COMPARING LEARNING MODELS WITH IDEAL MICRO-EXPERIMENTAL DATA SETS

YAW NYARKO AND ANDREW SCHOTTER

Abstract. This paper compares three learning models. The models compared are the Reinforcement Learning Model of Erev and Roth, the EWA model of Camerer and Ho and a Belief Learning model in which beliefs are elicited from subjects using a proper scoring rule (The Stated Belief Model of Nyarko and Schotter).

We find that the Stated Belief Model outperforms both of the others in dramatic fashion and is capable of predicting not only the behavior of subjects but the period-to-period changes in their behavior in a far superior manner. We suggest that the reason for this performance has to do with the fact that eliciting beliefs and using them as inputs into the belief learning model provides us with the "ideal data set" upon which to compare these models.

1. Introduction

In recent years there has been a number of articles each attempting to present and test a model that describes how individuals learn when playing repeated n-person games. The models investigated are variants of the basic belief-learning model prevalent among economists, the reinforcement learning model favored by psychologists and reintroduced into the economics literature by Erev and Roth (1998) , and more recently, a hybrid Experience Weighted Attraction (EWA) model introduced by Camerer and Ho (1996) which nests the other models a special cases.

One must, however, take care when comparing models empirically. First, it is important to choose the level of aggregation upon which these comparisons will be made. For example, it is ironic is that while all of the above models are formulated as models of individual learning, when they are tested and compared the authors too often aggregate the data, either over time or across individuals, and make their comparisons on the aggregate level.

Second, when comparing models one must make sure that the data used for the comparisons are comparable. More precisely, every model has an ideal data set which, if available, would allow it to perform at its peak efficiency. When such a data set is not available, one is forced to use proxies for those variables that are unobservable. For example, take two models, Model A and Model B, and assume that the ideal data set for Model A is observable to the analyst while this is not the case for Model B. If we were to then compare the performance of these two models and conclude that Model A was better, we would face the problem of not knowing

Date: December 10, 2000.

Key Words: Belief Learning, Game Theory, Experimental Economics
JEL Classification C92, D83, C72

This paper was written under NSF grant # SES-9905227. The authors would also like to thank the C.V. Starr Center for Applied Economics for its financial support. We owe a great deal of thanks to Jay Ezrielev, Meta Brown, and Judy Goldberg for their research assistance.

how much of its performance to impute to the model and how much to impute to the fact that it had its ideal data set available while Model B did not and had to use proxies. First-best comparisons of two models occur only when each has its ideal data set available.

The consequences of these ideas for the comparison of learning models are numerous and significant. For example, reinforcement models of learning are fortunate in often being able to have their ideal data sets observable. This is so because they typically rely only on the observable outcomes of play (actions and payoffs) as inputs into their models. In other words, in order to calculate a player's propensity to choose a given strategy, all that is needed is a time series of that player's payoffs from all previous plays of the game. The EWA model, as written, also has its ideal data set available since it posits a model in which strategies are chose in accordance with their attractions which, as defined by Camerer and Ho (1996) are functions of past observable outcomes. In general, ideal data sets for belief learning models are not available since their basic component, beliefs, are private and hence not observable. To rectify this experimentalists and theorists have used history based proxies as substitutes for these beliefs. These proxies are what we will later call $\gamma - weighted$ empirical beliefs - i.e., beliefs formed by taking the weighted average of the opponent's past play with weights which decline geometrically at the rate γ. In this paper we elicit subject beliefs using a proper scoring rule. If such beliefs are used in place of $\gamma - weighted$ empirical beliefs, they can provide the ideal data sets needed for belief models to be compared to reinforcement and EWA models.

In this paper we aim to compare three models on the basis of individual-by-individual comparisons using ideal data sets. We do this using a data set where subjects play a 2x2 constant sum game repeatedly for 60 periods against either a randomly chosen or fixed opponent. The first model is a belief learning model where the beliefs used are those elicited from subjects at every stage of the experiment using a quadratic scoring rule. We will call this the Stated-Beliefs Model since it uses stated or elicited beliefs as an input into its probabilistic choice rule. It is, however, simply a standard belief-learning model in which beliefs are observable and not estimated using $\gamma - weighted$ empirical beliefs as proxies. The two other models compared are the simple Reinforcement Learning Model of Erev and Roth (1998) and the EWA model of Camerer and Ho (1996).

We demonstrate two things. First, by disaggregating the data down to the individual level we gain new insights into the performance of the models. Second, because previous studies compared the EWA and Reinforcement Models to Belief learning models which used proxies for beliefs instead of direct elicitations, they have tended to draw the seemingly incorrect conclusion that Belief Learning models are inferior. We show the opposite. Not only is the Stated-Belief Model better calibrated, (i.e., a better predictor of both the actions of subjects and the changes in their actions), but it has a better resolution meaning that it makes stronger predictions about what type of actions we can expect to see. These facts are true on a subject-by- subject basis in the sense that out of the 28 subjects in the experiment run here where subjects played a 2x2 constant sum game 60 times with the same opponent (our Experiment 1), the Stated-Belief Model defined a better fit not only to the actions of subjects but also to changes in these actions for 22 of these 28 subjects. The same was true for 20 out of 28 subjects when comparing the Stated Belief and Reinforcement models. We will see shortly why it is important not only to have a model that fits the levels in the data well but also their first differences.

The properties of the Stated Beliefs model were investigated in Nyarko and Schotter (2000). That paper compared the Stated Belief model to the class of belief learning models employing $\gamma-$empirical beliefs as proxies. They demonstated that the Stated-Belief Model outperforms all models in this class. Hence, it represents the "best" belief learning model we are aware of, and therefore will be used in our comparisons with the EWA and Reinforcement models.

We will proceed as follows: In Section 2 we will further motivate our paper by displaying the time path of choices made by selected subjects in our experiments and comparing these time paths with the time paths predicted by our two competing learning models. In Section 3 we quickly review the Stated-Belief, Reinforcement and EWA models and present the experimental design. In Section 4 we present our formal comparison tests of these models using individual-by-individual data. Finally, in Section 5 we offer some conclusions.

2. What Is wrong With These Pictures

To illustrate the need for the micro-comparison of learning models, consider the following eight time paths of actions chosen by eight individual subjects in the 60 period 2x2 constant sum games run by Nyarko and Schotter (2000).

[Figures 1a-1h here]

In the experiment that generated this data subjects played the following constant-sum game for 60 periods.

	Green	**Red**
Green	6,2	3,5
Red	3,5	5,3

This game has a unique mixed strategy equilibrium where each player is suppose to mix his or her use of the Green and Red strategies with probabilities 0.4 and 0.6 respectively. The time series presented in Figure 1a-1d, present the results for four individuals who played this game repeatedly 60 times with the same subject, while Figures 1e-1h present the results for four subjects who played the same game 60 times with randomly drawn opponents.

[Figures 1a-1h here]

In each figure we have placed the round of the experiment on the horizontal axis and the probability with which the subject chooses the Red strategy on the vertical. The x's mark the pure strategy chosen in any given round while circles indicate the predictions of the Erev-Roth (1998) Reinforcement Learning Model for that round. The thick solid line presents the predictions of the EWA model while the thin solid lines present the predictions of the Stated Belief Model. Finally, the dotted line indicates the static Nash equilibrium prediction for the use of the Red strategy.

What is striking about all of these diagrams is the failure of both the EWA and Reinforcement models to track the period-to-period movement of actions for individuals. As we can see subjects change their strategies repeatedly over time as they play the game over their 60 round horizon yet neither the EWA or Reinforcement models capture this movement very well. To the extent that these models fit the data, they seem to do so by passing a relatively straight line through a gyrating time series in an effort to minimize the errors in prediction being made. The Stated-Belief Model, on the other hand, captures the qualitative movements of the data quite well. It predicts the abrupt changes seen in the data with a fair amount of precision. Figure 1c presents a striking example of what we are talking about. In that figure both the EWA and Reinforcement models quickly stop their movement while the Stated Belief model tracks the period-to-period changes in behavior surprisingly well. While this figure presents an extreme example, it is by no means rare. In the other figures, we have chosen patterns that are less pronounced, with Figure 1d presenting the EWA model's best shot. Still, none of the other models are capable of characterizing the movement in the data the way the Stated Belief Model does.

This same pattern, though slightly less pronounced, appears in Figures 1e-1h which uses data from Experiment 2 where subjects were randomly matched. While there are numerous cases where the Stated Belief model is clearly superior in matching the data movements, like the subject in Figure 1f, there are more cases like the subject in Figure 1g where the situation is more ambiguous. Still, as we demonstrate later when we present our goodness-of-fit measures, the Stated Belief model is far superior to either of the other two in its performance.

It is also interesting to include the time path of predicted actions for the Fictitious Play Belief Learning model where the beliefs used to generate predictions are based on $\gamma-$ weighted empirical proxies ($\gamma = 1$) of past actions. These are presented in Figures 1i-1p for both the fixed and random matching experiments.

[**Figures 1i-1l and 1m-1p here**]

As you can see, despite the fact that this is a belief leaning model, because its beliefs are historical constructs that are not free to move from period to period, it too does a relatively poor job of fitting the qualitative aspects of the model. For example, compare Figures 1a and 1i which presents the predictions of our models for Player 2 (Subject 2) in Experiment 1. In Figure 1a we see that for this subject the Stated Belief model does a reasonable job of tracking the movements in the data while, for this same subject, the Fictitious Play model predicts little movement as seen in Figure 1i. Even more striking is the comparison of Figures 1c and 1k (Subject 5). For this subject the Stated Belief model predicts pure strategies almost always while the Fictitious Play model predicts the unending use of a mixed strategy with approximately 50/50 probabilities. Hence it is no wonder that other investigators have concluded that these non-belief learning models are capable of outperforming this type of constrained historical based model.

As a result of these diagrams we can see that each model must be judged on two criteria. One, indicating how well it fits the actual choices made by the subjects (what we call "level calibration") and another indicating how well it captures the qualitative features of the data, in this case the first difference of the time series (what we call "change calibration"). We also compare our three models according to their resolution using the Sander's resolution measure as used by Feltovich (2000) (see Yates (1982)). Which model is best for a given decision maker depends on the relative weights one gives to predicting the levels or the changes in the behavior of one's opponent in his or her objective function. If one's loss function depends heavily on being able to predict changes in behavior from period to period, then a heavy weight would be given to the model's change calibration, while if one only cares about the long-run ability of the model to fit the data, then level calibration may be weighted more heavily. What we show in this paper is that for almost all subjects the Stated Belief Learning Model of Nyarko and Schotter (2000) is the best amongst all of these models on both measures. In short, it dominates the two models it is being compared to. Before we proceed to presenting our results, however, let us pause and quickly describe the learning models we will be comparing.

3. Learning Models and Experimental Designs

In this Section we will briefly outline the essential features of the learning models we compare. These models are the simple Roth-Erev (1998) Reinforcement Model, the Camerer-Ho (1998) EWA model, and the Stated Belief model. Let us outline these models one at a time.

3.0.1. **Learning Models.** Belief-learning and Reinforcement models are motivated by very different views of how people learn. In a Reinforcement Model, at any point in time strategies are chosen with probabilities which are functions of how profitable those strategies have been in the past. Those that have performed relatively well historically are chosen more frequently today. This is the Thorndike's (1898) "law of effect", a basic tenant of how psychologists view learning. Hence, reinforcement models are strictly backward looking models and do not conjecture at all about what strategies are likely to be chosen by one's opponent either today or in the future. In contrast, belief learning models view the process of learning as one in which at each time during the game people form beliefs about one's opponent and then, given these beliefs, choose strategies as an increasing function of their expected payoff. Such models are well-known in economics and include as special cases the Cournot and Fictitious Play models.

All of the learning models we consider here are composed of two parts. One part defines the attraction of a strategy for a player at each point in time while the other part takes these attractions and transforms them into choice probabilities, i.e., probabilities with which these strategies are chosen. Let us investigate the attraction formulations in each of these models first and then move on to their probabilistic choice models.

Let $A_i^j(t)$ denote the attraction of strategy j for player i in period t, (i.e., how appealing that strategy is to player i at that time). In the reinforcement learning model, the attraction of strategy j in period t is simply equal to a linear function of the cumulative payoffs that have accrued to that strategy when it was chosen in the past. More precisely, if we start with some initial attraction for strategy j $A_i^j(0)$, the attraction of strategy j in period t can be define recursively as

$$A_i^j(t) = \begin{cases} A_i^j(t-1) + \pi_i(s_i^j, s_{-i}(t)) & \text{if } s_i^j = s_i(t) \\ A_i^j(t-1) & \text{if } s_i^j \neq s_i(t) \end{cases}, \tag{3.1}$$

where $\pi_i(s_i^j, s_{-i}(t))$ is the payoff that i receives in period t when he or she uses strategy j and his or her opponents use the n-1 tuple, $s_{-i}(t)$. Hence, a strategy is only reinforced in this model in a period if it is actually used. If it is not used, its attractiveness is left unchanged.

In a belief learning model, the attractiveness of a strategy is equal to the payoff that the subject expects to receive if he or she decides to use that strategy. Since expected payoffs are what defines the attractiveness of any strategy, belief learning models are distinguished by how the beliefs upon which these expectations are calculated are defined. More precisely, if beliefs are not elicited then they have to be estimated using the historical data generated in the game. Hence, some model of belief formation is required that will define the probability that a subject thinks that his opponent will use a given strategy j in period t given the history his opponents choices up until period t-1.

A proxy for beliefs often used is class of $\gamma - weighted$ empirical distributions which, following Cheung and Friedman (1997), can be defined as follows: if γ in (-∞, ∞) is the weight which a subject gives in forming his beliefs to the actions of his opponents in previous period, then Player i's γ-weighted empirical beliefs (or, for simplicity, empirical beliefs) are then the sequence defined by

$$b_{it+1}^{j} = \frac{1_t(a^j) + \Sigma_{u=1}^{t-1} \gamma_i^u 1_{t-u}(a^j)}{1 + \sum_{u=1}^{t-1} \gamma_i^u} \tag{3.2}$$

where b_{it+1}^{j} is player i's belief about the likelihood that the opponent will choose action a^j in period t+1, 1(a^j) is an indicator function equal to 1 if a^j was chosen in period t and 0 otherwise, and γ_i^u is the weight given to the observation of action a^j in period t-u.

Fictitious play beliefs are those as above for the special case of $\gamma = 1$ and we call a belief learning model where beliefs are defined using fictitious play beliefs the Fictitious Play Model. We define the Cournot beliefs to be those which assign probability one to opponent's previous period play and the resulting model the Cournot Belief model (i.e. $\gamma = 0$). Finally, if γ were estimated simultaneously with the other parameters in a belief learning model, the resulting estimate of $\gamma, \hat{\gamma}$, would be used to define what we have called $\hat{\gamma}$-beliefs and the associated model the $\hat{\gamma}-$Belief Learning Model.

The Stated-Belief Model is a belief learning model where the beliefs are the elicited or stated beliefs. What Nyarko and Schotter (2000) have demonstrated, however, is their Stated-Belief Model, in which the beliefs used are those elicited using a proper scoring rule, outperforms all models in the class employing $\gamma-$empirical beliefs. Hence, in this paper when we compare belief learning models to either Reinforcement or EWA models, we do so using the "best" belief learning model we are aware of - - the Stated-Belief learning model, a belief-learning model employing an ideal data set since beliefs, if measured truthfully, are observable.

In the EWA model, the attraction of a strategy is specified by two key variables which are updated after each round. The first, N(t) is defined by Camerer and Ho define as "the number of 'observation-equivalents' of past experience" supporting any choice. In some sense, it is the strength of attraction of player i for choosing a strategy. The second variable in the EWA model, $A_i^j(t)$, is player i's attraction to strategy s$_i^j$, in period t. The variable N(t) and $A_i^j(t)$ begin the analysis with some prior values N(0) and $A_i^j(0)$. These reflect the values for these variables based on some preconceived notions of how good the strategies are based either on the players' experience with other, perhaps similar, games and on introspection. As time progresses N(0) and $A_i^j(0)$ are updated. N(0) is updated using the following rule:

$$N(t) = \rho \cdot N(t-1) + 1, t \geq 1. \tag{3.3}$$

$A_i^j(0)$ is updated using the following updating rule:

$$A_i^j(t) = \frac{\phi N(t-1) \cdot A_i^j(t-1) + [\delta + (1-\delta) \cdot I(s_i^j, s_i(t)] \cdot \pi_i(s_i(t), s_{-i}(t))}{N(t)}. \tag{3.4}$$

In this updating function, ϕ is a discount or decay function which depreciates the value of past history-weighted attractions, $N(t-1) \cdot A_i^j(t-1)$, $I(s_i^j, s_i(t))$ is an indicator function which takes the value of 1 if in period t the player's strategy choice $s_i(t)$ equals $s_i^j$and takes on a value of zero otherwise, and δ is a weight placed on the payoff that a player would have received if he had chosen strategy s_i^j when, in fact, he had not. So, in each period, the attraction to a strategy is updated differently depending on whether the strategy was chosen in that period or not. If the strategy was chosen, then after depreciating the attraction of the strategy by $\phi N(t-1) \cdot A_i^j(t-1)$ the attraction of the strategy is incremented by $(1-\delta) \cdot \pi_i(s_i(t), s_{-i}(t)$,i.e. by $(1-\delta)$ times the payoff received. If the strategy is not played, then its attractiveness is still incremented but the weight attached to it is now δ and not $1-\delta$. For the 2x2 model we have here, the EWA results in the following six parameters: $(A_i^{RED}(0), A_i^{GREEN}(0), N(0), \rho, \phi, \delta)$.

