encouraged by the ability of the RHB model to account for some of the heterogeneity exhibited by the subjects.

IV. Concluding Remarks

In this paper we test whether subjects can detect and exploit non-equilibrium play in a zero-sum game with a unique equilibrium in mixed strategies. In order to provide an informative test we conducted an experiment in which subjects repeatedly play against computer implemented mixed strategies. The mixed strategies were varied across subjects. We observe subjects, on average, doing remarkably well at adjusting their strategies towards a best response and achieving payoffs above their Nash equilibrium levels. However, there is substantial heterogeneity in subjects' behavior and performance. We formulated a single variable model of probabilistic belief formation that captures this heterogeneity and other features of the data.

There are several directions for further research. First, one can investigate whether subjects can detect and best respond when opponents deviate from Nash equilibrium strategies in other classes of games. Second, the adopted methodology of having humans play against preprogrammed strategies can be used to address other open questions in game theory such as how do alternative behavioral rules influence the convergence to equilibrium when there are multiple equilibria. Finally, the promising empirical performance of the simple Random Hierarchical Belief model should be tested on data sets from other game experiments and the properties of the dynamics of the model should be explored analytically.

Bibliography

Fox J (1972) The Learning of Strategies in a Simple, Two-Person Zero-Sum Game without Saddlepoint. *Behavioral Science* 17: 300-308

Kahneman, Daniel and Amos Tversky (1979) Prospect Theory: An Analysis of Decision Under Risk. *Econometrica* 47: 263-292

Lieberman B (1962) Experimental Studies of Conflict in Some Two-Person and Three-Person Games. In: Criswell JH, H Solomon, and P Suppes (eds.) *Mathematical Methods in Small Group Processes*. Stanford University Press, Stanford, pp. 203-220

Rosenthal RW, J Shachat, and M Walker (2001) Hide and Seek in Arizona. Technical Report, IBM TJ Watson Research Laboratory

Table 1: Maximum Likelihood Estimates of d for Row Subjects

Row Subject	MLE d	Col. Left Frequency	Avg. RHB Frequency Left	Proportion Left By Subject	Sum Squared Error RHB	Sum Squared Error	Difference
