In D. Reitter & F. E. Ritter (Eds.), Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016). University Park, PA: Penn State.

Explaining inter-individual variability in strategy selection: A cue weight learning approach

Hrvoje Stojic[1] (hrvoje.stojic@upf.edu),
Henrik Olsson[2] (olsson@santafe.edu), **Pantelis P. Analytis**[3] (analytis@mpib-berlin.mpg.de)
[1]Department of Economics and Business, Universitat Pompeu Fabra, [2]Santa Fe Institute
[3]Center for Adaptive Behavior and Cognition (ABC), Max Planck Institute for Human Development

Abstract

Do people integrate all the information at hand when they make choices or do they employ heuristics that ignore some of it? Recent research indicates that people's behavior should and does depend on the statistical properties of the environments within which cognition operates. However, in a single environment there are always decision makers who rely on less effective strategies. The source of this inter-individual variation has not been identified yet. In this article we postulate that it can be largely explained by differences in the speed of learning. We designed an experiment where participants first made choices between three multi-cue alternatives and received feedback about their quality. In a second stage, they predicted the quality of alternatives without receiving feedback. The quality was a linear combination of cue weights and cue values. To employ heuristics the participants had to learn at least weight directions and ranks, while for the integrative strategy they needed to learn the cue weights. We find that participants who showed evidence of learning cue weights rather than the ordering performed well in the estimation task that followed decisions, with cue weight knowledge being strongly related to decision performance. Further, we find that differences in how fast participants learn the cue weights explain the variability in regards to what strategy they adopted within an environment.

Keywords: decision making; heuristics; cue weight learning; function learning; strategy selection.

Introduction

Consider the following problem: you want to decide which hotel to book for your next vacation and you have access to information such as the facilities of the hotel, average reviews, cleanliness etc. To make an educated choice you could weight and add all the information at hand for each alternative and then choose the one that achieved the highest score. This is a weighted additive strategy (WADD; Payne et al., 1993). Alternatively, you could compare the hotels according to the most important cue and choose the one with the largest cue value. If some alternatives are tied on the first cue, you could move to the next cue in the ranking until you reach a decisive cue and stop your search. This corresponds to a heuristic strategy called take-the-best (TTB; Gigerenzer & Goldstein, 1996). On average take-the-best would ignore most of the information, as your decision would often be based on a single cue. Researchers have investigated theoretically the conditions under which it is well-advised to rely on integrative strategies such as WADD or heuristic strategies like TTB (e.g., Hogarth & Karelaia, 2007, 2005; Martignon & Hoffrage, 2002). Empirically, however, there is a large inter-individual heterogeneity and substantial proportion of people still seem to use an inferior strategy (Bröder, 2003; Rieskamp & Otto, 2006; Pachur & Olsson, 2012).

Strategy performance primarily depends on the statistical properties of the relationship between cues and alternative quality. TTB fares well in comparison to WADD when the most informative cues are much more valuable than the less informative ones (Hogarth & Karelaia, 2007), or when the cue inter-correlations are high (Hogarth & Karelaia, 2005). In environments with binary cue values, when the weights of the cues with higher weight rankings are larger or equal to the sum of weights of the cues with lower rankings, TTB cannot be outperformed by WADD. When this property does not hold, a WADD model with well-calibrated weights is expected to outperform TTB. The former environments are called non-compensatory and the latter compensatory (Martignon & Hoffrage, 2002).

Several experiments have demonstrated that over time most people converge to the best performing strategy. For example, people tend to adopt TTB in non-compensatory environments and WADD in compensatory environments (Bröder, 2003; Rieskamp & Otto, 2006). Similarly, in non-linear environments, when none of the aforementioned two strategies performs well, many people employ memory-based exemplar strategies (Pachur & Olsson, 2012). Further, people prefer heuristic strategies over integrative strategies when they are under time pressure or when the cost of learning cue values is high (Rieskamp & Hoffrage, 2008).