It can easily be demonstrated that this EWA model nests both the Reinforcement all $\gamma-$Belief Learning models as special cases and refer the reader to Camerer and Ho (1996) for a proof.

After these attractions are defined for any model, they must be incorporated into a choice function in order to determine the probability of choosing any given strategy in any given round. Two frequently used choice rules are the logit rule

$$P_i^j(t+1) = \frac{e^{\lambda(A_i^j(t))}}{\sum_{k=1}^{m_i} e^{\lambda(A_i^j(t))}}, \tag{3.5}$$

where j is a generic action of player i and m_i are the number of strategies available to player i with each strategy indexed by k;, and the a power function rule,

$$P_i^j(t+1) = \frac{(A_i^j(t))^\lambda}{\sum_{k=1}^{m_i} (A_i^j(t))^\lambda}. \tag{3.6}$$

Once the attractions in any period are defined, each model inserts them into either into either the logit choice function (3.5) or the power function (3.6) to determine choice. In the models tested below, we will use the logit model for our EWA and Stated Belief Learning model[1], and the power function with $\lambda = 1$ for the Reinforcement Model as was done by Erev and Roth (1998). We do this because these are the formulations used by the authors of these models.

[1] To keep the estimates for our Stated Belief model consistent with the Nyarko-Schotter (2000) paper, it is estimated with extra parameters λ^j so that $P_i^j(t+1) = \frac{e^{\lambda^j + \lambda(A_i^j(t))}}{\sum_{k=1}^{m_i} e^{\lambda^j + \lambda(A_i^j(t))}}$. For our 2x2 model with two actions RED and GREEN, since the probabilities must add up to one, this really means that we are adding only one extra parameter $\lambda_0 = \lambda^{RED} - \lambda^{GREEN}$. In our estimates λ_0 was almost always close to 0, and it had negligible effect on our MSD and resolution scores to be reported in section 4.

3.1. **Experimental Design.** The data which we use in this paper were generated by three different sets of experiments run using the experimental laboratory of the C.V. Starr Center for Applied Economics at New York University from the Fall of 1997 through the Summer of 2000. Inexperienced subjects were recruited from undergraduate economics courses and reported to the lab for experiments that took between $1\frac{1}{2}$ and 2 hours. No subjects had any training in game theory. In these experiments subjects played a 2x2 game 60 times with the same opponent under various treatments. Payoffs, which were equal to the sum of the subjects earnings in each round, were denominated in experimental dollars and converted into U.S. dollars at a rate of 1 pt. =$.05. Subjects, on average, earned approximately $15.00 for their participation which was paid to them at the end of the session. They were paid $3.00 simply for showing up.

The 2x2 game used in our experiments was presented above.

The program used to run the experiments was generously supplied to us by Jason Shachat and the Experimental Science Lab of the University of Arizona. [2]

In two separate experiments, Experiments 1 and 2, the identical 2x2 constant sum game was run under different subject-matching protocals. In Experiment 1, subjects were matched with the same opponent in each period for 60 periods while in Experiment 2, they were randomly matched in each round.

Before subjects chose strategies in any round, they were asked to write down on a worksheet a probability vector that they felt represented their beliefs or predictions about the likelihood that their opponent would use each of his of her pure strategies. To elicit truthful reporting of beliefs subjects were paid according to the accuracy of their predictions using a quadratic scoring rule and the procedures employed by Nyarko and Schotter (2000). We refer the reader to that paper for details of belief elicitation in the 2x2 games and to the Appendix of this paper for the actual instructions used to elicit beliefs.

Because critics might suggest that the act of elicitation might alter the behavior of subjects and perhaps lead them to best respond to the beliefs being elicited instead of focusing on their historical beliefs, we ran a third experiment where beliefs are not elicited. This was done as a robustness check. As was pointed out in Nyarko and Schotter (2000) the use of elicitation did not change the behavior of subjects when compared to a control experiment where such elicitation was not done.

Finally, note that our belief learning model is not nested in the EWA model because it does not specify a weighted historical model for belief formation as does EWA. Subjects may be forming their beliefs in ways that are totally unconnected to weighted historical proxies and in that sense our Stated-Belief Model is free to outperform both the EWA and Reinforcement Models as it does.

Table 1 presents our experimental design.

[Table 1 here]

4. Results

The data generated by our experiments define time paths of chosen actions for each of our subjects making choices over 60 rounds each. To compare the results of our three models for each player and each model we construct three indices which

[2]The instructions were computerized and are available upon request from the authors.

describe how well the model is calibrated and resolved on the data generated by that person. Our first measure is simply a mean Mean Squared Deviation (MSD score) calculated individual-by-individual in a particular experiment. We call this the "MSD in levels." To compare learning models on the basis of whether they can predict the **changes** in behavior observed in an accurate manner, we propose to simply use the MSD metric on the data after taking first differences. We will call this the "MSD in changes." Finally, we will compare the results of our models using a measure of their resolution. Resolution attempts to measure how sharp the predictions of the model are in attempting to predict the pure strategies of the subjects.

4.1. **Comparing Models via MSD in Levels and in First differences.** Our first measure is simply a mean Mean Squared Deviation (MSD score) calculated individual-by-individual in a particular experiment. In other words, let $p_{i,t}^m$ be the probability predicted by model m that player i chooses Red in a given round t. Let $a_{i,t}$ be the action chosen by subjects i in round t of the experiment he or she engaged in. In the 2x2 experiments we are discussing, $a_{i,t}$ will denote the probability weight placed on the Red strategy and will take on a value of 1 if Red was chosen and 0 if Green was. Then, the MSD score for individual i and model m, m =1,2,3 (representing the Stated Belief, EWA and Reinforcement model, respectively) is:

$$MSD_i^m = \frac{1}{T} \sum [p_{i,t}^m - a_{i,t}]^2,$$

where T is the number of observations (rounds) for individual i. This is your standard calibration score measuring the goodness of fit of the models to the data (level calibration). [3]

As Figures 1a-1h indicate, however, such a measurement fails to capture the movement in the data in the sense that a model may be well calibrated but achieve its good fit by passing a relatively flat time series of predictions through a constantly moving time series of actions. It may explain the levels in the data but not the changes in it. To correct for this, we take the first difference of both the actual choices made by our subjects and each model's predictions. These first differences record the actual period-by-period change in the choices of subjects and predictions made about these changes. Comparing these two time series indicates whether a model predicts changes in the behavior of subjects and not their levels.

More precisely, $\Delta_{i,t}^a = a_{i,t} - a_{i,t-1}$ represents the change in the choice of subject i between period t and t-1. Given the pure strategy choices of our subjects, Δ_{it}^a can take the values {-1, 0,+1} Similarly, $\Delta_{it}^m = p_{i,t}^m - p_{i,t-1}^m$ represents the change in the predictions of model m about the actions of subjects i between period t-1 and t.

To compare learning models on the basis of whether they can predict the changes in behavior observed in an accurate manner, we propose to simply use the MSD metric on this first difference data as follows:

$$MSD_{\Delta,i}^m = \frac{1}{T} \sum [\Delta_{i,t}^m - \Delta_{i,t}^a]^2.$$

Hence, for any individual i and any model m, we have two goodness-of-fit measures, one measuring levels and one changes in the level of the data.

[3] See Selten (1998) for an axiomatic justification for this MSD score.

The results of our exercises are presented in Tables 2 and Table 3, where we present the mean and individual-by-individual level and change calibration scores for our models, as well as in a series of diagrams that present the results in a more disaggregated manner. Let us look at the means first.

[Table 3 here]

As we can see in Table 2, in terms of the mean level and change MSD's the Stated-Belief Model easily outperforms all of the other models. In fact, it dominates all of the other models we have tested here by having both lower mean level and change MSD scores than any other model. This can easily be seen in Figures 2a and 2b where these mean level and change MSD scores are compared and can be verified using a set of binary Wilcoxon Matched-Pairs Signed-Rank tests presented in Tables 4a -4c which test the null hypothesis that the sample MSD scores of different models come from the same population. We have included in Table 2 not only the EWA and Reinforcement models, but also two belief learning models that use $\gamma-$empirical beliefs. In one of the models, the Fictitious Play Model (FP), the weights attached to previous observations are all equal to 1 ($\gamma = 1$). In the other, the $\hat{\gamma}-$model,the weights used are those estimated jointly with the other parameters in the Logit choice model using Maximum Likelihood techniques. The Stated-Belief Model outperforms these models too.

[Figures 2a and 2b and Tables 4a-4c here]

As Tables 4a-4c indicate, we can reject the hypothesis that the sample of level and change MSD scores calculated from the Stated Belief Model come from the same population as those of any other model at significance levels below 1% for all experiments. As indicated, since the test statistics for these comparisons are all negative, the Stated Belief Model consistently yields lower MSD level and change scores than the other models. For the other models, the results are mixed with us not being able to find consistently different MSD scores between these models at the 5% level.

If there is a second-best learning model it is the EWA model but it is not consistently the second best on either score. Note, however, that all models are more accurate, both in terms of their level and change MSD scores, when used in explaining the data from Experiment 2 where the random matching protocol was used than in Experiment 1 where subjects played against fixed partners. We suspect that the "play the field" environment induced by this treatment environment is an easier one to predict behavior in since it avoids the double and triple guessing that occurs when the same subjects repeatedly play against each other over time.

Another interesting feature of the data is that all of our relevant models perform equivalently when used in experiments with and without elicitation (i.e., Experiment 1 and Experiment 3). In fact we could not find any difference in the performance of our models across these experiments using a series of pair-wise Wilcoxon tests.[4] We take this as evidence that the performance of these models was not changed when elicitation was removed from the experimental protocols as it was in Experiment 3. This result is in line with the results of Nyarko and Schotter (2000) which found no change in behavior resulting from the removal of elicitation.

As Figures 1a-1p indicate there is a great deal of the story missing if one does not dig deeply into the individual-by-individual components of aggregate measures such a mean MSD scores. Table 3 presents the individual level and change MSD scores for all of our subjects in all of our experiments. To demonstrate how superior the Stated Belief model is to the others, consider Figures 3a -3d which presents the MSD level and change scores for each subject in Experiments 1 and 2.

Figures 3a -3d here

As you can see, for almost all subjects the MSD level and change scores for the Stated-Belief Model are lower than they are for any other model. For example, in Experiment 1 for 22 of the 28 subjects the Stated-Belief Model had MSD scores which dominated the EWA model. An identical comparison to the Reinforcement Model demonstrated this was true for 20 of the 28 subjects. For the random matching experiment (Experiment 2) 17 subjects had level and change MSD scores which dominated those of the EWA model while 23 subjects had such scores for the Reinforcement Model.

The picture changes if one compares the performance of the $\hat{\gamma}$ -empirical Belief Learning Model (the best performing $\gamma-$empirical belief-learning model according to Nyarko and Schotter (2000)). In Experiment 1 for only 5 subjects did the $\hat{\gamma}$ model dominate the EWA model. It did so 25 times when compared to the Reinforcement Model, however. For Experiment 2 the results were similar. There were only 9 subjects for whom the $\hat{\gamma}$ model dominated the EWA model while there were 26 subjects for whom the $\hat{\gamma}$ model dominated the Reinforcement model.

To get picture of the relationship of these models on an individual by individual basis, consider Figure 4a- 4d. These figures present a pair-wise comparison of our learning models for each experiment. In each figure, we place the MSD_i score for each individual on the horizontal axis and the corresponding $MSD_i[\Delta]$ scores on the vertical. If we then take an individual each model will define a vector in this $MSD_i - MSD_i[\Delta]$ space depicting how well that model does for that individual on both of these measures. For each experiment and each pair of models we present 28 such vectors.

[Figures 4a-4d here]

As you can see in these figures, there is a distinct clustering of MSD_i and $MSD_i[\Delta]$ scores for the EWA and Reinforcement models in the Northeast section of the figures indicating that these models generate component wise higher score vectors in the $MSD_i - MSD_i[\Delta]$ space.

---

[4]

	MSD Difference Exp 1 vs. Exp 3	
	Elicitation vs No Elicitation	
	MSD Level	**MSD Change**
FP	W=1.346, p=.178	W=1.33, p = .182
$\hat{\gamma}$	W=-.190, p=.849	W=.063, p=.949
EWA	W=-1.968, p=.049	W=.444, p=.656
Rein	W =1.549, p=.1213	W=1.346, p=.1783

In summation, it should be clear that the Stated-Belief Model has dramatically outperformed both the EWA of the Reinforcement learning models not only in its abilities to match the levels in the data but their changes as well.

4.2. **Comparing Models via Resolution.** Finally we will compare the results of our models using a measure of their resolution. Resolution attempts to measure how sharp the predictions of the model are in attempting to predict the pure strategies of the subjects. For example, in our experiment where subjects make binary strategy choices (Red or Green), a model is well resolved if its predictions are at the extreme ends of the unit interval of probability estimates concerning an opponents use of Red. A model that predicts the use the Red strategy with probability of 0.5 all the time, whatever its calibration, is a poorly resolved model since its predictions are fairly uninformative about the one-period-ahead actions of an opponent. Following Feltovich (2000) we use Sanders resolution metric (see Yates (1982):

$$R_i^j = (\frac{1}{T})\sum_{t=1}^{T}[pred(j) \times (1 - pred(j))].$$

Note that a model is better resolved for person i the lower is that person's R_i. Hence, a learning model will be penalized if it does not go out on a limb and predict that a strategy will be used with high probability. Learning models that always predict probabilities of .50/.50 for 2x2 games will have the worst resolution and a Sander's score of .25. Those that repeatedly predict the use of pure strategies will be best resolved although they may have poor calibration if they are consistently wrong. Hence, it is very possible that relatively well calibrated models can be poorly resolved and visa versa.

As we have seen in Figures 1a-1p, EWA, Reinforcement and Fictitious Play models all make predictions which vary only slightly from period to period and, more importantly, tend to be in the middle of the unit interval defining the probability that a subject will choose he Red strategy in our 2x2 game. In short, they tend to give little indication of the actions of opponents since they offer little precision in their probability estimates. The Stated-Belief Model, on the other hand, because it does not use historical proxies for beliefs, makes strong predictions with probabilities at the extreme ends of the unit interval quite often. To illustrate this point, we present Table 5 and Figures 5a and 5b which present the scores of each subject using the Sanders's resolution index. For the actual predictions of our subjects see Appendix A which provides the predictions of each model for each individual in each experiment.

[**Table 5 and Figures 5a and 5b here**]

As can be seen in either Table 5 or Figures 5a and 5b, the Stated-Belief Model is a far better resolved model than either the EWA or Reinforcement Model. For example, the Sander's index is smaller for the Stated-Belief Model for 22 of the 28 subjects in Experiment 1 and 20 of the 28 subjects in Experiment 2. Using a set of pair-wise Matched-Pairs Signed-Rank Wilcoxon tests comparing the Stated Belief to the EWA and Reinforcement models we can reject the null hypothesis that the indices for the Stated-Belief Model come from the same population as either the indices for the EWA or Reinforcement models at levels below 1%. In addition, the

EWA demonstrates a better resolution than does the reinforcement model in all comparisons. [5]

In summation, the Stated Belief model has proven to be a far better resolved model than either the EWA of reinforcement models when tested on individual data generating a Sanders resolution index.

5. Conclusions

This paper has attempted to make what we feel is a precise and rigorous comparison of three well discussed learning models: The Reinforcement model of Erev and Roth (1998) , the EWA model of Camerer and Ho (1996) , and the Stated belief-learning model of Nyrako and Schotter (2000) where the beliefs used are those elicited from the subjects using a proper scoring rule. We have made our comparisons on a micro level comparing these models individual by individual and round by round.

Our most striking result indicates that both the EWA and Reinforcement models are poor predictors of the qualitative nature of the time series of subjects actions. While individual actions over the 60 round horizon of the experiment are tremendously variable and capable of changing direction dramatically from period to period, only the Stated Belief Model is capable of tracking these changes. It does so because the beliefs elicited from the subjects are themselves quite variable and to the extent that subjects best-respond to these beliefs, their actions vary as well. For both the EWA and Reinforcement model, predicted behavior is very smooth and changes little from period to period.