1	0.457	0.735	0.672	0.680	22.762	43.520	-20.758
2	0.459	0.195	0.132	0.205	23.604	32.595	-8.991
3	0.531	0.560	0.454	0.535	35.177	49.755	-14.578
4	0.538	0.275	0.178	0.250	30.403	37.500	-7.097
5	0.604	0.170	0.078	0.110	18.660	19.580	-0.920
6	0.649	0.230	0.102	0.210	36.422	33.180	3.242
7	0.657	0.350	0.210	0.255	34.927	37.995	-3.068
8	0.705	0.450	0.275	0.485	69.218	49.955	19.263
9	0.723	0.210	0.067	0.065	9.081	12.155	-3.074
10	0.753	0.935	0.935	0.890	17.852	19.580	-1.728
11	0.754	0.740	0.656	0.535	41.659	49.755	-8.096
12	0.767	0.480	0.285	0.255	28.456	37.995	-9.539
13	0.792	0.915	0.917	0.850	22.670	25.500	-2.830
14	0.794	0.560	0.388	0.380	40.830	47.120	-6.290
15	0.797	0.485	0.256	0.290	46.104	41.180	4.924
16	0.804	0.840	0.815	0.850	24.414	25.500	-1.086
17	0.805	0.245	0.054	0.070	12.987	13.020	-0.033
18	0.806	0.825	0.812	0.485	75.230	49.955	25.275
19	0.807	0.810	0.774	0.625	44.342	46.875	-2.533
20	0.816	0.745	0.669	0.570	51.106	49.020	2.086
21	0.820	0.595	0.432	0.490	53.230	49.980	3.250
22	0.824	0.605	0.445	0.515	68.138	49.955	18.183
23	0.826	0.320	0.089	0.085	12.431	15.555	-3.124
24	0.832	0.345	0.090	0.125	22.654	21.875	0.779
25	0.835	0.260	0.040	0.040	7.011	7.680	-0.669
26	0.845	0.740	0.660	0.395	57.490	47.795	9.695
27	0.846	0.475	0.240	0.160	29.176	26.880	2.296
28	0.854	0.885	0.901	0.835	27.447	27.555	-0.108
29	0.859	0.865	0.890	0.835	30.497	27.555	2.942
30	0.864	0.795	0.754	0.620	57.314	47.120	10.194
31	0.885	0.755	0.715	0.515	52.167	49.955	2.212
32	0.888	0.915	0.952	0.950	8.712	9.500	-0.788
33	0.889	0.900	0.918	0.945	12.343	10.395	1.948
34	0.891	0.405	0.111	0.090	11.534	16.380	-4.846
35	0.891	0.740	0.677	0.630	33.710	46.620	-12.910
36	0.893	0.925	0.958	0.905	16.885	17.195	-0.310
37	0.913	0.505	0.192	0.195	31.129	31.395	-0.266
38	0.914	0.870	0.920	0.885	18.104	20.355	-2.251
39	0.916	0.600	0.367	0.480	57.041	49.920	7.121
40	0.930	0.775	0.764	0.730	35.082	39.420	-4.338
41	0.931	0.405	0.058	0.070	10.077	13.020	-2.943
42	0.935	0.830	0.881	0.925	8.648	13.875	-5.227
43	0.940	0.380	0.044	0.085	11.480	15.555	-4.075
44	0.945	0.920	0.979	0.975	4.596	4.875	-0.279
45	0.963	0.505	0.124	0.110	19.636	19.580	0.056
46	0.970	0.540	0.101	0.135	21.436	23.355	-1.919
47	0.975	0.680	0.607	0.445	56.132	49.395	6.737
48	0.978	0.695	0.637	0.800	31.863	32.000	-0.137
49	0.982	0.355	0.020	0.040	3.110	7.680	-4.570
50	0.990	0.640	0.250	0.270	42.797	39.420	3.377
51	1.010	0.655	0.562	0.640	46.527	46.080	0.447