Within a single environment, however, there is always a substantial portion of participants that use inferior strategies. For example, in a non-compensatory environment there are always participants that continue using WADD, or TTB in the compensatory environment. The source of this inter-individual variation has not been identified yet, although it is widely reported (e.g., Brehmer, 1994; Einhorn, 1970; Bröder, 2003; Rieskamp & Otto, 2006). Bröder (2012) provides a summary of existing research on inter-individual differences in adoption of TTB and WADD strategies. The only variable that shows some correlation is the intelligence score. TTB users in the non-compensatory environment tend to score higher on an intelligence test than WADD users, although the effect is rather small. None of the personality measures, such as the "Big Five", show a substantial correlation with strategy adoption. Similarly, motivational variables, cognitive styles, working memory capacity, and working memory load do not seem to influence adoption of TTB or WADD. Hence, the variation *within* an environment remains largely unexplained.

In this article we propose a solution to this puzzle. Strategies like TTB and WADD rely on cue weights. While in some experiments participants are given the cue validity weights directly (e.g., Rieskamp & Otto, 2006), in most of them participants have to learn the weights (e.g., Bröder, 2003; Bergert & Nosofsky, 2007). Hence, besides figuring out which strategy to use, they also need to learn the statistical properties that are input to the strategies. Importantly, strategies differ with respect to the amount of knowledge they require about the validity weights. While WADD requires exact quantitative estimates, TTB only requires the ranking and directions. Under reasonable theoretical assumptions, heuristic strategies like TTB are largely insensitive to the gap between estimated and objective validity weights, while performance of WADD is heavily affected (Hogarth & Karelaia, 2007; Katsikopoulos et al., 2010). As a result, in many environments people can leverage WADD's improved performance only after some learning has occurred, and the estimated weights are relatively close to the objective ones. When coupled with usual individual differences in speed of learning, this explanation can address the observed variability in strategy selection. For example, in an environment favoring WADD, this leads to the prediction that slower learners will stick longer to the TTB heuristic, while faster learners will have more precise knowledge about the cue validity weights and will adopt WADD in greater numbers.

Our article suggests a novel approach in the study of decision making strategy by examining decision processes and cue weight learning in tandem. In our experiment, participants complete two tasks, a decision making and an estimation task. By adding an estimation task where participants make predictions about values of alternatives we can model their cue weight learning and infer the evolution of their knowledge about cue weights. Thus, we can identify the role of cue weight learning in strategy selection and test the predictions made above.

Method[1]

Participants

Seventy-eight participants (49 women, 29 men, $M_{age} = 21.8$, age range: 17–54 years), recruited from the Universitat Pompeu Fabra subject pool, took part in the study. They were paid a show-up fee of five euros and a performance dependent bonus of 6.8 euros on average. The experiment lasted 43 minutes on average.

Stimuli and procedure

The experiment consisted of two tasks: the participants first completed a decision making task and then an estimation task. In the decision task they repeatedly faced three alternatives, each described by the same four cues (Figure 1, left). The task was presented as a cheese game. Each alternative represented a cheese, the cues were "Lactic", "Acetic", "Casein" and "Texture", while the alternative values represented enjoyment units (EU).

[1]The raw data is publicly available on Figshare: http://dx.doi.org/10.6084/m9.figshare.1609680.

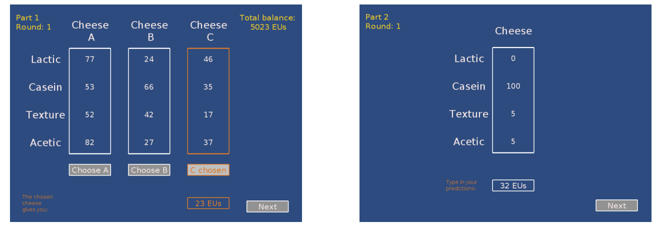

Figure 1: Left: Decision task. Right: Estimation task.

The criterion value, Y, of each alternative was a noisy linear combination of cue values and cue validity weights with weights fixed at 4,-3,2 and -1. These cue validity weights strongly favors WADD over TTB. Cue values were sampled from uniform distribution $U(10,90)$. A normally distributed error term, $e \sim N(0,30)$, was added to each alternative. We created 480 unique alternatives in this manner and allocated them randomly across 160 trials, three alternatives per trial. Cue inter-correlations were zero on average. The stimuli were drawn only once and all participants received the same stimuli. The earnings were determined by the criterion value Y of the chosen alternative, which was also shown as feedback in each trial.