We have found that at the micro level of disaggregation, (and all level of aggregation above it) the Stated Belief Model outperforms both of the other models both in its ability to predict the behavior of subjects and the changes in their behavior. It is also a better resolved model. We feel that the Stated Belief model achieves such a decisive victory here, while its performance has been less successful elsewhere, because we have been able to compare these models using the ideal data sets that are required. More precisely, others have failed to pick up the superiority of belief learning models because, due to the absence of Stated or elicited beliefs in their experiments, they have had to rely on weighted historical proxies as belief surrogates. What we have shown here is that these proxies are inadequate and that the process of belief formation is a complex one that can not be easily captured by merely weighting history using geometrically declining weights.

---

[5]

Wilcoxon Tests: Resolution Experiment 1
Format = Test Statistic, p-value

Model	Stated Beliefs	EWA	Reinforcement
Stated Beliefs	–	-3.689, 0.0002	-4.623, 0.000
EWA		–	-4.623, 0.0000
Reinforcement			–

Wilcoxon Tests: Resolution Experiment 2
Format = Test Statistic, p-value

Model	Stated Beliefs	EWA	Reinforcement
Stated Beliefs	–	-2.209, 0.0272	-4.600, 0.000
EWA		–	-4.623, 0.0000
Reinforcement			–

Finally, our results suggest that if one is going to make progress in describing human behavior in strategic situations, it will be necessary to properly model the process of belief formation. While we assign this task to future research we have learned that any model that is to be successful will have to depart from the type of history-based models that generate smooth and non-varying belief time series.

6. Bibliography

References

[1] Camerer, C. and Ho, T. H., "Experience -weighted Attraction Learning in Games: A Unifying Approach", mimeo, California Institute of Technology, 1996.

[2] Cheung, Y. W., Friedman, D., "Individual Learning in Normal Form Games: Some Laboratory Results", **Games and Economic Behavior**, vol., 19, 1997, pp. 46-76.

[3] Erev, I. and Roth, A., "Predicting How People Play Games: Reinforcement Learning in Experimental Games With Unique Mixed Strategies" **American Economic Review,** vol. 88, (1998), pp. 848-881.

[4] Feltovich, N. "Reinforcement-Based vs. Belief Based Learning Models in Experimental Asymmetric-Information Games", **Econometrica**, vol. 68, No. 3, May 2000, pp. 605-641.

[5] Fudenberg, D. and Levine, D., Theory of Learning in Games, Cambridge MA. MIT Press. 1998.

[6] Nyarko, Yaw and Schotter, Andrew, "An Experimental Study of Beliefs Learning Using Elicited Beliefs", mimeo, New York University, Revised October 2000.

[7] O'Neill, B., "Non-metric Test of the Minimax Theory of Two Person Zero-sum Games", **Proc. Natl. Acad. Sci. USA**, vol. 84, 1987, pp. 2106-2109.

[8] Selten, R.,"Axiomatic Characterization of the Quadratic Scoring Rule", Experimental Economics, Vol. 1, No., 1, 1998, pp. 43-62.

[9] Shachat, J., "Mixed Strategy Play and the Minimax Hypothesis", UCSD Economics Discussion Paper 96-37, University of California at San Diego, November 1996.

[10] Thorndike, E.L., "Animal Intelligence: An Experimental Study of the Associative Processes in Animals", Psychological Monographs **, (1898)** 2.

[11] Yates, J.P., "External Correspondence: Decompositions of the Mean Probability Score", **Organizational Behavior and Human Performance**, vol. 30, (1982), pp. 132-156.

Table 1: Experimental Design			
Experiment	**Matching**	**Elicitation**	**Number of Subjects**
1	fixed	yes	28
2	random	yes	28
3	fixed	no	26

Table 2: Mean Level and Change MSD's						
			Level MSD's			
			Models			
Experiment		**SB**	**FP**	$\hat{\gamma}$	**EWA**	**Reinforcement**
1	*Mean*	.1261	.2281	.206	.198	.247
	Std. Dev	.079	.027	.056	.025	.005
2	*Mean*	.121	.198	.182	.170	.235
	Std. Dev	.075	.051	.049	.050	.016
3	*Mean*	*NA*	.222	.216	.207	.243
	Std. Dev	*NA*	.026	.022	.032	.012
			Change MSD's			
			Models			
Experiment		**SB**	**FP**	$\hat{\gamma}$	**EWA**	**Reinforcement**
1	*Mean*	.255	.512	.436	.477	.527
	Std. Dev	.167	.135	.126	.140	.139
2	*Mean*	.219	.354	.320	.344	.370
	Std. Dev	.126	.123	.124	.124	.128
3	*Mean*	*NA*	.452	.440	.452	.460
	Std. Dev	*NA*	.100	.100	.098	.099

Table 3: Individual Level and Change MSD's

Experiment 1: Level MSD

	Model				
Player	**SB**	**FP**	**gamma ^**	**EWA**	**Reinforcement**
1	0.0336	0.1057	0.0649	0.0834	0.2358
2	0.1464	0.2311	0.2303	0.1523	0.2381
3	0.1396	0.2439	0.2319	0.1781	0.2468
4	0.0133	0.2258	0.2309	0.1904	0.2380
5	0.0277	0.2473	0.0557	0.2008	0.2522
6	0.0793	0.2348	0.2353	0.2060	0.2406
7	0.2391	0.2275	0.2259	0.2099	0.2519
8	0.2464	0.2408	0.2495	0.2134	0.2477
9	0.0368	0.2423	0.2083	0.2105	0.2472
10	0.0920	0.2275	0.2034	0.2122	0.2500
11	0.0726	0.2168	0.0541	0.1989	0.2505
12	0.0233	0.2110	0.2180	0.2007	0.2505
13	0.1002	0.2309	0.2076	0.2002	0.2471
14	0.0832	0.2180	0.2198	0.2016	0.2472
15	0.1758	0.2327	0.2437	0.2037	0.2495
16	0.1733	0.2372	0.2409	0.2049	0.2523
17	0.1164	0.2456	0.2080	0.2052	0.2487
18	0.2487	0.2218	0.2218	0.2051	0.2489
19	0.0914	0.2328	0.1676	0.2036	0.2500
20	0.0000	0.2474	0.2332	0.2049	0.2500
21	0.2049	0.2488	0.2351	0.2065	0.2526
22	0.0526	0.2094	0.1906	0.2057	0.2357
23	0.1120	0.2465	0.2315	0.2065	0.2552
24	0.2017	0.2495	0.2497	0.2076	0.2517
25	0.1986	0.2470	0.2426	0.2090	0.2512
26	0.2450	0.2398	0.2348	0.2095	0.2470
27	0.2253	0.2150	0.2230	0.2101	0.2398
28	0.1508	0.2091	0.2137	0.2093	0.2420
Mean	0.1261	0.2281	0.2061	0.1982	0.2471
std Dev.	0.0794	0.0272	0.0552	0.0255	0.0055

Experiment 1: Change MSD

Model

	Player	**SB**	**FP**	**gamma ^**	**EWA**
			Reinforcement		
1	0.0680	0.2099	0.1503	0.1685	0.2545
2	0.2786	0.4816	0.4608	0.4738	0.4771
3	0.1904	0.4807	0.4595	0.4706	0.4769
4	0.0312	0.3588	0.3599	0.4117	0.4097
5	0.0564	0.9247	0.1330	0.9286	0.9243
6	0.1538	0.4309	0.4215	0.4247	0.4263
7	0.5636	0.5551	0.5548	0.5454	0.5644
8	0.5864	0.6005	0.6070	0.5929	0.6153
9	0.1006	0.3690	0.4180	0.3972	0.3759
10	0.2319	0.7118	0.5812	0.7115	0.7168
11	0.1385	0.7778	0.1184	0.1910	0.8846
12	0.0468	0.4604	0.4410	0.4893	0.4953
13	0.2027	0.5129	0.4575	0.4452	0.5443
14	0.1796	0.4272	0.3866	0.4488	0.4262
15	0.2961	0.4432	0.4252	0.4711	0.4439
16	0.4178	0.6450	0.6280	0.6187	0.6661
17	0.2425	0.5594	0.4531	0.4506	0.5615
18	0.4408	0.4371	0.4369	0.4184	0.4425
19	0.1857	0.5075	0.3860	0.3869	0.5434
20	0.0000	0.5074	0.5060	0.4986	0.5116
21	0.4300	0.4910	0.4854	0.4885	0.4943
22	0.1074	0.4630	0.4075	0.3824	0.4761
23	0.1874	0.5353	0.5481	0.5466	0.5463
24	0.3795	0.4905	0.4895	0.5014	0.4961
25	0.3828	0.5397	0.5546	0.5437	0.5466
26	0.5249	0.5346	0.4898	0.4823	0.5281
27	0.4153	0.4060	0.4044	0.4081	0.4278
28	0.3110	0.4796	0.4561	0.4782	0.4788
Mean	0.2553	0.5122	0.4364	0.4777	0.5270
std Dev.	0.1675	0.1350	0.1263	0.1396	0.1389

Experiment 2: Level MSD

			Model		
Player	**SB**	**FP**	**gamma ^**	**EWA**	**Reinforcement**
1	0.0162	0.2194	0.2074	0.2120	0.2471
2	0.0859	0.1599	0.1224	0.1250	0.2083
3	0.1882	0.2254	0.1972	0.1980	0.2503
4	0.2490	0.1679	0.1749	0.1495	0.2467
5	0.2191	0.1983	0.1980	0.2014	0.2430
6	0.0333	0.1844	0.1818	0.1639	0.2254
7	0.1141	0.1179	0.1217	0.1193	0.2142
8	0.1298	0.2604	0.1365	0.1319	0.2479
9	0.0360	0.2040	0.1933	0.1940	0.2409
10	0.1867	0.1917	0.1620	0.1762	0.2386
11	0.1235	0.2400	0.2383	0.2255	0.2420
12	0.1955	0.1943	0.1874	0.1696	0.2166
13	0.0803	0.2272	0.2183	0.2228	0.2422
14	0.2411	0.2494	0.2222	0.2097	0.2516
15	0.0932	0.0663	0.0629	0.0609	0.1954
16	0.0437	0.1150	0.1000	0.1151	0.2174
17	0.1166	0.1960	0.1743	0.1818	0.2432
18	0.1120	0.1183	0.1099	0.1171	0.2153
19	0.2165	0.2396	0.2389	0.0729	0.2511
20	0.1207	0.2370	0.2324	0.1739	0.2498
21	0.1578	0.2166	0.1655	0.1627	0.2426
22	0.1692	0.2461	0.2149	0.1966	0.2486
23	0.0376	0.2475	0.2484	0.2389	0.2495
24	0.2430	0.2410	0.2309	0.2460	0.2532
25	0.0000	0.2322	0.2255	0.2256	0.2438
26	0.0613	0.2135	0.2207	0.2266	0.2350
27	0.1089	0.1135	0.1140	0.0924	0.1998
28	0.0159	0.2213	0.2103	0.1704	0.2330
Mean	0.1213	0.1980	0.1825	0.1707	0.2354
std Dev.	0.0757	0.0507	0.0494	0.0506	0.0168

Experiment 2: Change MSD

	Model				
Player	**SB**	**FP**	**gamma ^**	**EWA**	**Reinforcement**
1	0.0173	0.4428	0.3750	0.3619	0.4764
2	0.1845	0.2370	0.1467	0.2215	0.2374
3	0.3403	0.4520	0.3850	0.3658	0.5107
4	0.3998	0.3316	0.2574	0.2694	0.4075
5	0.4705	0.4328	0.4312	0.4205	0.4428
6	0.0678	0.2718	0.2735	0.2743	0.2725
7	0.1708	0.1618	0.1613	0.1692	0.1695
8	0.2609	0.3570	0.2437	0.2579	0.4761
9	0.0728	0.4943	0.4077	0.4833	0.4801
10	0.3035	0.3633	0.3445	0.3566	0.3751
11	0.2321	0.4746	0.4770	0.4762	0.4786
12	0.3564	0.3633	0.3433	0.3786	0.3589
13	0.1836	0.5767	0.5715	0.5768	0.5812
14	0.4518	0.4554	0.4227	0.4200	0.4609
15	0.1248	0.0953	0.0920	0.1104	0.1356
16	0.0751	0.2266	0.1915	0.2273	0.2548
17	0.2153	0.3756	0.3421	0.3678	0.3910
18	0.2064	0.2343	0.2357	0.2349	0.2554
19	0.1925	0.1927	0.1444	0.1499	0.1869
20	0.2075	0.2696	0.2600	0.3196	0.2728
21	0.2982	0.3771	0.2914	0.3683	0.4084
22	0.3440	0.3874	0.3339	0.3853	0.3916
23	0.0441	0.4222	0.4207	0.4597	0.4264
24	0.5935	0.6052	0.5406	0.5819	0.5965
25	0.0000	0.4924	0.4618	0.5016	0.4949
26	0.1176	0.3862	0.3884	0.4037	0.3929
27	0.1773	0.1798	0.1790	0.1853	0.1700
28	0.0314	0.2676	0.2466	0.3178	0.2723
Mean	0.2193	0.3545	0.3203	0.3445	0.3706
std Dev.	0.1492	0.1268	0.1239	0.1238	0.1277

Experiment 3: Level MSD

			Model		
Player	**SB**	**FP**	**gamma ^**	**EWA**	**Reinforcement**
1	Na	0.2192	0.2183	0.2077	0.2459
2	Na	0.2429	0.1945	0.1852	0.2480
3	Na	0.1543	0.2148	0.0933	0.2398
4	Na	0.2402	0.2274	0.2212	0.2493
5	Na	0.2160	0.2078	0.2043	0.2422
6	Na	0.2489	0.2365	0.2310	0.2504
7	Na	0.2363	0.2336	0.2301	0.2451
8	Na	0.1854	0.1795	0.1806	0.2348
9	Na	0.2278	0.2183	0.2156	0.2465
10	Na	0.1457	0.1458	0.1413	0.1963
11	Na	0.2026	0.1915	0.1876	0.2360
12	Na	0.2374	0.2339	0.2407	0.2483
13	Na	0.2258	0.2077	0.2042	0.2411
14	Na	0.2448	0.2322	0.2277	0.2502
15	Na	0.2322	0.2212	0.2216	0.2428
16	Na	0.2399	0.2268	0.2128	0.2489
17	Na	0.2290	0.2384	0.2323	0.2486
18	Na	0.2330	0.2241	0.2309	0.2496
19	Na	0.2124	0.2111	0.1847	0.2432
20	Na	0.2424	0.2272	0.2273	0.2511
21	Na	0.2380	0.2408	0.2316	0.2407
22	Na	0.2236	0.1939	0.1939	0.2428
23	Na	0.2175	0.2172	0.2146	0.2433
24	Na	0.2125	0.2244	0.2144	0.2385
25	Na	0.2204	0.2149	0.2195	0.2526
26	Na	0.2397	0.2475	0.2369	0.2521
Mean	Na	0.2218	0.2165	0.2073	0.2434
std Dev.	Na	0.0256	0.0218	0.0323	0.0108

Experiment 3: Change MSD

		Model			
Player	**SB**	**FP**	**gamma ^**	**EWA**	**Reinforcement**
1	Na	0.4983	0.4466	0.5005	0.4949
2	Na	0.4084	0.3921	0.3650	0.4080
3	Na	0.2029	0.2029	0.2037	0.2205
4	Na	0.5290	0.5243	0.5244	0.5281
5	Na	0.3523	0.3465	0.3892	0.3742
6	Na	0.5738	0.5625	0.5539	0.5789
7	Na	0.4895	0.4983	0.4914	0.4941
8	Na	0.4044	0.3995	0.3881	0.4095
9	Na	0.4487	0.4478	0.4433	0.4778
10	Na	0.3558	0.3557	0.3492	0.3590
11	Na	0.4097	0.4009	0.3952	0.4265
12	Na	0.5917	0.5900	0.5791	0.5977
13	Na	0.4309	0.4507	0.4301	0.4269
14	Na	0.5305	0.4813	0.5287	0.5299
15	Na	0.3716	0.3287	0.4048	0.3744
16	Na	0.4117	0.3782	0.4180	0.4081
17	Na	0.4655	0.4804	0.4831	0.5109
18	Na	0.5592	0.5304	0.5591	0.5622
19	Na	0.3171	0.3129	0.2974	0.3231
20	Na	0.4353	0.3890	0.4155	0.4422
21	Na	0.5538	0.5510	0.5589	0.5637
22	Na	0.3636	0.3551	0.3744	0.3752
23	Na	0.5556	0.5568	0.5540	0.5646
24	Na	0.4631	0.4445	0.4564	0.4615
25	Na	0.3853	0.3736	0.4414	0.3933
26	Na	0.6453	0.6490	0.6460	0.6672
Mean	Na	0.4520	0.4403	0.4519	0.4605
std Dev.	Na	0.0993	0.1007	0.0979	0.0992