Table 2: Maximum Likelihood Estimates of d for Column Subjects

Column Subject	MLE d	Row Left Frequency	Avg. RHB Frequency Left	Proportion Left By Subject	Sum Squared Error RHB	Sum Squared Error	Difference
1	0.619	0.200	0.904	0.885	17.252	20.355	-3.103
2	0.629	0.670	0.423	0.545	35.712	49.595	-13.883
3	0.666	0.375	0.776	0.660	45.452	44.880	0.572
4	0.675	0.350	0.799	0.735	34.648	38.955	-4.307
5	0.689	0.515	0.633	0.650	44.354	45.500	-1.146
6	0.703	0.285	0.871	0.820	23.814	29.520	-5.706
7	0.716	0.790	0.271	0.195	24.344	31.395	-7.051
8	0.722	0.185	0.946	0.935	11.444	12.155	-0.711
9	0.741	0.490	0.688	0.605	57.472	47.795	9.677
10	0.748	0.175	0.959	0.920	13.946	14.720	-0.774
11	0.749	0.400	0.809	0.650	49.261	45.500	3.761
12	0.749	0.250	0.918	0.895	13.204	18.795	-5.591
13	0.761	0.725	0.368	0.445	32.713	49.395	-16.682
14	0.764	0.345	0.858	0.780	35.151	34.320	0.831
15	0.782	0.760	0.307	0.325	33.183	43.875	-10.692
16	0.794	0.580	0.599	0.580	50.298	48.720	1.578
17	0.797	0.295	0.917	0.895	17.512	18.795	-1.283
18	0.799	0.510	0.695	0.560	64.783	49.280	15.503
19	0.811	0.820	0.209	0.270	33.273	39.420	-6.147
20	0.817	0.885	0.126	0.275	38.506	39.875	-1.369
21	0.828	0.800	0.222	0.530	69.519	49.820	19.699
22	0.832	0.925	0.068	0.290	46.031	41.180	4.851
23	0.835	0.505	0.726	0.635	46.294	46.355	-0.061
24	0.838	0.760	0.300	0.555	56.473	49.395	7.078
25	0.847	0.510	0.729	0.680	42.803	43.520	-0.717
26	0.849	0.865	0.126	0.165	23.458	27.555	-4.097
27	0.863	0.830	0.177	0.340	45.894	44.880	1.014
28	0.865	0.720	0.352	0.575	59.895	48.875	11.020
29	0.883	0.890	0.078	0.120	19.166	21.120	-1.954
30	0.883	0.215	0.982	0.985	1.165	2.955	-1.790
31	0.891	0.840	0.140	0.150	23.869	25.500	-1.631
32	0.893	0.850	0.117	0.185	26.458	30.155	-3.697
33	0.893	0.310	0.953	0.960	6.736	7.680	-0.944
34	0.902	0.875	0.078	0.110	15.371	19.580	-4.209
35	0.917	0.710	0.363	0.590	54.410	48.380	6.030
36	0.920	0.435	0.920	0.905	16.699	17.195	-0.496
37	0.922	0.705	0.350	0.395	50.160	47.795	2.365
38	0.937	0.530	0.808	0.800	36.797	32.000	4.797
39	0.940	0.745	0.280	0.375	41.222	46.875	-5.653
40	0.942	0.930	0.026	0.030	2.665	5.820	-3.155
41	0.945	0.940	0.994	0.990	0.964	0.990	-0.026
42	0.966	0.830	0.095	0.065	10.020	12.155	-2.135
43	0.967	0.580	0.785	0.820	32.952	29.520	3.432
44	0.970	0.865	0.023	0.040	5.730	7.680	-1.950
45	0.978	0.470	0.962	0.975	5.115	4.875	0.240
46	1.010	0.335	0.997	0.990	1.012	1.980	-0.968
47	1.010	0.665	0.330	0.460	48.883	49.680	-0.797
48	1.010	0.430	0.971	0.990	3.766	1.980	1.786
49	1.010	0.525	0.990	0.995	1.295	0.995	0.300
50	1.010	0.615	0.765	0.980	19.375	3.920	15.455
51	1.010	0.540	0.972	1.000	1.307	0.000	1.307

Figure 1: Joint Left Frequencies of Human Row Players vs Fixed Mixtures
First and Second 100 Stage Games for each Pair