After every 40 trials in the decision task participants answered questions that probed their knowledge about the cue weights. Following Speekenbrink & Shanks (2010), we asked them to rate the strength of the relation between each cue and the value of the cheese on a scale from -10 (highly negative) to 10 (highly positive). Questions for all four cues were shown on the same screen, in the same order that was used to present the stimuli.

In the estimation task participants received a single alternative in each trial and their task was to predict the criterion value (Figure 1, right). No feedback was provided. We incentivized truthful reporting by computing the payoff as a function of a difference between the prediction P and the criterion value, $200 - |P - Y|$.

The stimuli for the estimation task were generated with the same cue validity weights as in the decision task. We generated 20 alternatives for interpolation trials by drawing cue values from the same range as in the decision task, $U(10,90)$, and multiplying them with weights. We generated extrapolation trials in an analogous way by drawing cue values from two intervals at the extreme ends, $U(0,10)$ and $U(90,100)$ that have not been experienced during the decision task. After a single draw was made, trials were randomly ordered and all participants received the same set of stimuli.

In the decision task the participants were informed about the cues and the range of values they could take, and that they could use this information in making their choices. They were not told about the functional relationship between cue values and value of the cheese, nor that the weights differ for different cues. It was stressed that in each trial they would get three new cheeses that differ in their cue values. The estimation task was announced at the beginning in the instructions, but without specifying details.

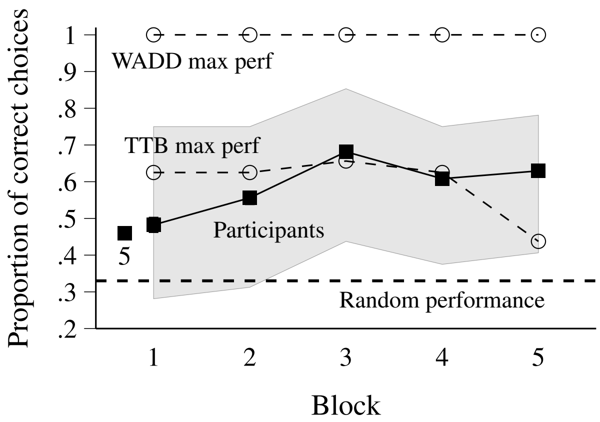

Figure 2: Mean accuracy over trial blocks. Each block result is a mean of individual means across 32 trials.

We told participants that it takes 60 minutes on average to complete the experiment. Each participant was presented with a unique random order of alternatives and cues. The four cue labels were also randomly attached to underlying cues separately for each participant.

Behavioral results

Choices in the decision task

Participants' performance, measured as percentage of correct choices per block, improved over time (Figure 2). Choice accuracy is much higher than the random level of 0.33 already in the first five trials (marked with number five in the figure), with 46% accuracy. People have a strong prior for positive linear relationships (Brehmer, 1994), which matches well the function that we used to construct the stimuli. Participants achieved a mean accuracy of 0.48 in the first block and by the end of the training phase they were close to choosing correctly the alternative with the highest criterion value two out of three times, 0.63. Although mean choice accuracy is similar to the accuracy achieved by TTB with ideal knowledge, 0.59 on average, the variance in individual choice accuracy curves is quite large. The shaded region around the mean performance indicates the range of accuracies, from 10^{th} to 90^{th} percentile. Hence, there are many individuals with accuracies far above what could be achieved with TTB.

Insight questions provide us with a first indication of how well participants have learned the cue validity weights. Previous research using such questions has shown that people have good insight into what they have learned (Speekenbrink & Shanks, 2010). Figure 3 shows mean ratings for all four cues. Participants got the relative ordering and directions right on average already after 40 trials and it got clearer as the training progressed. They learned that the second cue has a larger weight (although negative) than the third cue only at the end, and failed to detect that the fourth cue had a small negative weight. This is not surprising as negative linear relationships are more difficult to learn than positive linear ones (Brehmer, 1994). Although insight questions use an arbitrary scale and it is difficult to identify exact cue weights that participants have acquired, they do suggest that people learn more than ordering and directions. This is supported by changes in ratings over the course of the decision task, even though the ordering and directions were mostly established already after first time participants answered the insight questions.