Table 4a: Experiment 1, Pairwise Wilcoxon										
Tests on MSD Levels and Changes										
Format = Level MSDs Test Statistic, p-value										
Change MSDs Test Statistic, p-value										
		Stated								
Model		**Beliefs**		**Fictitious Play**		**þ**		**EWA**		**Reinforcement**
Stated		–		-4.281, 0.0000		-4.190, 0.0000		-3.644, 0.0003		-4.623, 0.0000
Beliefs		–		-4.486, 0.0000		-4.258, 0.0000		-4.304, 0.0000		-4.623, 0.0000
Fictitious				–		2.516, 0.0119		4.600, 0.0000		-4.623, 0.0000
Play				–		3.416, 0.0006		1.503, 0.1329		-3.530, 0.0004
þ						–		2.357, 0.0184		-4.600, 0.0000
						–		-1.480, 0.1388		-4.0531, 0.0001
EWA								–		-4.623, 0.0000
								–		-2.926, 0.0034
Rein-										–
forcement										–

Table 4b: Experiment 2, Pairwise Wilcoxon										
Tests on MSD Levels and Changes										
Format = Level MSDs Test Statistic, p-value										
Change MSDs Test Statistic, p-value										
		Stated								
Model		**Beliefs**		**Fictitious Play**		**þ**		**EWA**		**Reinforcement**
Stated		–		-3.757, 0.0002		-2.869, 0.0041		-2.368, 0.0179		-4.600, 0.0000
Beliefs		–		-3.803, 0.0001		-2.209, 0.0272		-3.188, 0.0014		-4.190, 0.0000
Fictitious				–		3.735, 0.0002		4.031, 0.0001		-4.418, 0.0000
Play				–		4.076, 0.0000		0.956, 0.3389		-3.029, 0.0025
þ						–		1.537, 0.1243		-4.623, 0.0000
						–		-3.621, 0.0003		-4.463, 0.0000
EWA								–		-4.623, 0.0000
								–		-2.118, 0.0342
Rein-										–
forcement										–

Table 4c: Experiment 3, Pairwise Wilcoxon

Tests on MSD Levels and Changes

Format = Level MSDs Test Statistic, p-value

Change MSDs Test Statistic, p-value

Model		**Stated Beliefs**		**Fictitious Play**		**bγ**		**EWA**		**Reinforcement**
Stated		–		–		–		–		–
Beliefs		–		–		–		–		–
Fictitious				–		2.603, 0.0092		4.076, 0.0000		-4.457, 0.0000
Play				–		2.883, 0.0039		0.673, 0.5009		-3.518, 0.0004
bγ						–		2.998, 0.0027		-4.432, 0.0000
						–		-1.257, 0.2087		-3.975, 0.0001
EWA								–		-4.457, 0.0000
								–		-2.273, 0.0230
Rein-										–
forcement										–

Table 5: Resolutions: Sander's Index

Experiment 1			Experiment 2		
	Model			Model	
SB	**EWA**	**Rein**	**SB**	**EWA**	**Rein**
0.040128	0.00162	0.24897	0.016249	0.193653	0.249225
0.160112	0.221357	0.247738	0.082693	0.13068	0.243672
0.142182	0.231357	0.249568	0.196508	0.197993	0.249945
0.013335	0.228338	0.247058	0.24896	0.137417	0.248785
0.008216	0.242555	0.2496	0.219205	0.202733	0.249182
0.083514	0.23273	0.245569	0.033325	0.168828	0.247585
0.239147	0.234448	0.249826	0.102849	0.096108	0.246418
0.246426	0.239628	0.247974	0.126931	0.138465	0.249404
0.044507	0.19311	0.249098	0.037134	0.193417	0.247801
0.096339	0.22743	0.249906	0.185074	0.141285	0.24699
0.068273	0.0219	0.248506	0.130496	0.224247	0.247232
0.031574	0.221235	0.24977	0.195492	0.16381	0.243833
0.111395	0.187373	0.248741	0.084915	0.223335	0.24722
0.078763	0.22356	0.24982	0.241068	0.212022	0.249327
0.181971	0.232868	0.249513	0.08527	0.067428	0.243139
0.180495	0.223675	0.249838	0.043251	0.122437	0.245701
0.116545	0.210695	0.249314	0.109702	0.186095	0.24821
0.248704	0.19134	0.24978	0.118201	0.123038	0.244369
0.019341	0.172802	0.249903	0.214975	0.084875	0.246193
3.12E-05	0.230325	0.249883	0.117947	0.158885	0.249762
0.205509	0.238512	0.249526	0.154316	0.153062	0.247475
0.052631	0.186038	0.246511	0.169687	0.194602	0.249629
0.111707	0.209988	0.247145	0.032943	0.239277	0.249754
0.203307	0.232457	0.249686	0.243048	0.246473	0.249787
0.200403	0.240938	0.249859	1.80E-07	0.226373	0.249065
0.244968	0.224197	0.248911	0.059295	0.225855	0.245084
0.225276	0.226902	0.246418	0.108491	0.095787	0.242178
0.150002	0.174028	0.247978	0.015874	0.17013	0.246493

Appendix 1:
Experiment 1: Belief Elicitation Instructions

You are about to partake in an experiment on decision making. Various research foundations have given money to support this research and depending upon the choices you make you may be able to earn a significant amount of money which we will arrange to pay you at the end of the experiment.

All of the instructions for the experiment will be presented to you on the computer terminal. However, we are making one modification to these instructions which we will explain to you after you have read the instructions. Please read the computerized instructions now and wait for further instructions when you have finished.

Instructions

The experiment you will perform will be identical to the one you have just had explained to you except for one change. Before each round we will ask you to write down on a worksheet your prediction about how many cards you expect your opponent to allocate to each choice he/she has. In addition to your earnings from the game we will pay you an extra amount depending upon how good your prediction is about your opponent.

Predicting Other People's Choices

At the beginning of the decision problem, before you make your choise of green or red, you will be given an opportunity to earn additional money by predicting the choices of your pair member in the decision problem.

For example, the game you will be playing will have two different strategies labeled green, red. Before each round we will ask you to predict what the probability is that you think your pair member will choose either one of his/her two choices -- green or red.

To make a prediction we will supply you with a prediction as follows:

Prediction Form

Probability your pair member chooses green	Probability your pair member chooses red

This form allows you to make a prediction **of the choice of your pair member** by indicating what the chances are that your pair member will choose red or blue. For example, say you think there is a 10% chance that your pair member will choose red, and hence a 90% chance that green will be chosen. This indicates that you believe that red is less likely to be chosen than green by a considerable margin. If this is your belief about the likely choice of your pair member, then write 10 next to the red entry and 90 next to the blue entry. Note that the numbers you write must sum up to 100. For example, if you think there is a 67% chance that your pair member will choose red and a 33% chance he/she will choose blue, write 67 in the space next to the red entry and 33 in the space next to the green entry.

At the end of the decision problem, we will look at the choice actually made by your pair member and compare his/her choice to your predictions. We will then pay you for your prediction as follows:

Suppose you predict that your pair member will choose Green with a 90% chance and Red with a 10% chance. In that case you will place 90 next to the Green entry and 10 next to the Red entry. Suppose now that your pair member actually chooses Red. In that case your payoff will be **Prediction Payoff = $[2-(1-0.10)^2 - (0.90)^2]$** . In other words, we will give you a fixed amount of 2 points from which we will subtract an amount which depends on how inaccurate your prediction was. To do this when we find out what choice your pair member has made we will we will take the number you assigned to that choice, in this case 10% on Red, subtract it from 100% and square it. We will then take the number you assigned to the choice not made by your pair member, in this case the 90% you assigned to Green, and square it also. These two squared numbers will then be subtracted from the 2 points we initially gave you to determine your final point payoff. Your point payoff will then be converted into francs at the rate of 1 point = $0.05.

Note that the worst you can do under this payoff scheme is to state that you believe that there is a 100% chance that a certain action is going to be taken and assign 100% to that choice when in fact the other choice is made. Here your payoff from prediction would be 0. Similarly, the best you can do is to guess correctly and assign 100% to that choice which turns out to be the actual choice chosen. Here your payoff will be 2 or $0.10.

<u>However since your prediction is made before you know what your pair member actually will choose, the best thing you can do to maximize the expected size of your prediction payoff is to simply state your true beliefs about what you think you pair member will do. Any other prediction will decrease the amount you can expect to earn as a prediction payoff.</u>

Prediction Worksheet

Round	Green	Red
1		
2		
.	.	.
.	.	.
.	.	.
60	.	.