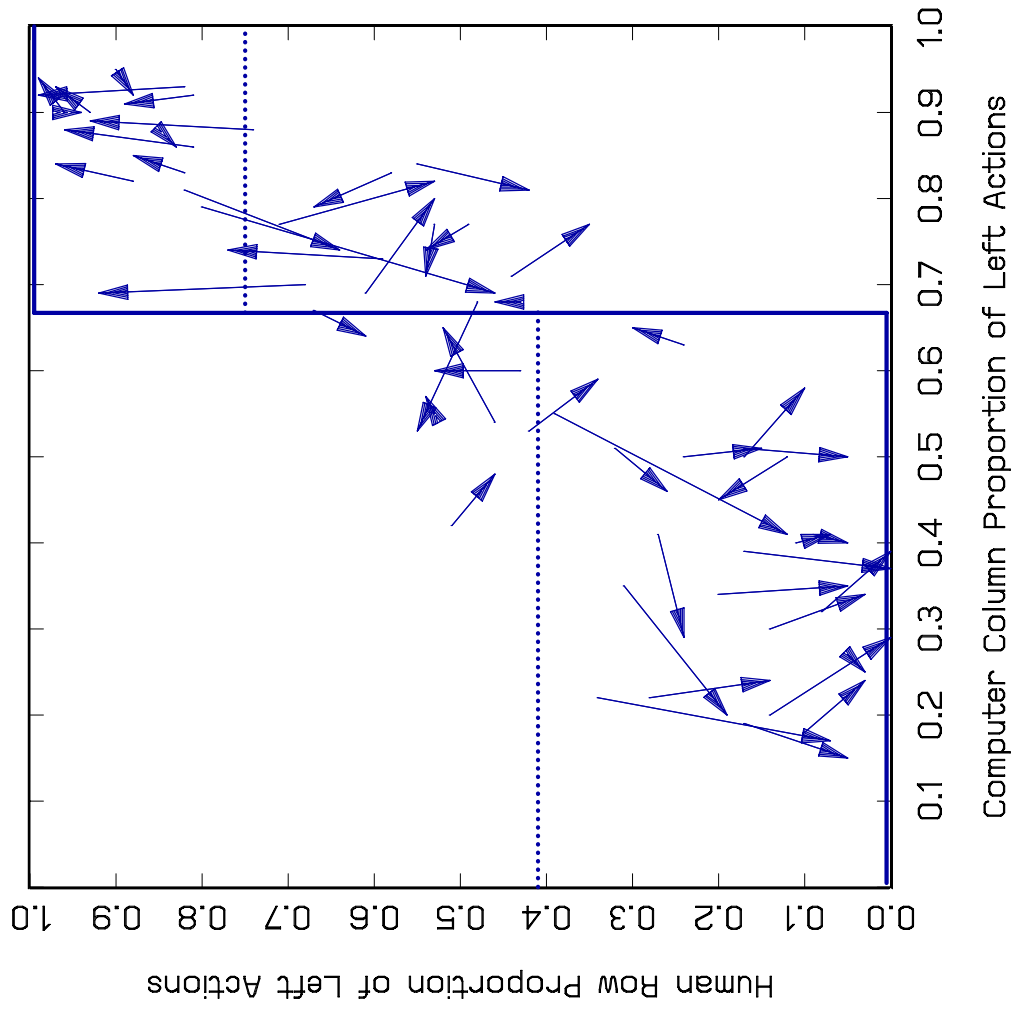


Figure 2: Joint Left Frequencies of Human Column Players vs Fixed Mixtures
First and Second 100 Stage Games for each Pair