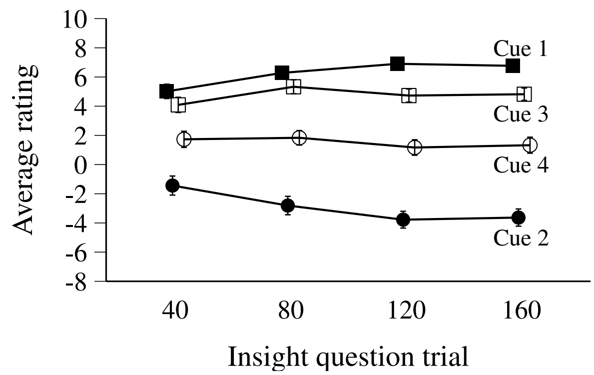

Figure 3: Average insight ratings across trials. Error bars represent standard errors of means across participants.

Predictions in the estimation task

We can also assess knowledge about cue validity weights by examining the performance in the estimation task. We computed mean absolute deviation (MAD) and correlation between participants' predictions and criterion values as a measure of performance. Mean MAD across participants is 120 (SD = 30.8), which means that on average predictions were 120 EU's away from criterion values. Mean (median) Spearman correlation is 0.63 (0.70; SD = 0.24). The participants are doing a good job in predicting criterion values of test items, but as expected, inter-individual variation in learning is substantial, with MAD ranging from 51 to 189. While most people are doing quite well, having very high correlations and low MAD's, some people do very poorly.

How would a decision maker that only learned the ranking of cues fare in the estimation task? Such a decision maker could take a mean of the criterion values experienced in the decision task and use it as a fixed prediction for all items in the estimation task. This is our baseline prediction performance. The MAD between baseline predictions and criterion values was 172, much larger than for observed MAD.

We get more complete insight by examining mean predictions across participants for each of the 40 items in the estimation task. Figure 4 shows that in the range of item values from about zero to 200, mean predictions correspond very closely to the criterion values. More deviations occur for more extreme values, with somewhat poorer predictions for extrapolation items than interpolation items. Importantly, predictions correspond much better to criterion values than baseline predictions. Thus, most participants do acquire more precise knowledge about cue validity weights, rather than only the ordering and directions.

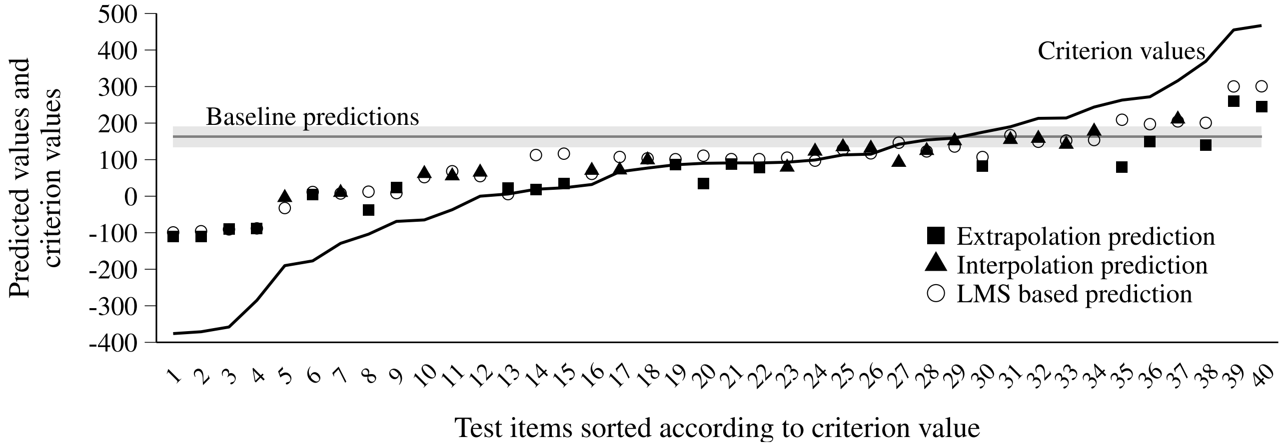

Figure 4: Mean prediction for each item in the estimation task. Black line that diagonally goes from lower left to upper right corner represents the criterion value of the items, while the gray horizontal line is the baseline prediction – mean value of the items experienced in the decision task.