Figure 1a: Model Predictions: Experiment 1, Player 2

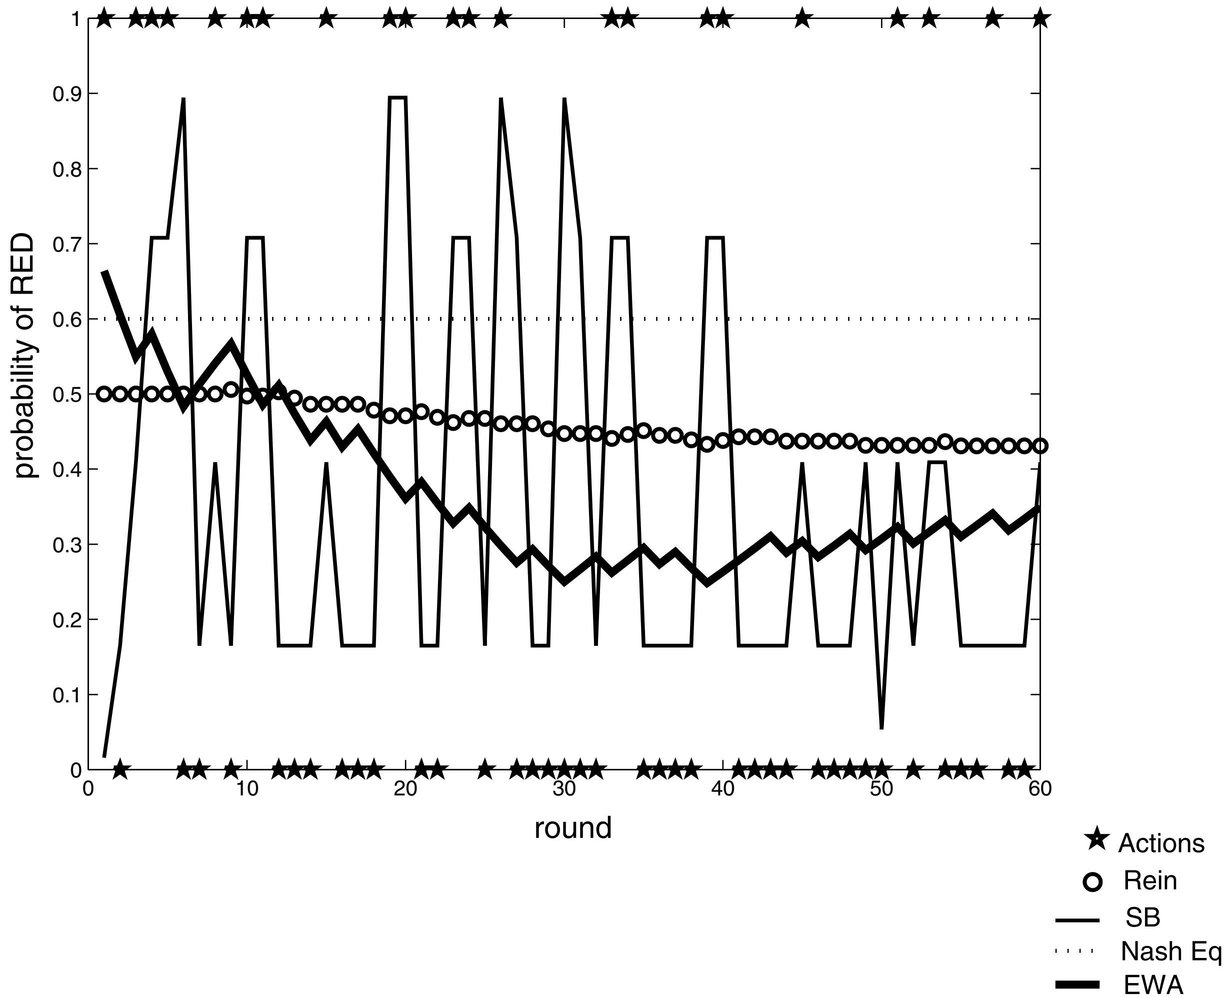

Figure 1b: Model Predictions: Experiment 1, Player 4

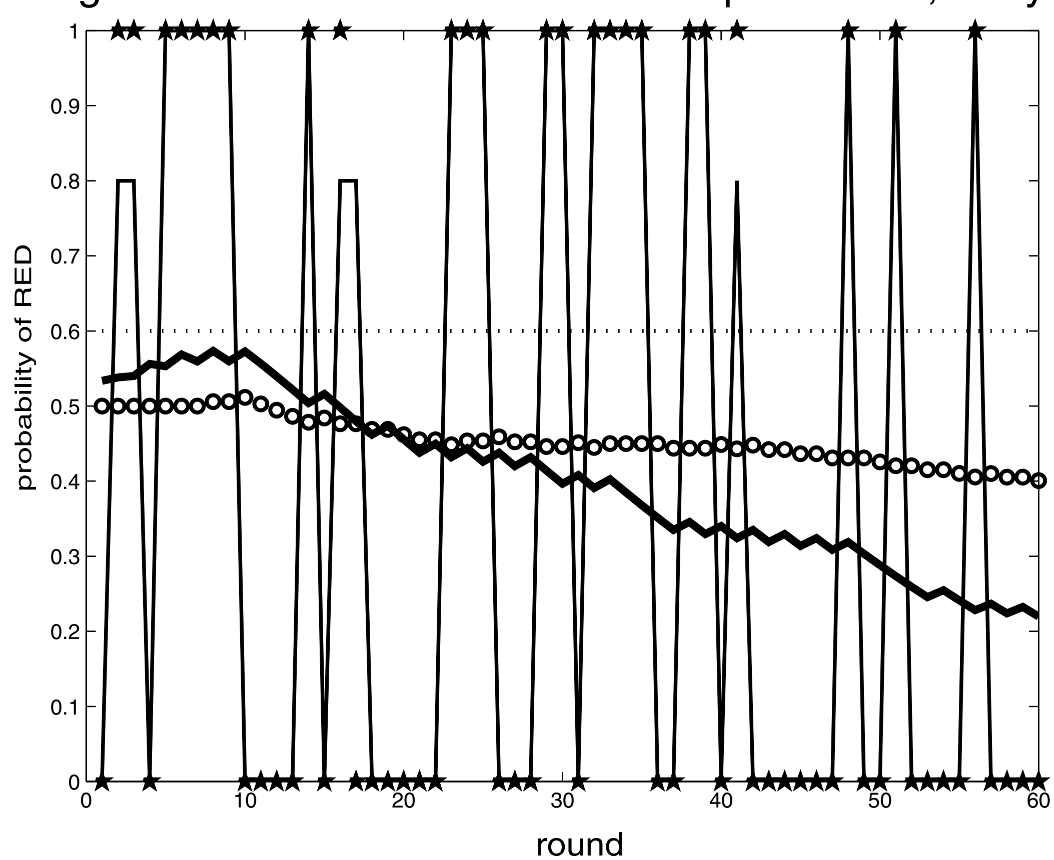

Figure 1c: Model Predictions: Experiment 1, Player 5

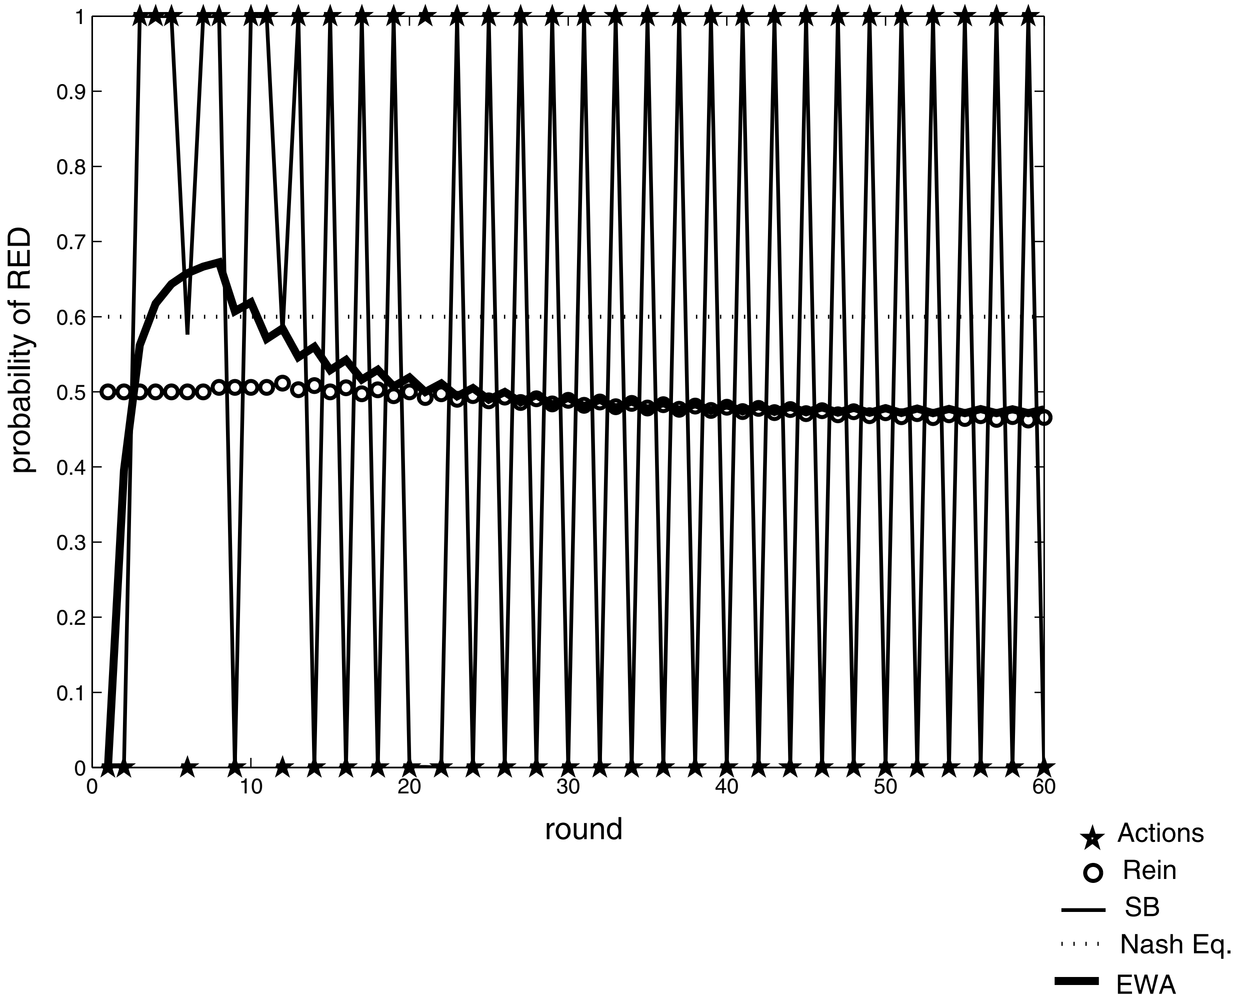

Figure 1d: Model Predictions: Experiment 1, Player 9

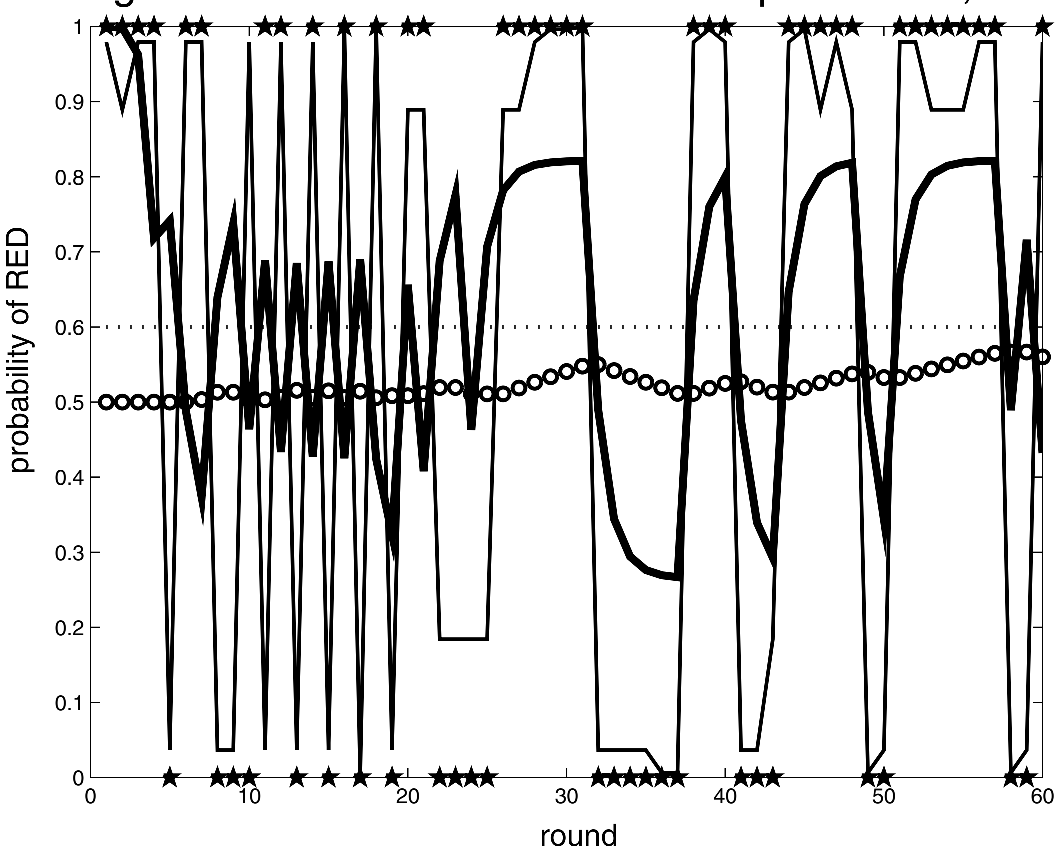

Figure 1e: Model Predictions: Experiment 2, Player 9

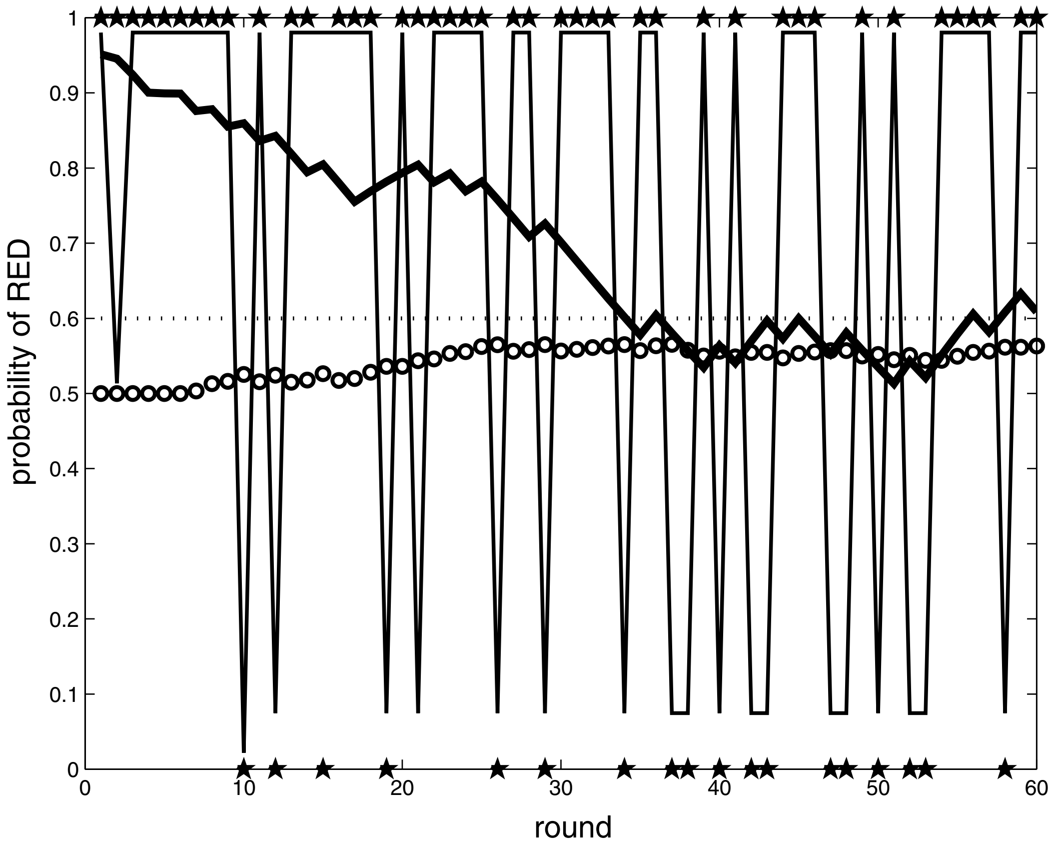

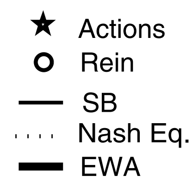

Figure 1f: Model Predictions: Experiment 2, Player 13

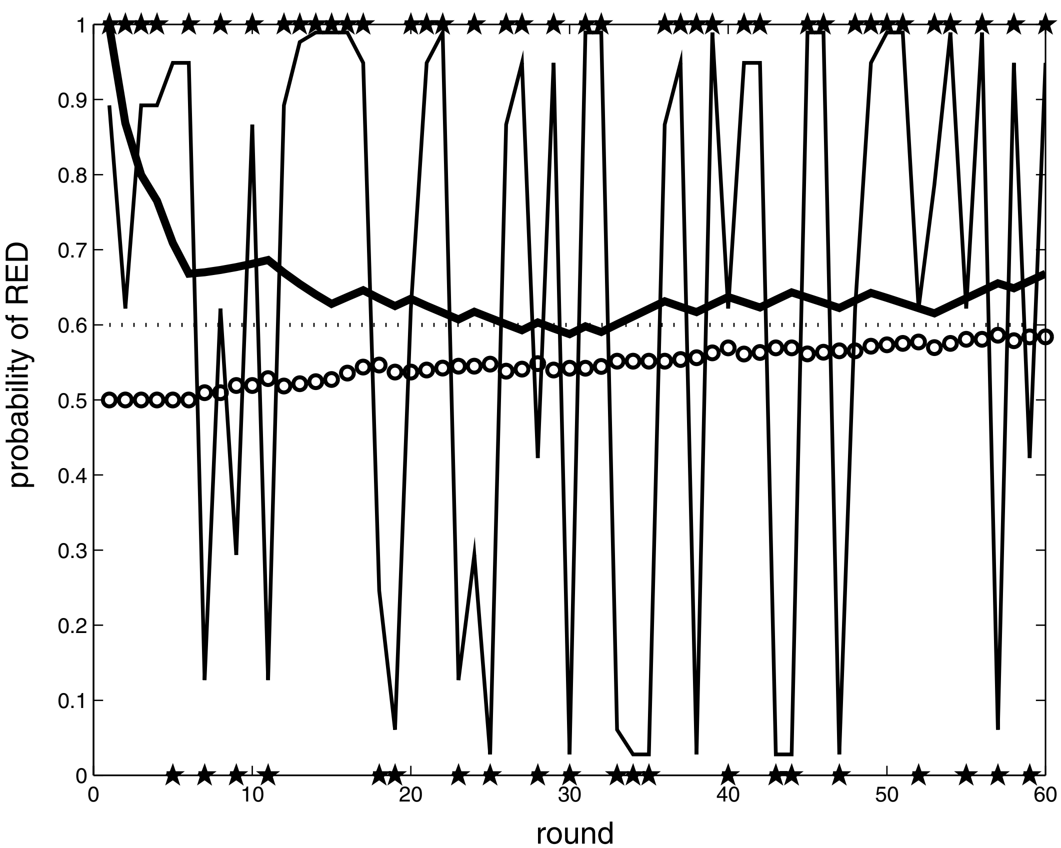

Figure 1g: Model Predictions: Experiment 2, Player 22