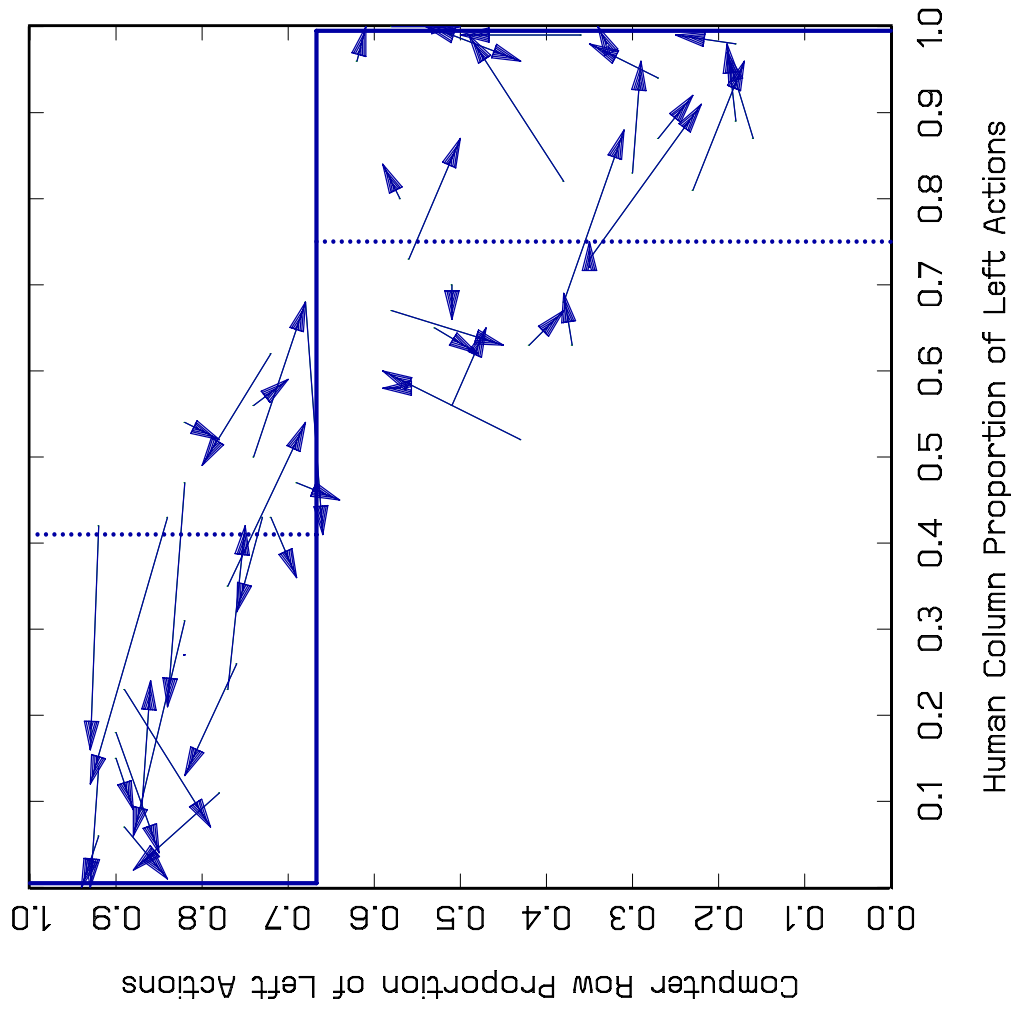


Figure 3: Average Stage Game Payoffs for Human Row Players
(Last 100 Stage Games)