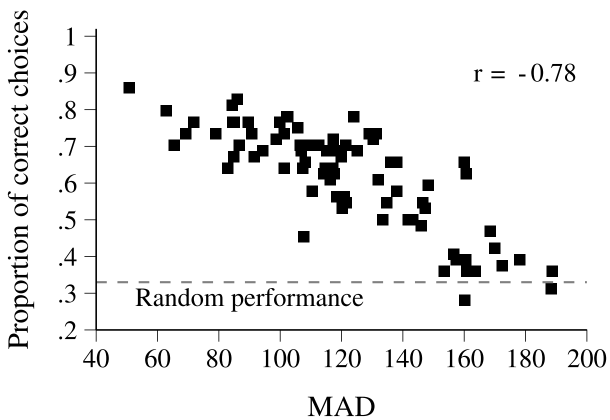

Figure 5: Relation between performance in the estimation task (mean absolute deviation (MAD) between predictions and criterion values) and the decision task (proportion of correct choices in the last two blocks).

We examine the relationship between individual performances in the two tasks to obtain model-free evidence that cue weight learning plays an important role in strategy selection. We find a strong relationship between choice accuracy in the decision task and MAD in the estimation task, as indicated by a Spearman correlation of -0.78 (Figure 5). This suggests that participants with good prediction performance know the cue weights well, which allowed them to employ WADD and achieve good decision performance. Surprisingly, many participants who had poor prediction performance also had decision performance far below 0.59 which is possible to achieve with very little knowledge for TTB in this environment. They either relied on WADD in spite of their poor knowledge or those participants simply paid less attention and performed close to random in both tasks.

Modeling

Next we turn to identifying the strategies used by each participant in the decision task. We first describe the cue weight learning model that will produce trial-by-trial predictions of participants' knowledge of cue weights. These weights will in turn be used in fitting TTB and WADD models to participants choice data. Finally, we will examine whether participants that were best fitted by TTB have less developed knowledge of cue weights than those best fitted by WADD, as predicted.

Modeling the cue weight learning

We used a least mean squares model to model the cue weight learning process (Gluck & Bower, 1988). The LMS model predicts the criterion value of an alternative on trial t as

$$P_t = \sum_{i=1}^{4} x_{i,t} u_{i,t},$$

where $u_{i,t}$ are cue utilization weights and $x_{i,t}$ are cue values of cue i in each trial t. Utilization weights are updated in every trial through the delta rule, based on a prediction error defined as the difference between the predicted criterion value, P_t and the true criterion value, Y_t, that a participant receives as a feedback in the decision task

$$u_{i,t+1} = u_{i,t} + \frac{\eta}{t^{\gamma}}(Y_t - P_t)x_{i,t},$$

where $0 \leq \eta \leq 1$ is a learning rate parameter shared by all four cues and $\gamma \geq 0$ is a decay parameter. We initialized the weights to $u_{i0} = 0, i = 1, \ldots, N$. Note that the cue weight learning process is based only on the alternative for which participants receive feedback, the rest is ignored by the LMS model.

We fitted two different versions of LMS model. LMS_d where both η and γ are free parameters and LMS where γ is set to 0. Parameters were initialized at the beginning of the decision task and in each trial cue values and criterion of the chosen alternative were used to update the weights. The weights from the last trial were used to make model based predictions in the estimation task. To estimate the model parameters we minimized the mean squared error between the participant's and model's predictions. The LMS model was fitted separately from the choice models.

Modeling the choices

Random Choice Model We used a random choice model (RCM) as a baseline. RCM predicts the same probability, .33, for each alternative.