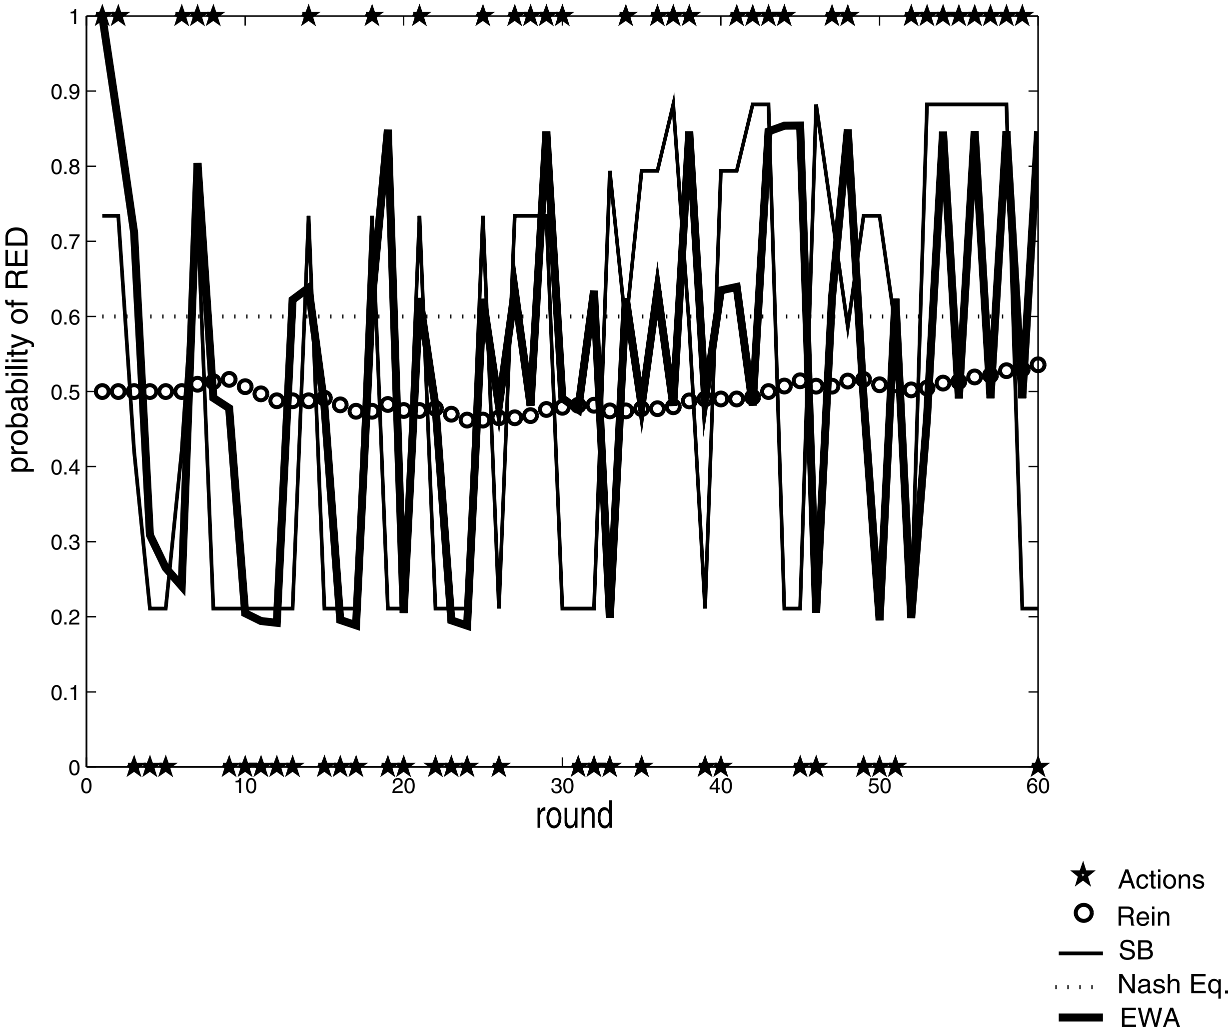

Figure 1h: Model Predictions: Experiment 2, Player 23

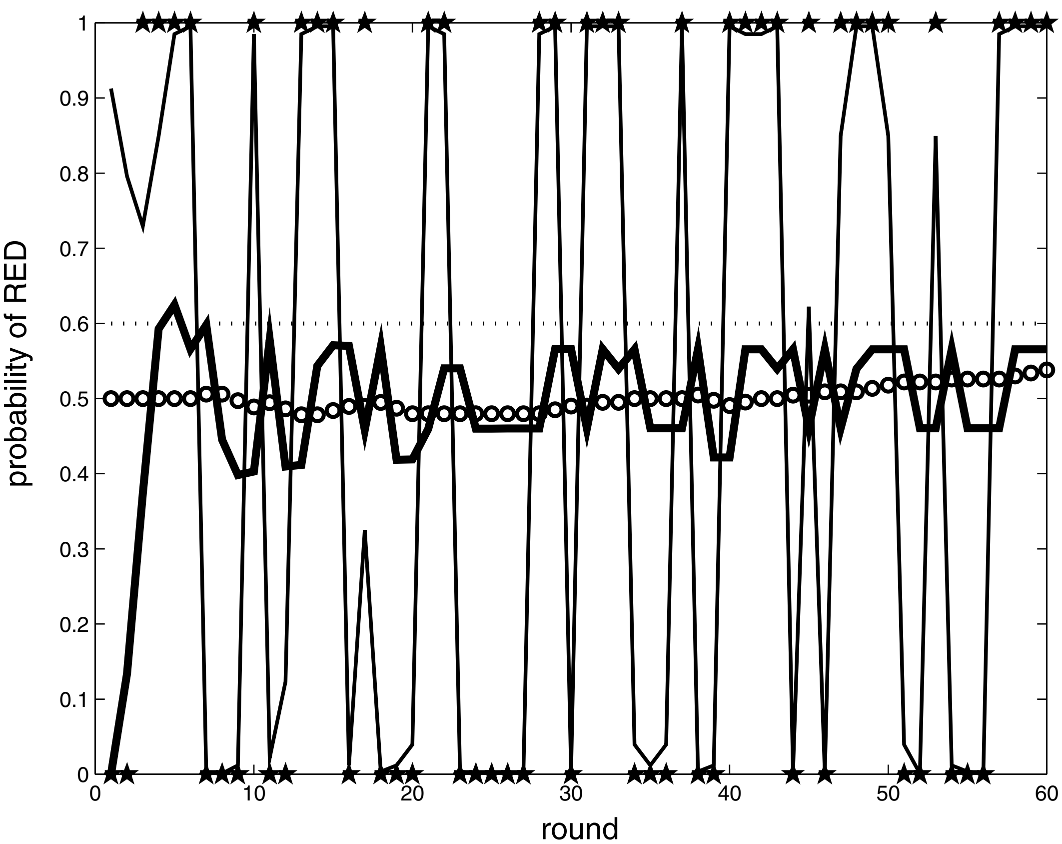

Figure 1i: Model Predictions: Experiment 1, Player 2

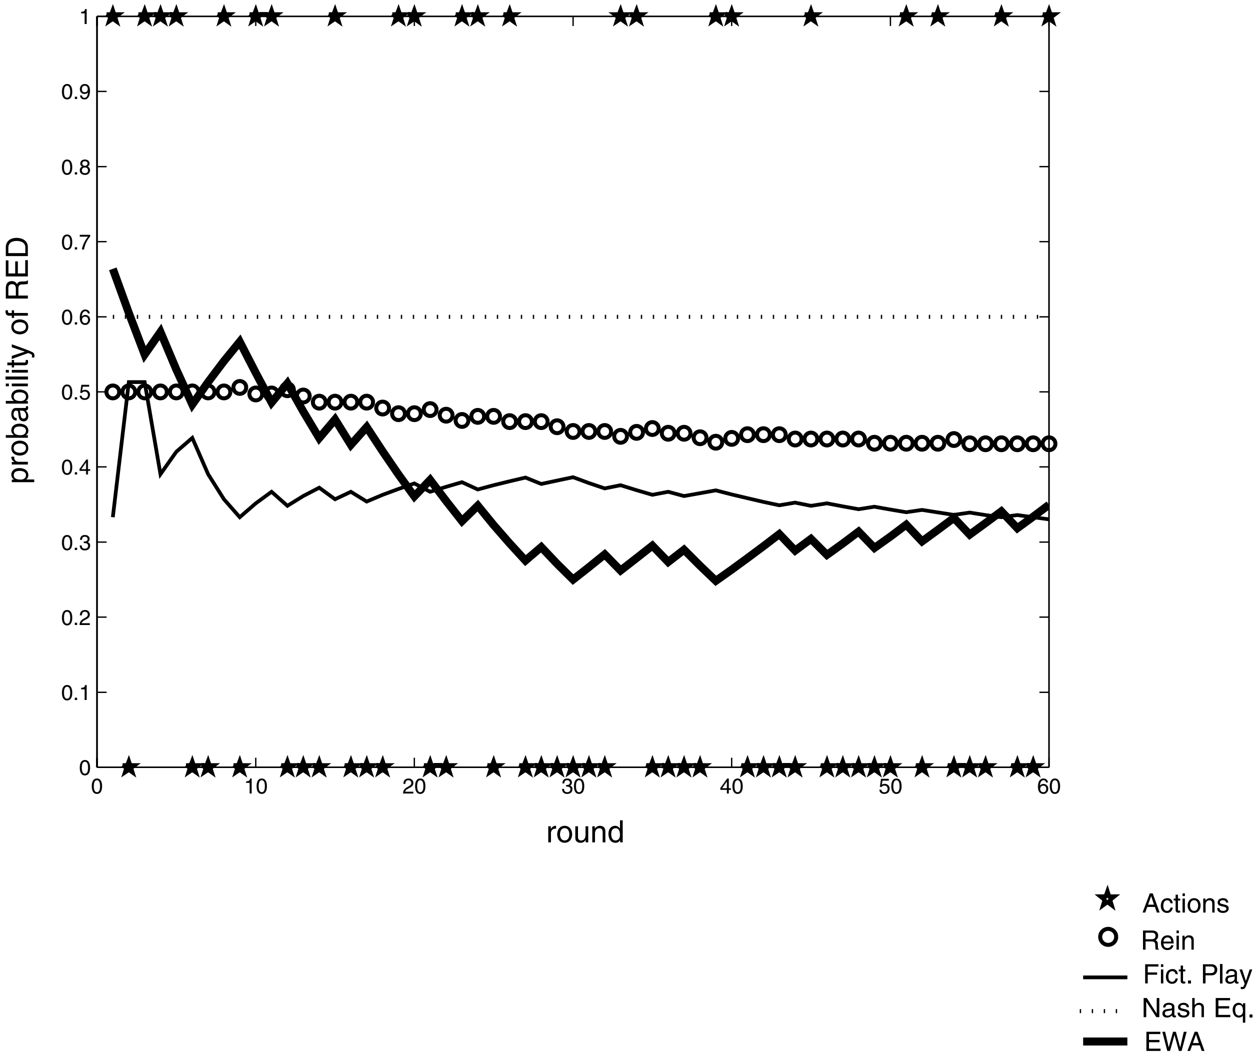

Figure 1j: Model Predictions: Experiment 1, Player 4

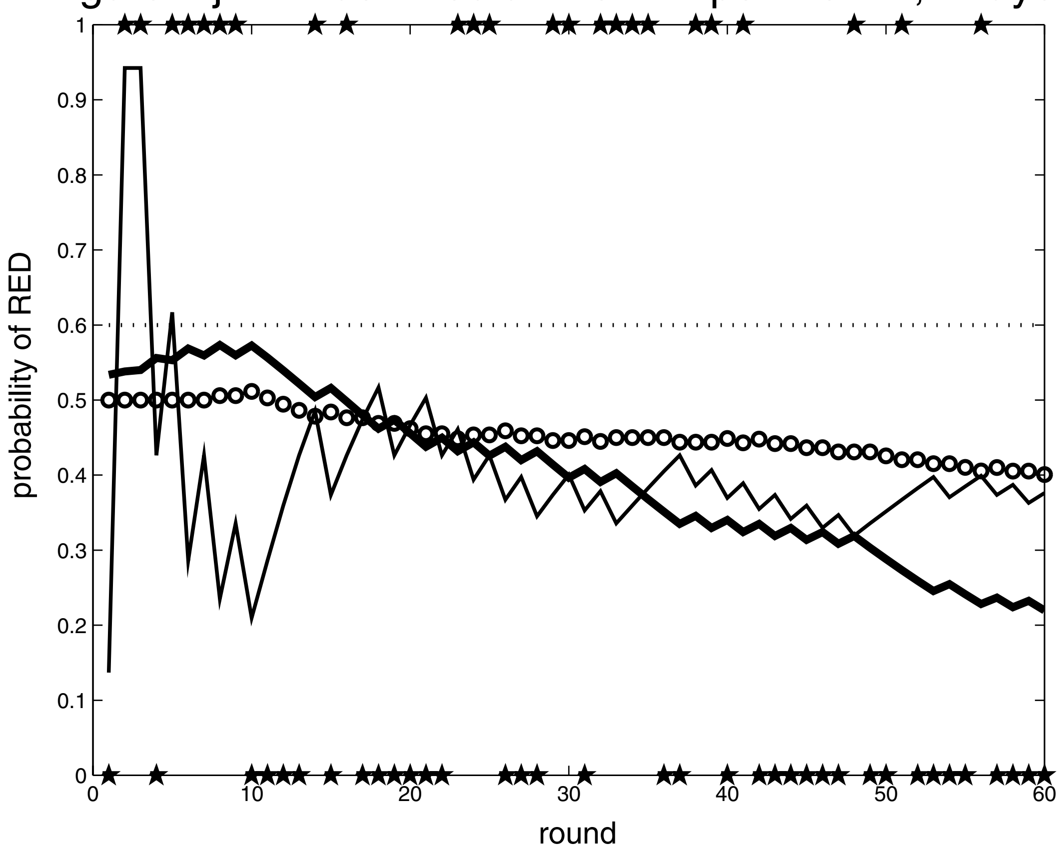

Figure 1k: Model Predictions: Experiment 1, Player 5

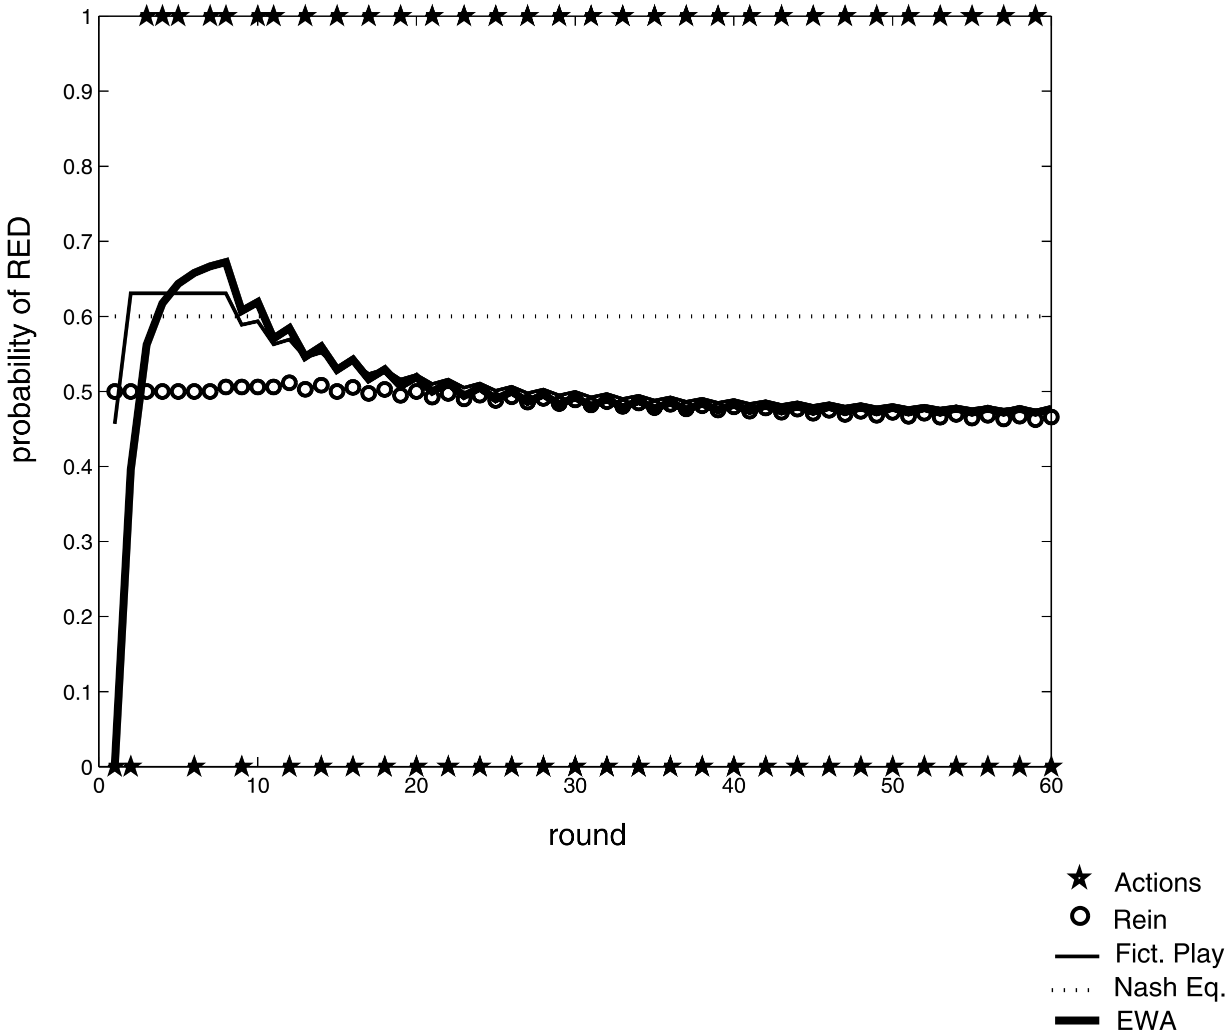

Figure 1l: Model Predictions: Experiment 1, Player 9

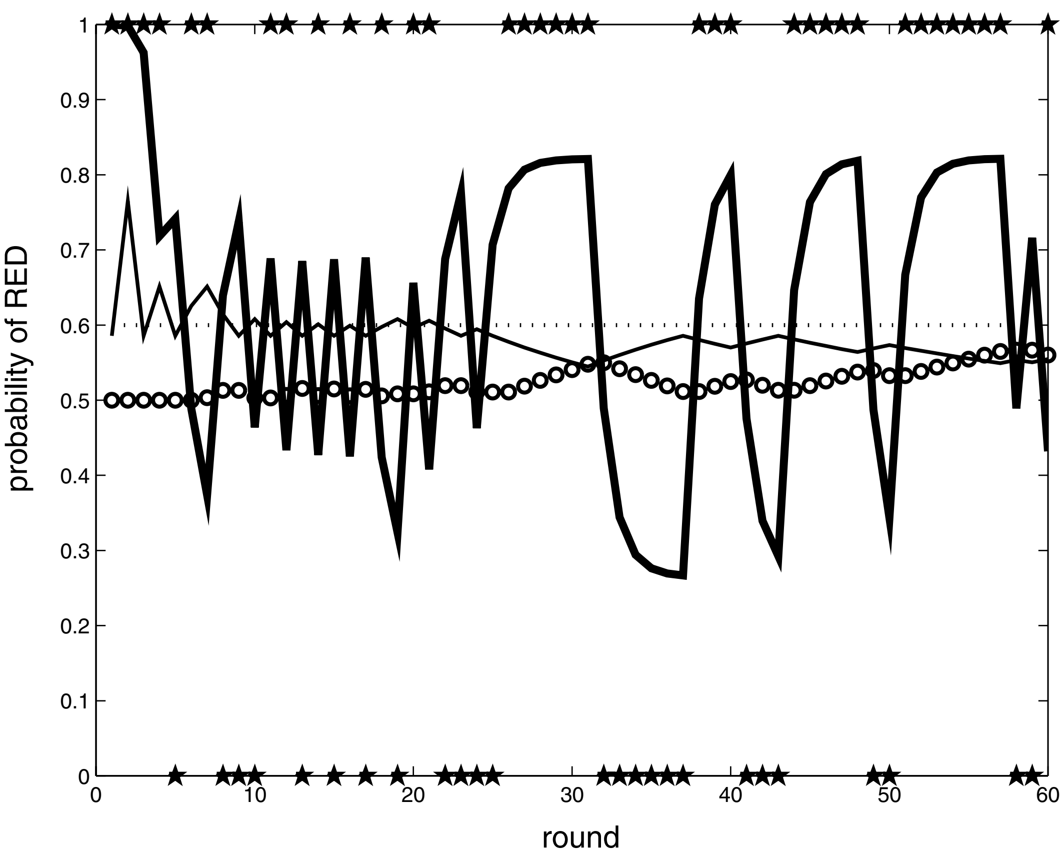

Figure 1m: Model Predictions: Experiment 2, Player 9

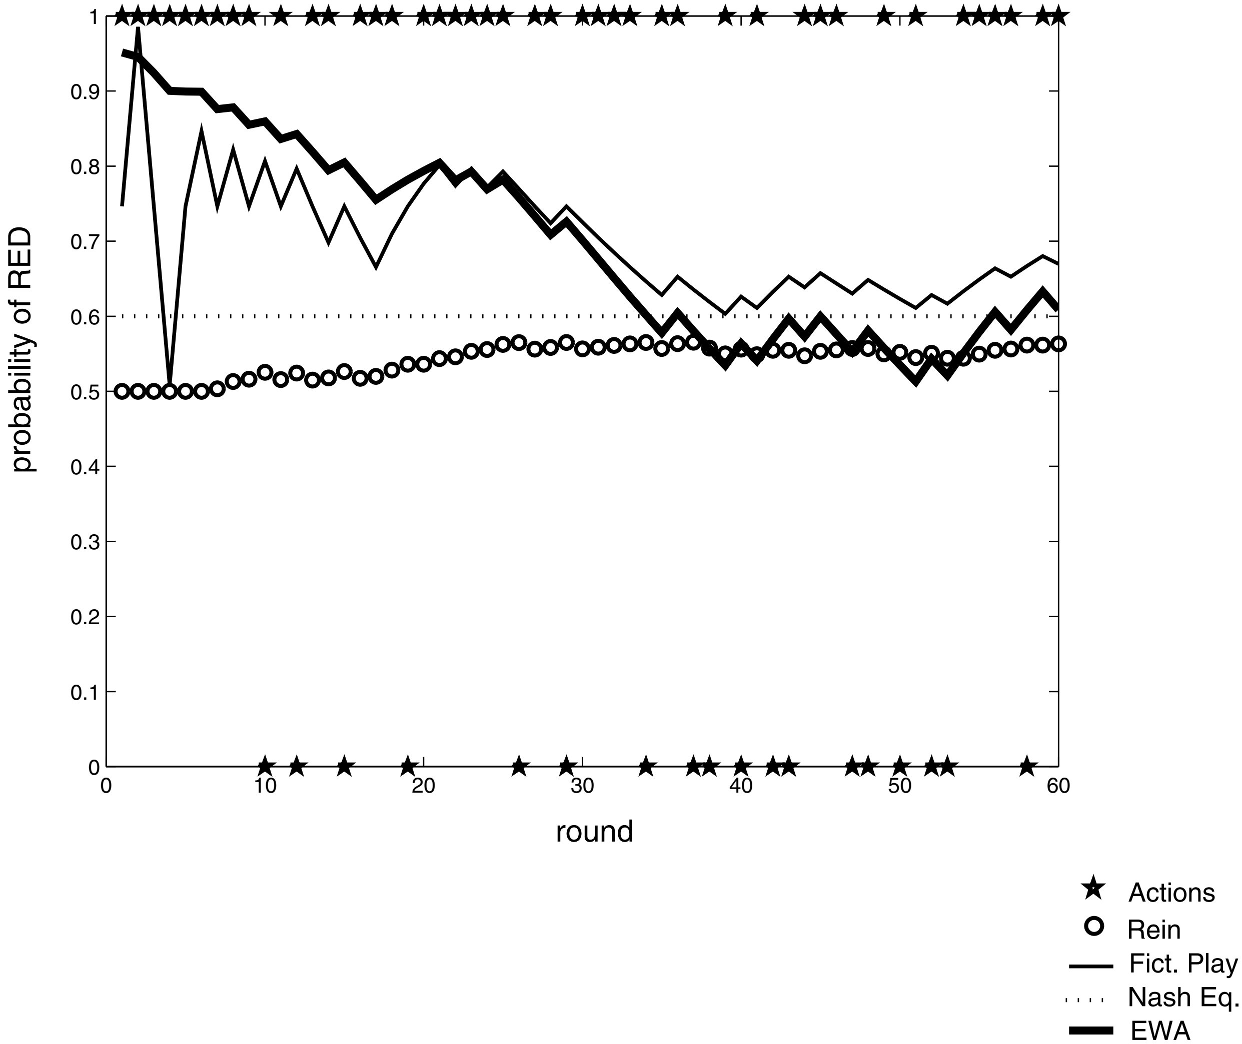

Figure 1n: Model Predictions: Experiment 2, Player 13

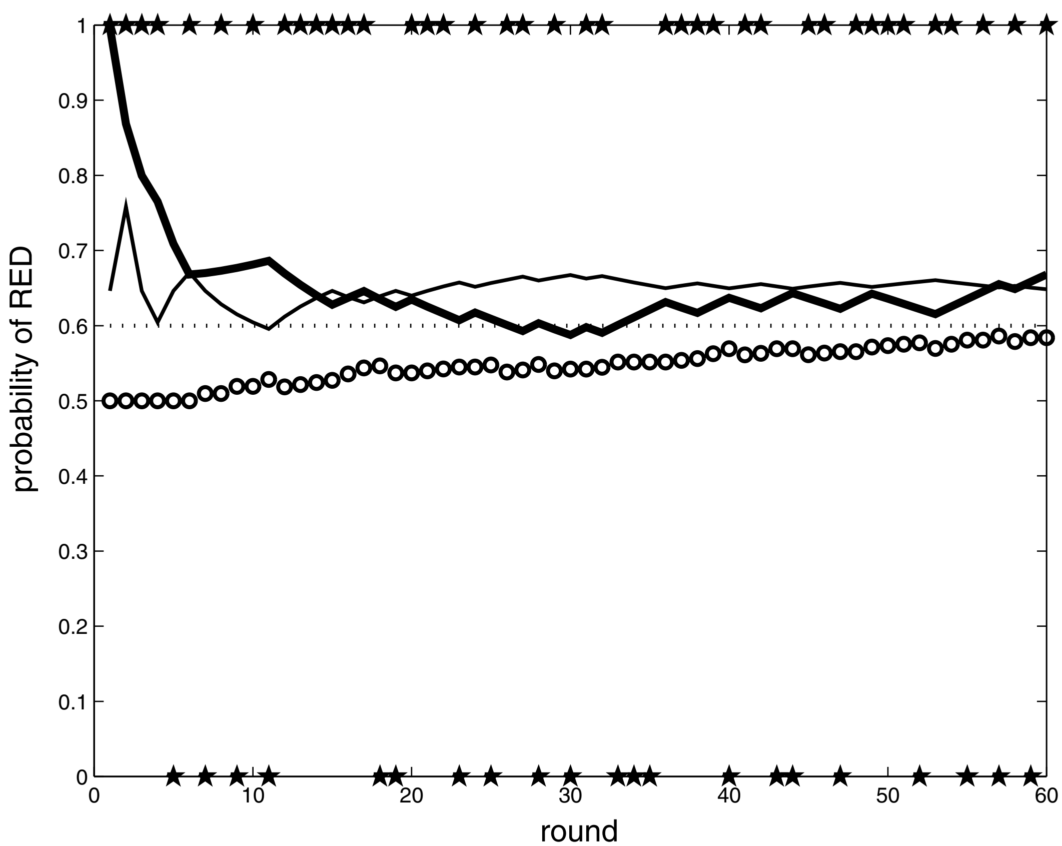