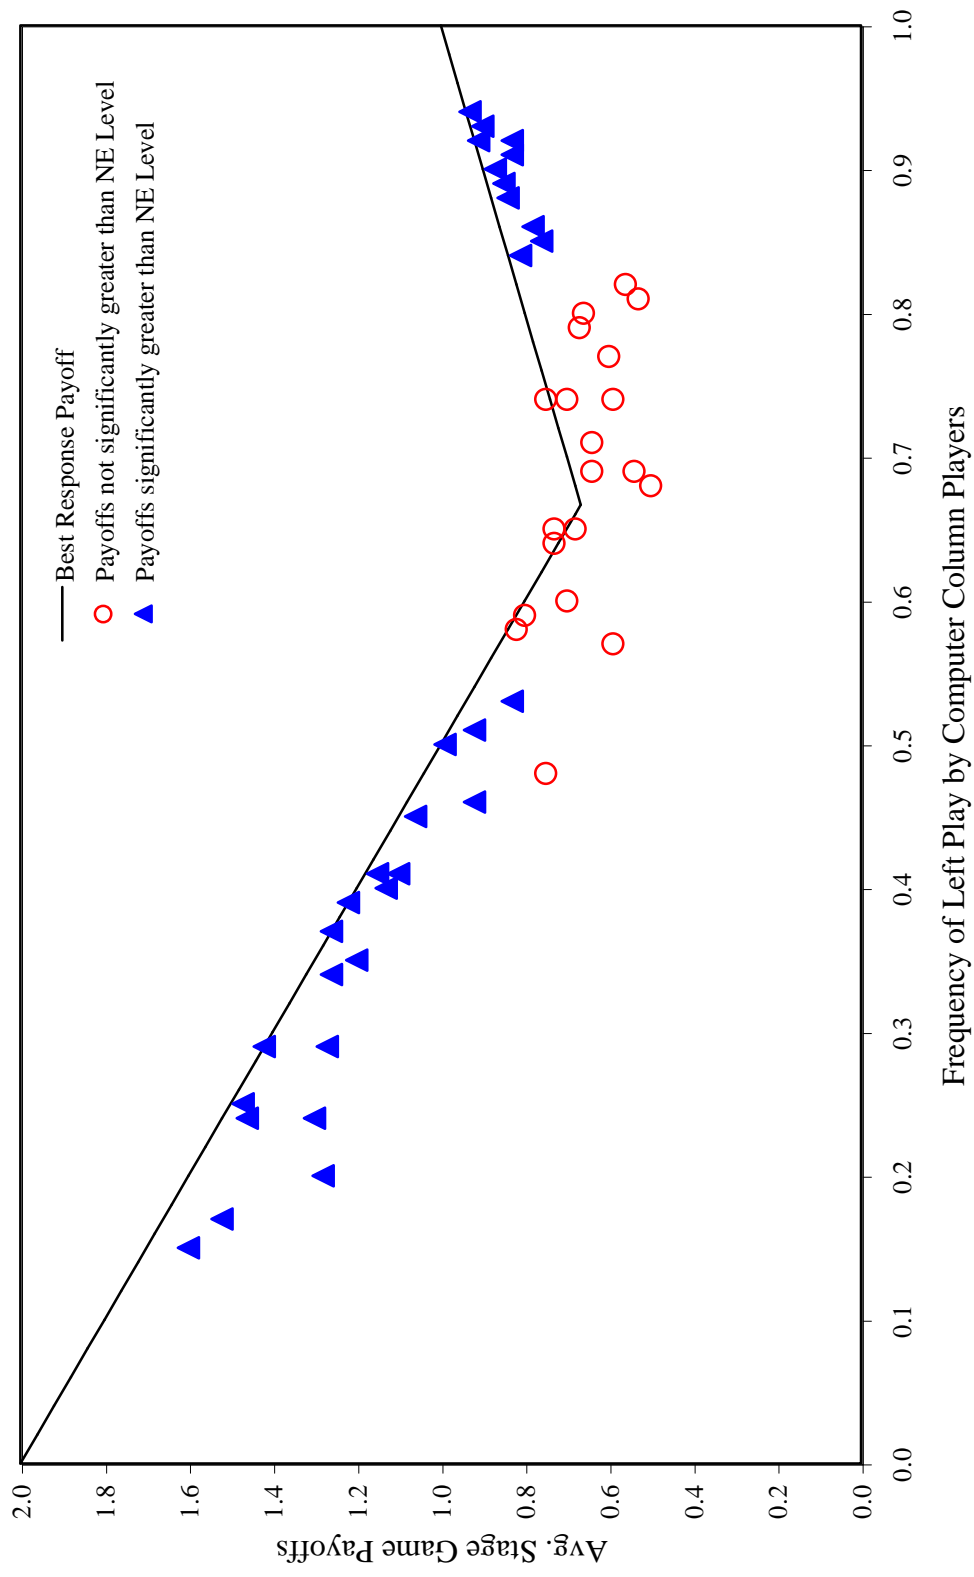


Figure 4: Average Stage Game Payoffs for Human Column
Players
(Last 100 Stage Games)