WADD Model Our version of WADD linearly combines the cue utilization weights learned by the LMS model with cue values to produce predicted value of each alternative k in trial t

$$R_t^k = \sum_{i=1}^{4} x_{i,t}^k u_{i,t},$$

where $u_{i,t}$ are cue utilization weights learned by the LMS model based on trials $1 : t-1$. WADD then deterministically decides by maximizing among the alternatives. To fit WADD to data we assume an additional "tremble" error. If a strategy produces a probability that alternative k is chosen, $P(C = k)$, then the probability of choosing k after taking into account the tremble error, ε, is given by

$$P(C_t = k; \varepsilon) = (1 - \varepsilon) \times P(C_t = k) + \frac{\varepsilon}{3}$$

TTB Model Our version of TTB uses the cue weight information from the LMS model, $u_{i,t}$, to order the absolute value of the weights from the largest weight to the lowest, producing a ranking r_t. The ranking is done on absolute values because a strong negative weight is as predictive as a strong positive weight. TTB then chooses an alternative with the largest cue value of the most predictive cue according to ranking r_t. If values of the first cue according to the ranking are the same for all alternatives[2], TTB inspects the second cue and so on, until it finds a cue that discriminates between the alternatives. If no cue discriminates, a choice is made at random. If the deciding cue had a negative weight according to the u_t, cue values of all three alternatives were multiplied with -1, to maintain the correctness of the rule of choosing the alternative with larger cue value. Same as in the WADD model, we add a "tremble" error term to arrive at the final choice probability, $P(C_t = k; \varepsilon)$.

Modeling results

Table 1 shows the mean Bayesian Information Criterion (BIC) score across participants for LMS models and choice models. Both *LMS* and LMS_d fit the predictions equally well, both in terms of mean BIC (368 and 366) and number of participants best fitted (39 for both). However, LMS_d fits results better in a qualitative sense. It emulates better the insight questions results where most people acquire ordering and directions very fast. Hence, we used weights from LMS_d in the choice models. Moreover, LMS_d based predictions for estimation task items correspond closely to participants' predictions (Figure 4).

[2]Ties are rare in environments with continuous cue values, making this version of TTB quasi-equivalent to a single-variable strategy, which uses only the most important cue.

Table 1: Mean Bayesian Information Criterion (BIC) scores of models (standard deviation in the parenthesis), number of participants best fitted the model and mean parameter values.

Model	#	BIC	N	η	γ	ε
LMS	1	368 (25)	39	2e-5	-	-
LMS_d	2	366 (24)	39	2e-4	.61	-
WADD	1	283 (40)	56	-	-	.62
TTB	1	302 (40)	11	-	-	.74
RCM	0	355 (2)	11	-	-	-

Note. # = Number of parameters in the model; N = number of participants best fitted by the model; η = learning rate in LMS; γ = decay rate in LMS; ε = tremble error.

In terms of choice models, as expected, WADD has better mean BIC score (283) than TTB (302). Similarly, most participants were best fitted by WADD (56), followed by TTB (11) and RCM (11). As has been widely observed in previous studies, although it pays better to adopt WADD, and indeed most people do so, there is substantial inter-individual variability.There are substantial differences between the three groups. As expected, WADD users reached the highest accuracy, they were choosing the best alternative on average in 0.63 proportion of trials. TTB users performed worse, having a choice accuracy of 0.55. Although RCM users were the worst, reaching mean accuracy of 0.42, their performance is somewhat higher than the random level and they do exhibit some learning by the end of the training phase.

Next we examine our prediction that participants best fitted with TTB are those that learn slower and did not manage to arrive at sufficiently good utilization weights to switch to WADD. We plot the evolution of utilization weights estimated with the LMS_d model, separately for participants best fitted with each model (Figure 6). We see that WADD users have a well developed knowledge of all four cues, while TTB users have less developed knowledge. Notably, TTB users have very good estimates for the most important cue and do not distinguish that well between the other three cues. Their adoption of the TTB strategy is well justified by their subjective knowledge of the cue weights. RCM users' knowledge is very poor, capturing unmotivated or inattentive participants.