Figure 1o: Model Predictions: Experiment 2, Player 22

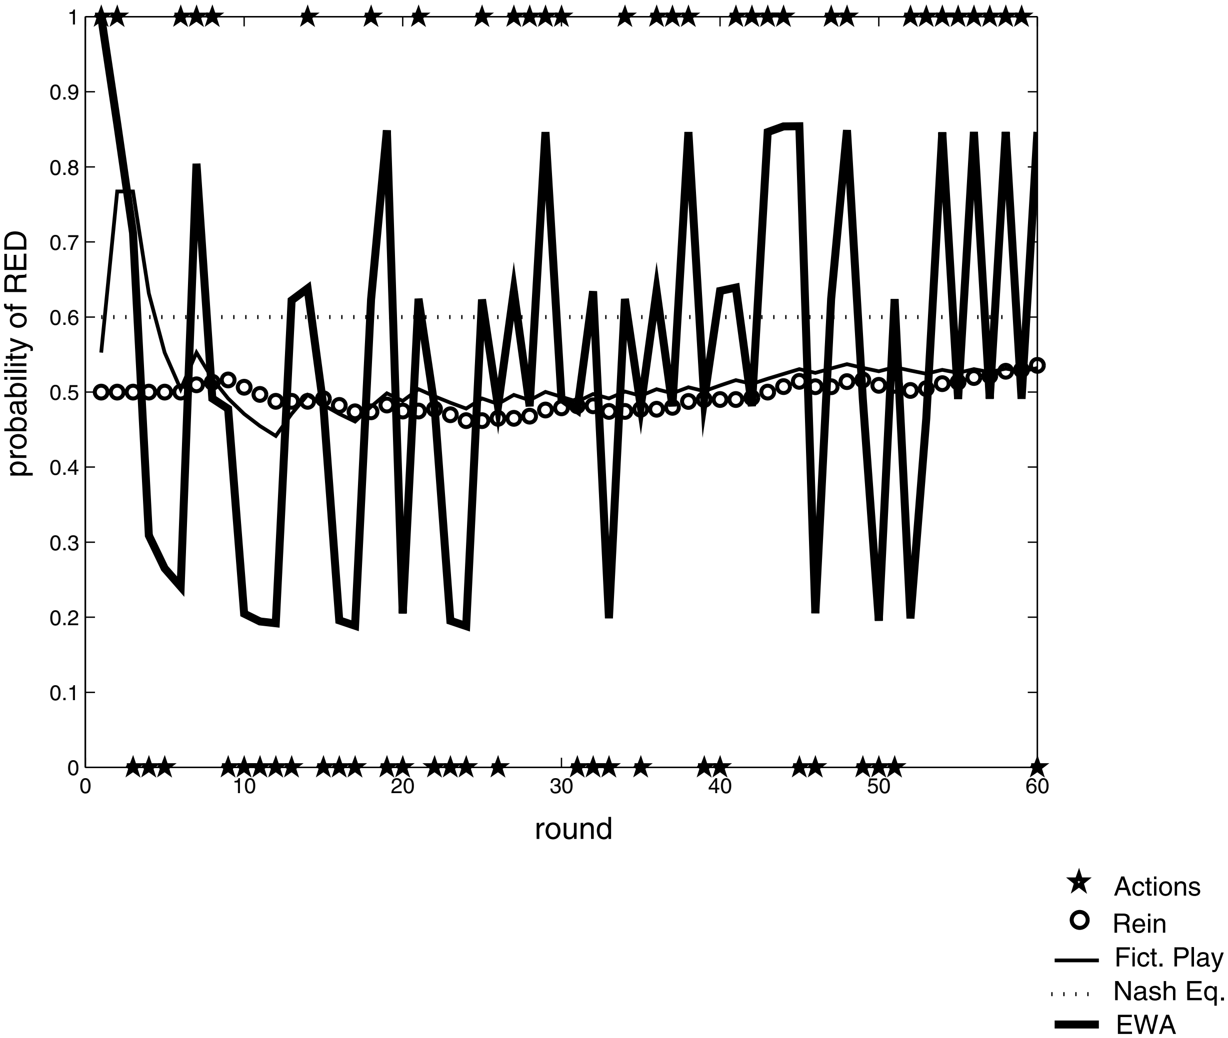

Figure 1p: Model Predictions: Experiment 2, Player 23

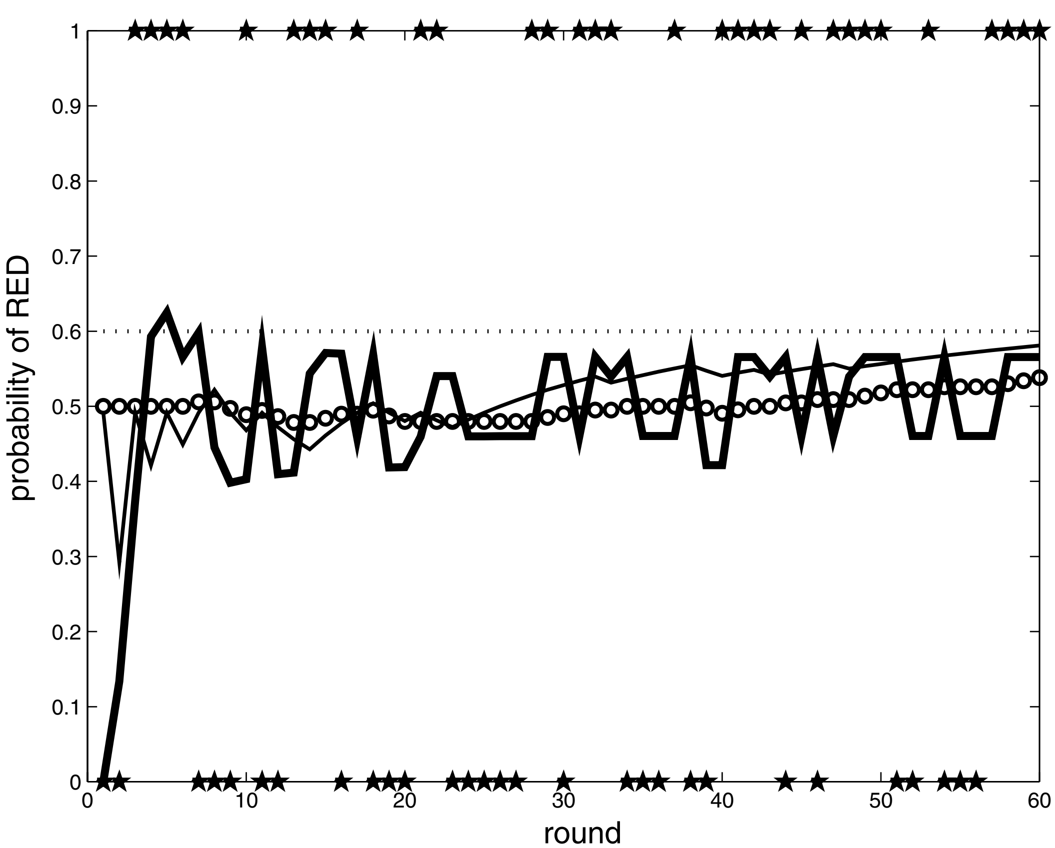

Figure 2a: Level MSD's Experiment 1

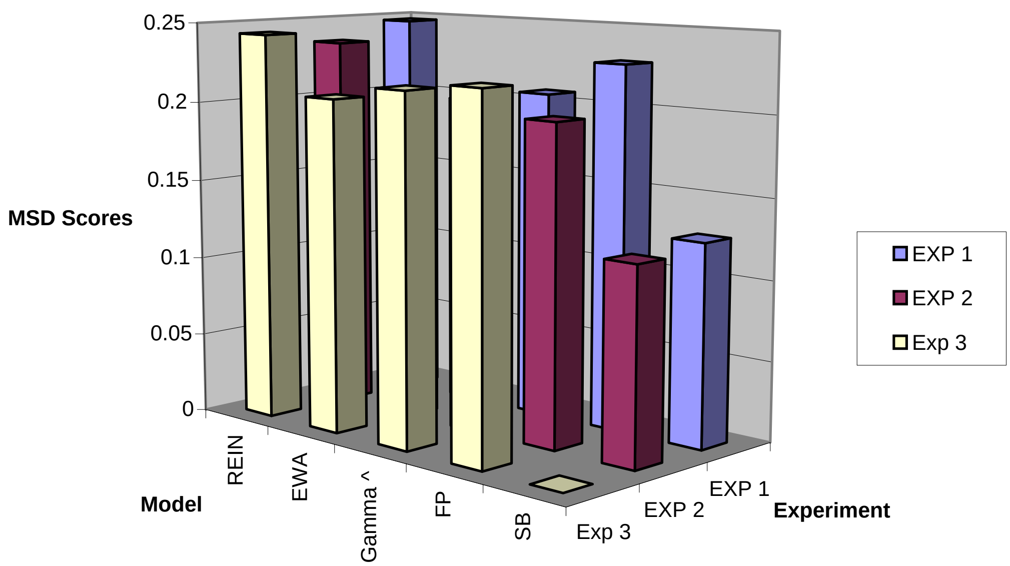

Figure 2b: Change MSD's-- Experiment 1

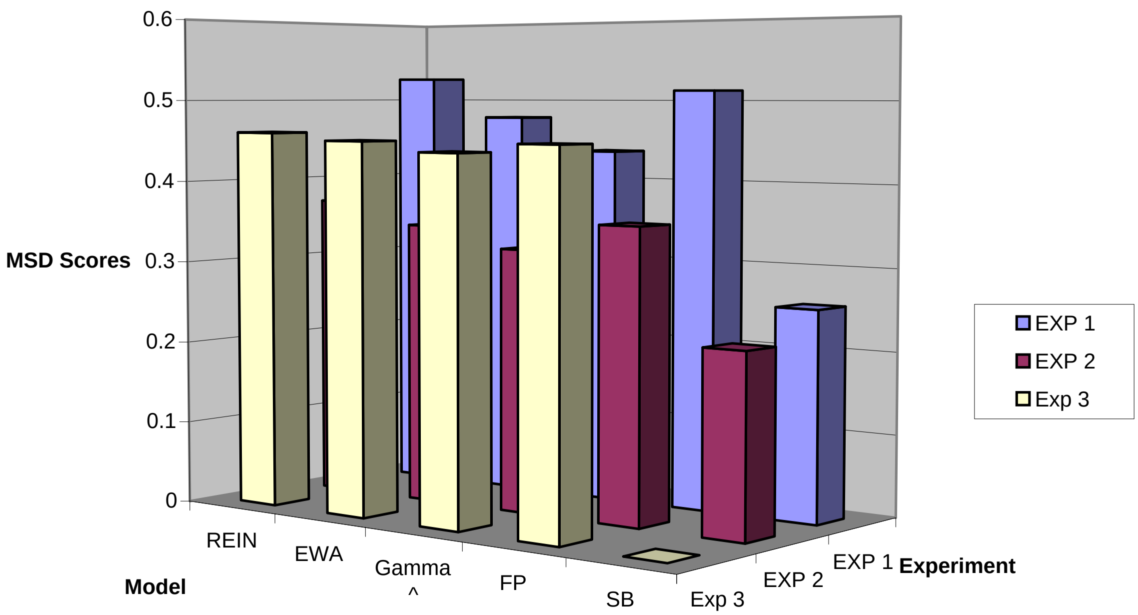

Figure 3a: Experiment 1 Individual Level MSD Scores

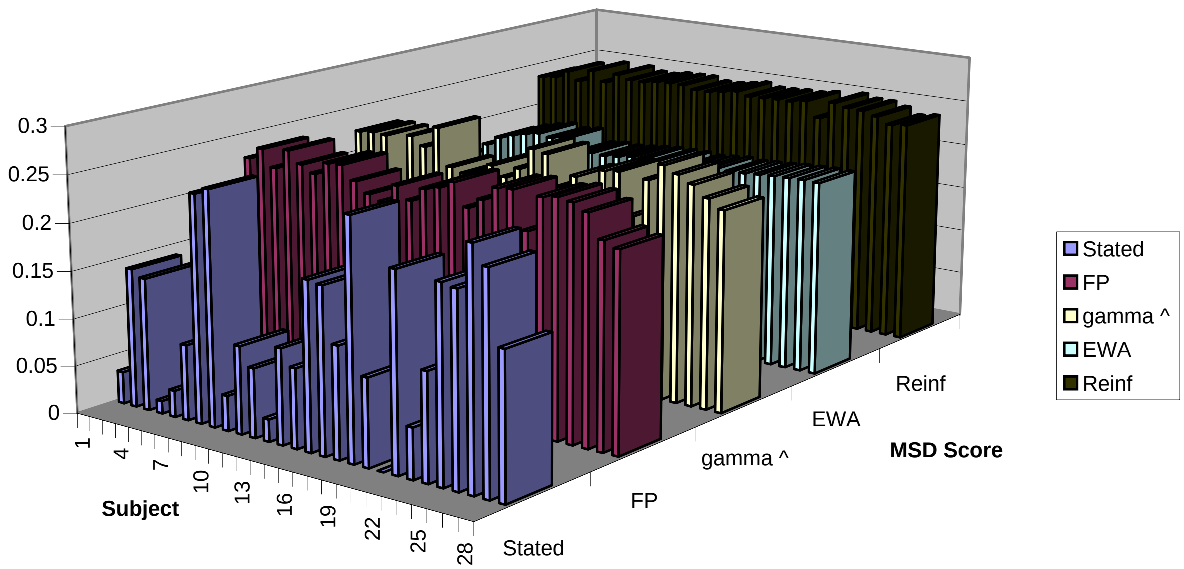

Figure 3b: Experiment 1 Individual Change MSD Scores

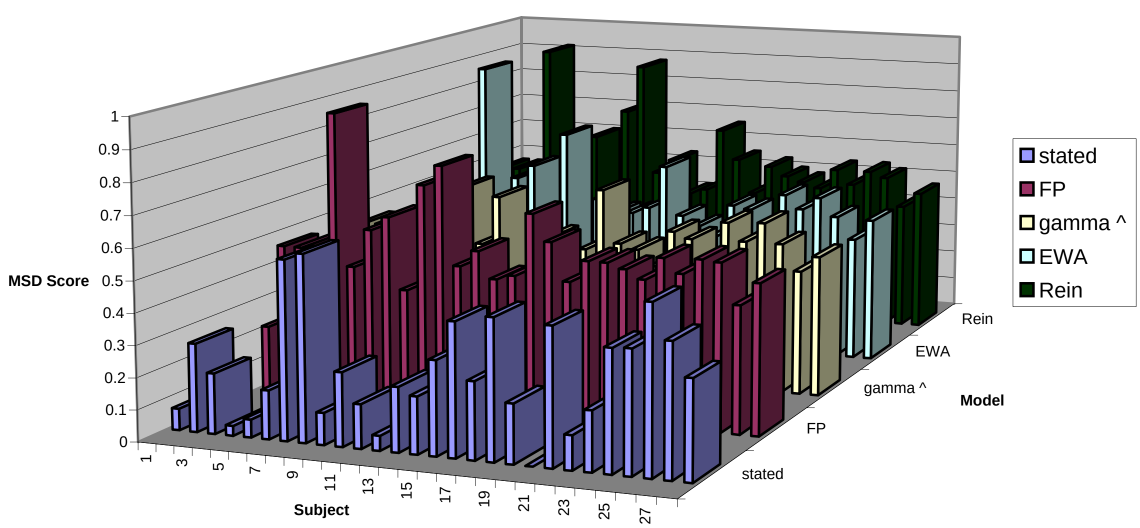

Figure 3c: Experiment 2 Individual Level MSD Scores

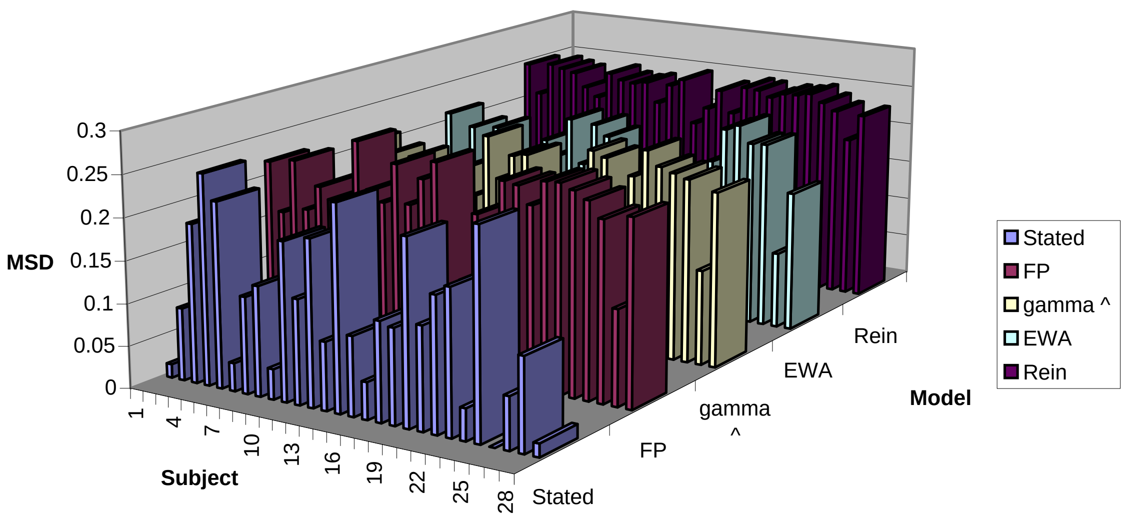

Figure 3d: Experiment 2 Individual Change MSD Scores