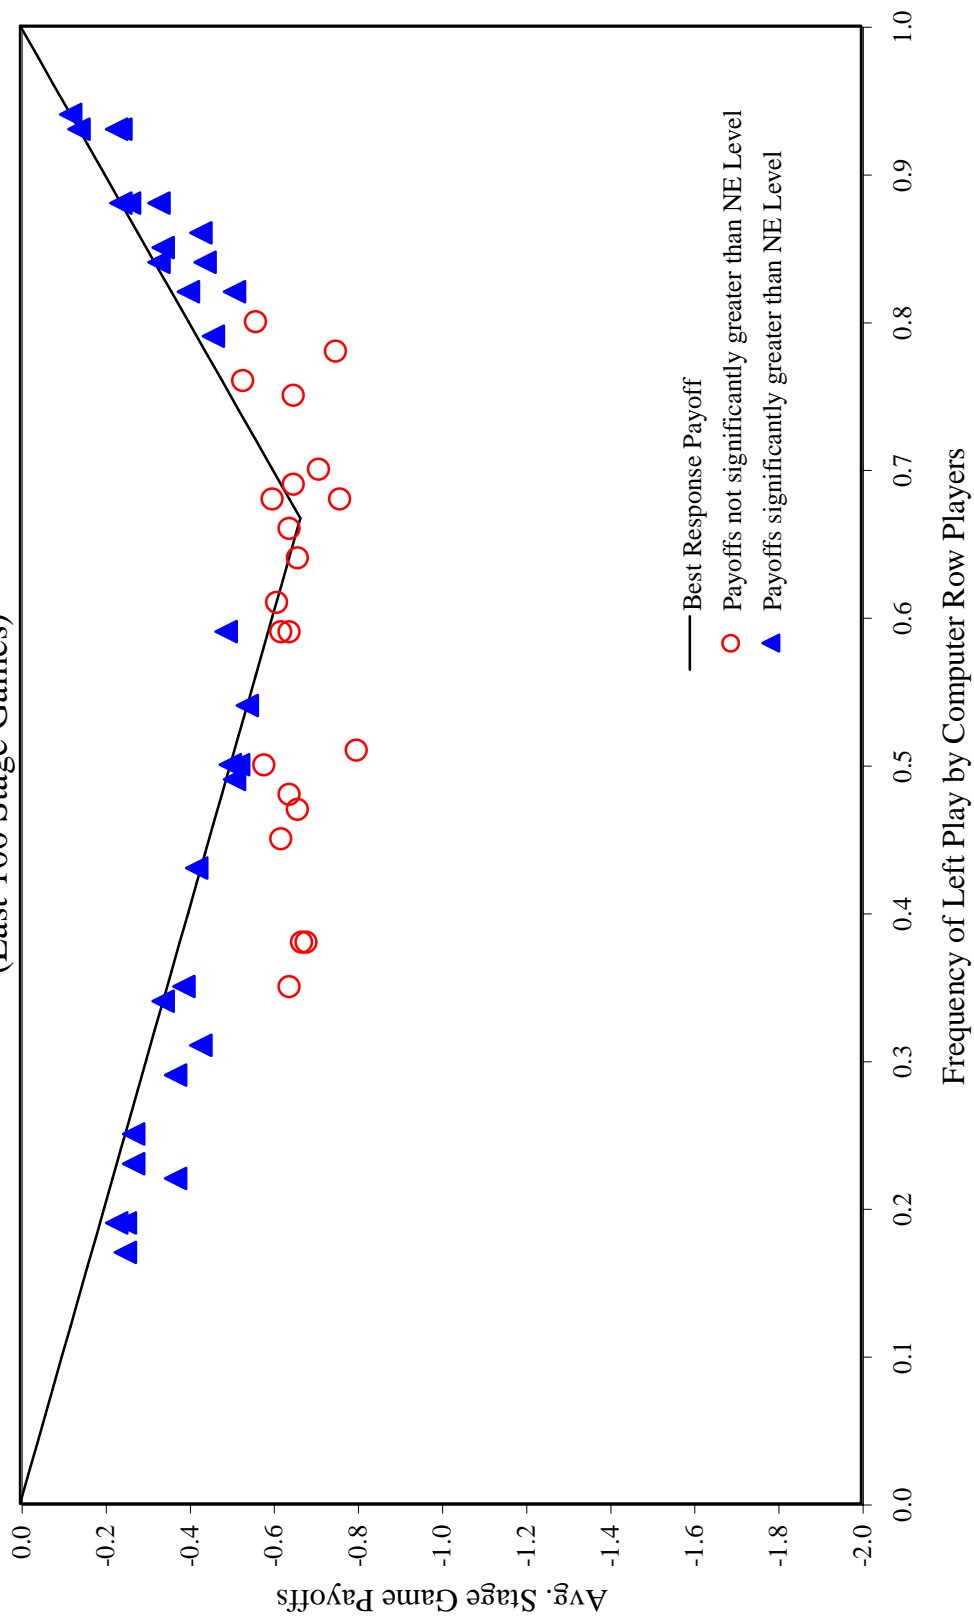


Figure 5 Row Player RHB Response Relationships - Last 100 Stage Games