We can also examine estimated learning rate parameters of the LMS_d model. Learning rates are higher for WADD users than TTB users, and lowest for RCM users (Figure 6). Median learning rate for WADD users was 0.00015, while for TTB users it was lower for an order of magnitude, 0.000027. Median decay rates are correspondingly higher for the WADD users, 0.69, than for the TTB users, 0.54. Performance of TTB users in the estimation task ($M_{MAD} = 133$) was expectedly worse than that of WADD users ($M_{MAD} = 112$), but importantly, substantially better than of RCM users ($M_{MAD} = 155$) or baseline ($M_{MAD} = 172$). Similar differences can be seen in the insight questions results, with knowledge of TTB users evolving over time. This suggest that even a TTB user learns more than just the ordering and the direction of cues.

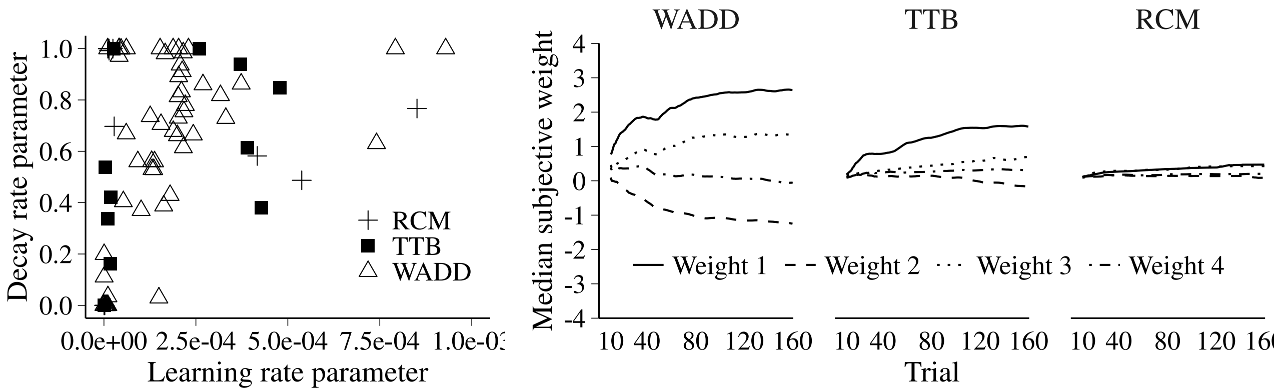

Figure 6: Left: Estimated learning η and decay rate γ parameters for the LMS_d model. Right: Smoothed median cue validity weights estimated with the LMS model for participants best fitted with WADD, TTB, and RCM.

Finally, we conducted a logistic regression with mean absolute difference between LMS obtained utilization weights in the last block and objective weight as a predictor of strategy use. We obtained a negative coefficient, as predicted, at a value of -1.466, with $p = 0.0139$ and $CI[-2.743, -0.367]$ (WADD users were coded as 1 and TTB users as 0, while RCM users were not included). In odds ratio terms, for one unit increase in mean difference, the odds of using WADD decrease by 76%. Odds of using WADD for the perfect knowledge (zero difference) is very high, 50.85, which amounts to a probability of 0.975. Although this outcome was already suggested by behavioral results illustrated in Figure 5, this analysis establishes the link between the knowledge of cue weights and strategy selection more clearly, in a model based manner. Since WADD users achieve greater decision performance, it explains the large correlation between estimation and decision performance seen in Figure 5.

Discussion & Conclusion

In our experiment participants differed in how fast they acquired knowledge of cue weights, and we predicted this heterogeneity to be responsible for the variability in strategy selection. Our results showed support for our predictions – WADD users had better developed knowledge of cue weights than TTB users and the performance in the estimation task is consistent with the strategy adoption. Our learning rate account suggests that, given time, TTB users would learn the weights sufficiently well and switch to the better performing WADD strategy.

Where do inter-individual differences in learning rates come from in the first place? These differences might be akin to traits like intelligence or personality factors investigated by Bröder (2012). This would require the learning rates to be stable across time and tasks within people. To our knowledge, there is no study that examines the stability of learning rates and is difficult to generalize beyond our task.