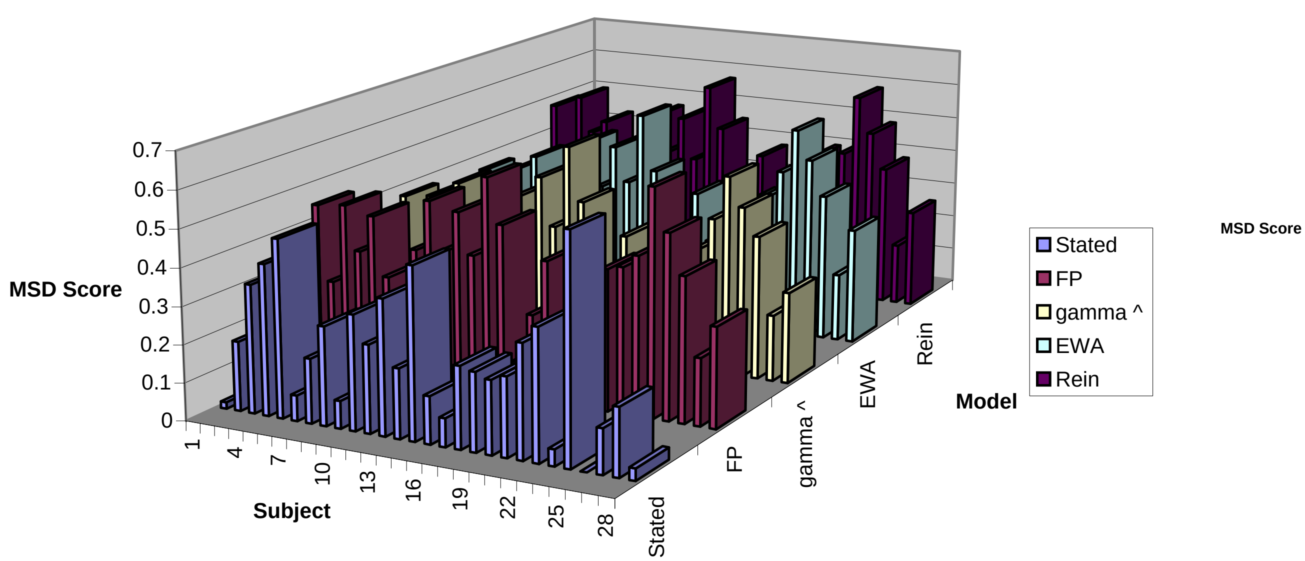

Figure 4a: Experiment 1
Individual Level MSD and Change MSD
Stated Beliefs v. Reinforcement

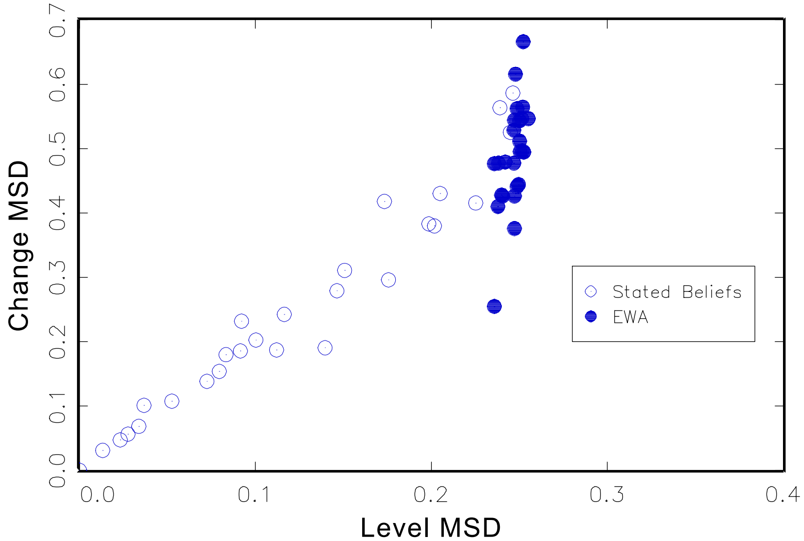

Figure 4b: Experiment 1
Individual Level MSD and Change MSD
Stated Beliefs v. EWA

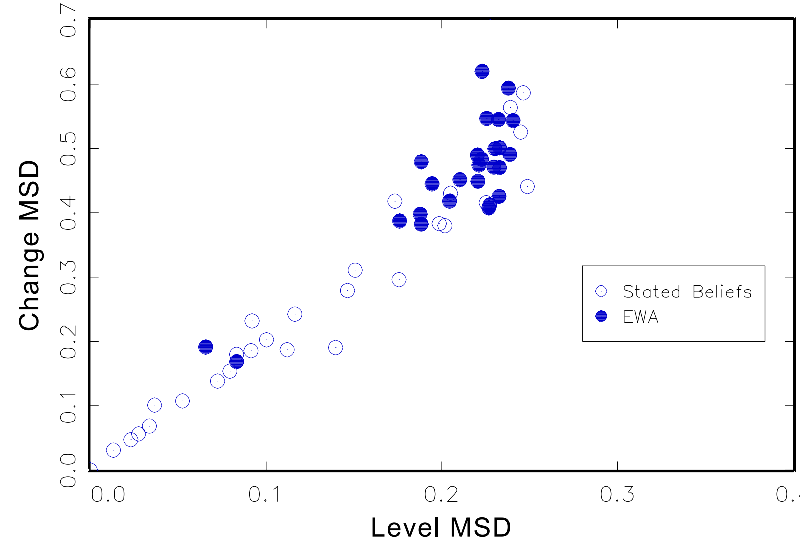

Figure 4c: Experiment 2
Individual Level MSD and Change MSD
Stated Beliefs v. Reinforcement

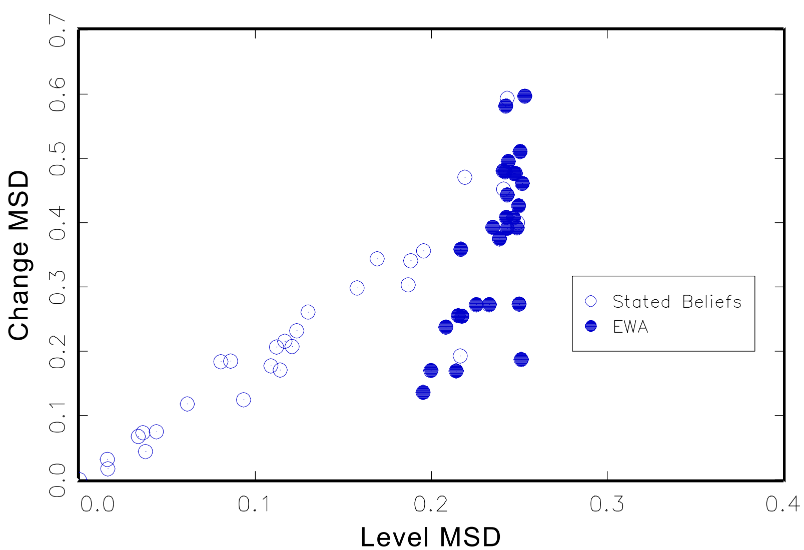

Figure 4d: Experiment 2
Individual Level MSD and Change MSD
Stated Beliefs v. EWA

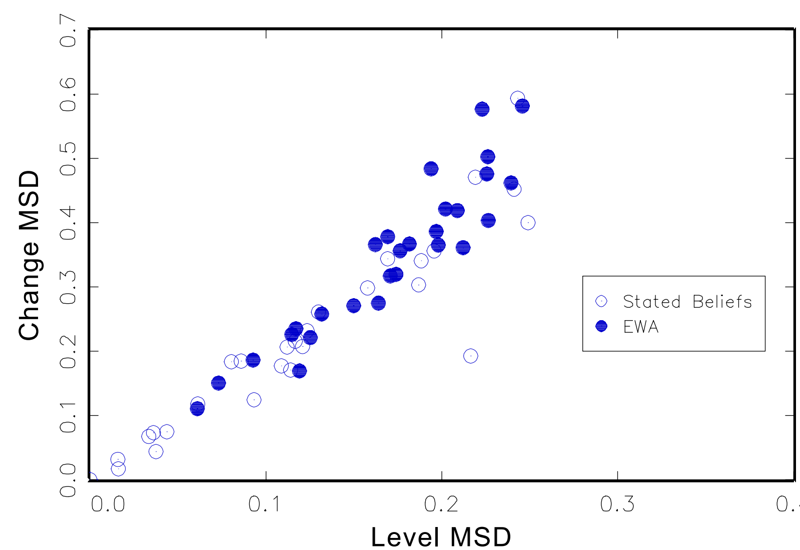

Figure 5a: Experiment 1: Resoluton Scores

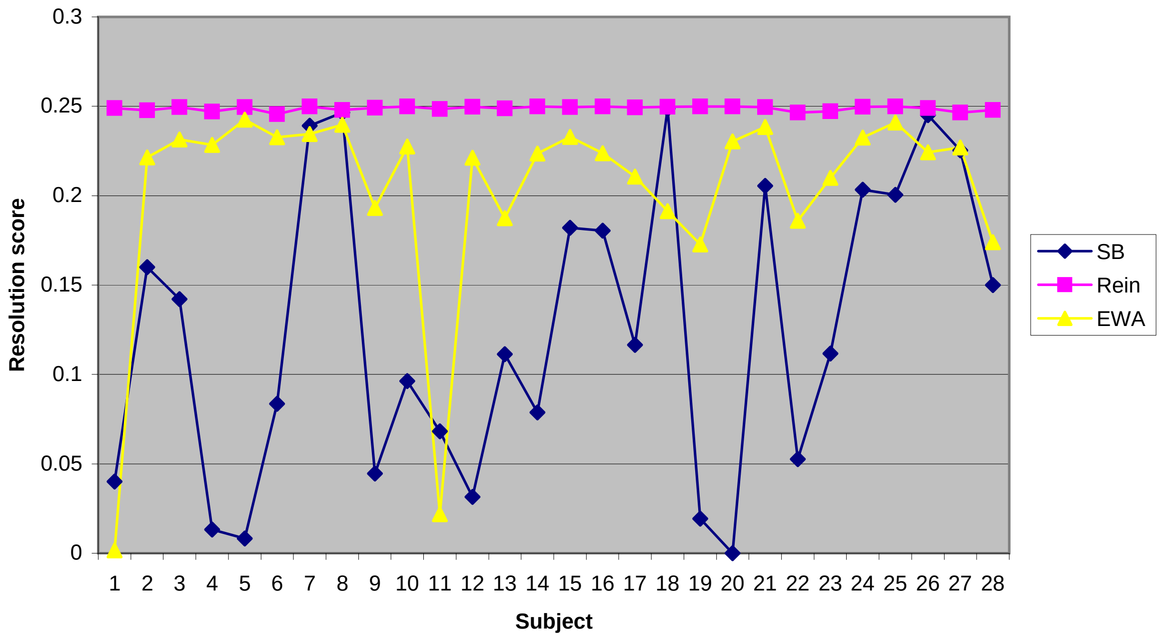

Figure 5b: Experiment 2: Resolution