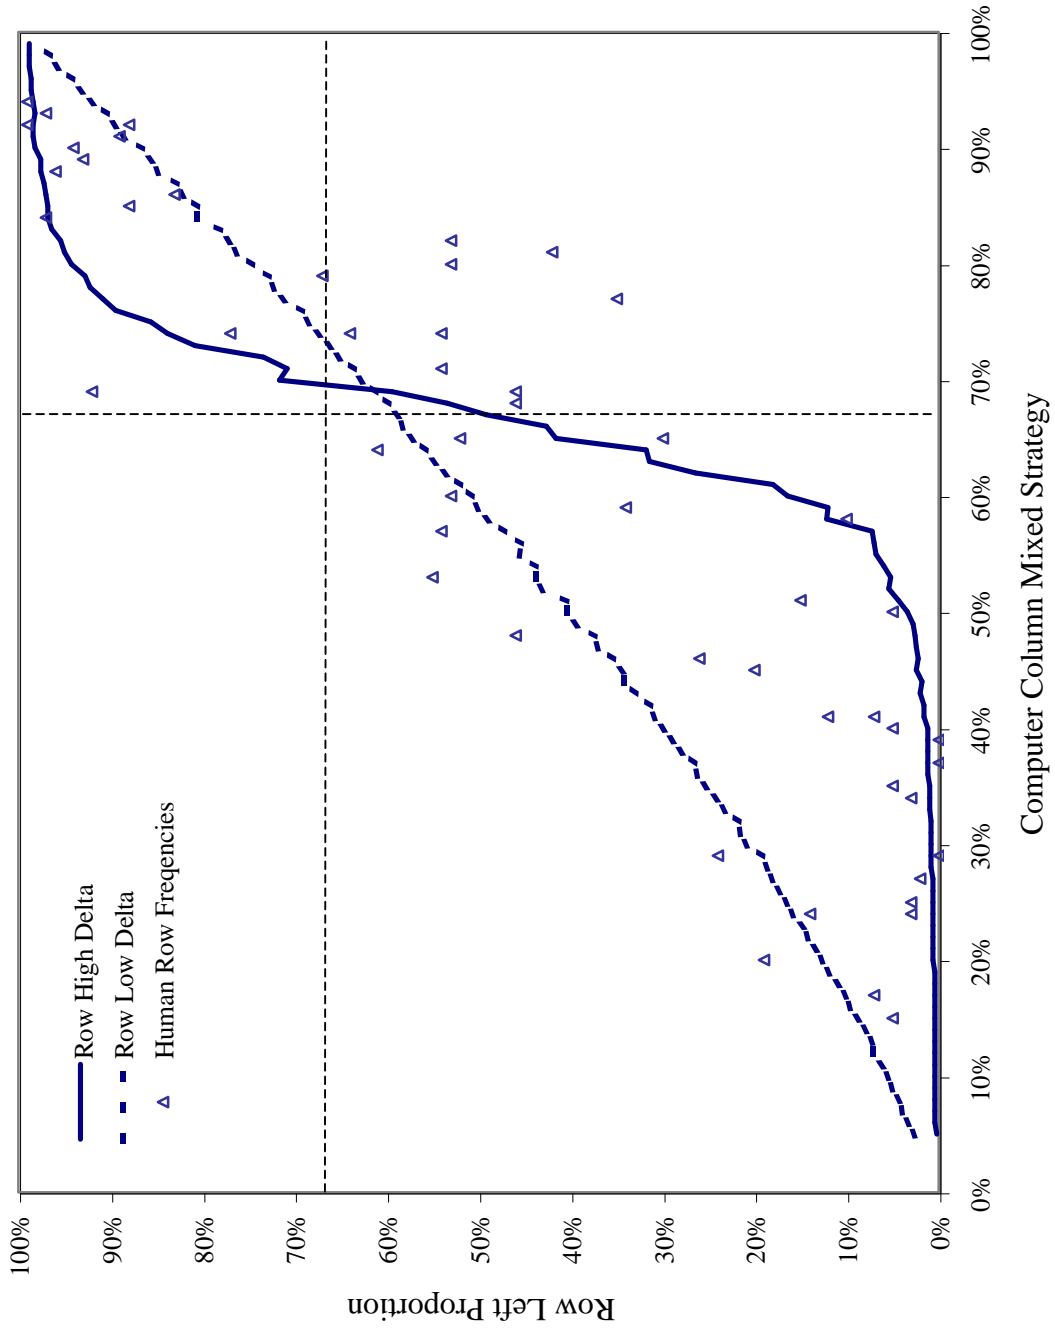


Figure 6: Column Player RHB Response Relationships - Last 100 Stage Games