In our study we set out to test a specific hypothesis and to inform the debate on whether people are better described by the WADD or TTB model. We have to note that the models do not perform particularly well in our task. This can be witnessed in the high values of the ε parameter in Table 1, meaning that models on average predict the choices of the participants half of the time. Given our modest goals we did not try to look for models that would explain behavior even better. Our results, however, indicate that we should look for such models within the probabilistic rather than deterministic class of models (Bergert & Nosofsky, 2007).

Our results could be also explained if some participants first adopted TTB and as a consequence learned cue weights differently. With our current experimental design we cannot, unfortunately, determine the direction of the causal arrow. However, our evidence indicates that TTB users acquire more than ordinal information about cue weights and that this knowledge becomes more precise over time. This suggests that, if such interdependence exists, at most it slows down the learning. This evidence comes from three sources – the insight questions, the estimation task and the joint modeling of cue weight learning and decision making. The continuous evolution of our participants' knowledge of cue weights goes against the frugality and robustness justifications of TTB. The argument against using cue weights hinges on their vulnerability to overfitting – relying on ordinal information instead leads to better generalization. From our perspective, TTB and other heuristic strategies are used either due to cognitive limitations or when the structure of the environment is known better and these strategies are the rational thing to do (also see Davis-Stober et al., 2010; Davis-Stober, 2011).

In this decision-making task, our evidence suggests that learning the properties of the environment is predominant, and strategy selection is influenced by it. Different decision making tasks, however, may lead to distinct linkages between cue weight learning and decision making processes. Exploring the nature of these interactions opens an exciting direction for future research (see Stojic et al., 2015, 2016).

Acknowledgments

Hrvoje Stojic has been supported by FPU grant awarded by Ministry of Education, Culture and Sport of Spain (FPU12/05859) and by the Barcelona Graduate School of Economics.

References

Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 107–129.

Brehmer, B. (1994). The psychology of linear judgement models. *Acta Psychologica*, *87*, 137–154.

Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 611–625.

Bröder, A. (2012). The quest for take the best - Insights and outlooks from experimental research. In P. M. Todd, G. Gigerenzer, & the ABC Research Group (Eds.), *Ecological rationality: Intelligence in the world* (pp. 216–240). New York, NY, US: Oxford University Press.

Davis-Stober, C. P. (2011). A geometric analysis of when fixed weighting schemes will outperform ordinary least squares. *Psychometrika*, *76*(4), 650–669.

Davis-Stober, C. P., Dana, J., & Budescu, D. V. (2010). Why recognition is rational: Optimality results on single-variable decision rules. *Judgment and Decision Making*, *5*(4), 216.

Einhorn, H. J. (1970). The use of nonlinear, noncompensatory models in decision making. *Psychological Bulletin*, *73*, 221–230.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227–247.

Hogarth, R. M., & Karelaia, N. (2005). Ignoring information in binary choice with continuous variables: When is less “more”? *Journal of Mathematical Psychology*, *49*(2), 115–124.

Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: matching rules and environments. *Psychological Review*, *114*, 733–758.

Katsikopoulos, K. V., Schooler, L. J., & Hertwig, R. (2010). The robust beauty of ordinary information. *Psychological Review*, *117*, 1259–1266.

Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision*, *52*, 29–71.

Pachur, T., & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. *Cognitive Psychology*, *65*, 207–240.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker.* New York, NY, US: Cambridge University Press.

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: how opportunity costs affect strategy selection. *Acta Psychologica*, *127*, 258–76.

Rieskamp, J., & Otto, P. E. (2006). SSL: a theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.

Speekenbrink, M., & Shanks, D. R. (2010). Learning in a changing environment. *Journal of Experimental Psychology: General*, *139*, 266–298.

Stojic, H., Analytis, P. P., & Speekenbrink, M. (2015). Human behavior in contextual multi-armed bandit problems. In *Proceedings of the 37th annual conference of the cognitive science society.*

Stojic, H., Olsson, H., & Speekenbrink, M. (2016). Not everything looks like a nail: Learning to select appropriate decision strategies in multiple environments. *Submitted for publication*.