Title
How people use the past as cues to the present

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors
Luthra, Mahi K
Todd, Peter M.

Publication Date
2022

Peer reviewed

How people use the past as cues to the present

Mahi Luthra (mkluthra@iu.edu)
Cognitive Science Program, Indiana University Bloomington (IUB);
Department of Psychological and Brain Sciences, IUB,
1001 E 10th St. Bloomington, IN 47405 USA

Peter M. Todd (pmtodd@indiana.edu)
Cognitive Science Program, Indiana University Bloomington (IUB);
Department of Psychological and Brain Sciences, IUB,
1001 E 10th St. Bloomington, IN 47405 USA

Abstract

Humans must often make decisions in temporally autoregressive environments (e.g., weather, stock market). Here, current states of the environment regress on their previous states (either across consecutive timesteps or from several timesteps back in a patterned fashion). The current work investigates people's abilities to utilize previous states of autoregressive sequences as cues to its current state. In Experiment 1 we determine whether utilization of autoregressions reduces as the temporal distance of the predictive timestep increases; and in Experiment 2 we explore whether participants' utilization of previous timesteps in predictions compete such that they reduce utilization of one timestep when increasing utilization of another timestep. We also fit data from both experiments with a trial-by-trial decision model. Overall, we find that participants significantly reduced utilization of a cue with its increased temporal distance. However, we obtained less conclusive results on competition among timestep cues. These results can explain people's predictions in sequential decision tasks (e.g., their tendencies to perceive clumpiness in random environments).

Keywords: decision making; autoregression; cue competition; recency

Introduction

Many real world systems are autoregressive—that is, past states of a system often provide important cues about its current state. For instance, we can predict aspects of the weather at a location (e.g., temperature) quite accurately using the weather report from the past week or even from a year ago. The same can be said about stock markets and clothing trends—such systems regress across consecutive timesteps and/or a certain number of timesteps back in a patterned fashion. In fact, analysts frequently use autoregressive functions to make predictions about the states of these systems (Salisu et al., 2022).

There is considerable work to suggest that humans both assume that systems exhibit autoregressive characteristics and are good at identifying and utilizing existing autoregressions to make predictions. For instance, Luthra and Todd (2021b) found that when given a two-alternative forced choice (2AFC) task, participants used a default recency strategy—on each new trial, they chose alternatives that were correct in recent trials, hence assuming autoregression across consecutive trials. Further, participants altered the weight they gave to recent trials depending on what was optimal for the condition they were given—on conditions where previous trials were less predictive of future ones, participants appropriately reduced their weighting of those trials when making each new prediction. Other research also suggests that participants perform well when the current state regresses on states from a certain number of timesteps back (e.g., three timesteps ago)—this is clear from people's ability to quickly identify patterns and has been demonstrated widely across the statistical learning literature (Saffran et al., 1996).

A default recency strategy, described above, is effective in dynamic autoregressive structures frequently observed in the real world—here only a small number of previous states will be predictive of the current state. However, research finds that people are usually more recency driven than what is optimal in their environment (da Silva et al., 2017; Luthra & Todd, 2021b). For example, Luthra and Todd (2021b) found that participants facing three environmental structures with varying degrees of optimal recency always responded to the task with slightly greater recency (i.e., using fewer previous timesteps) than what was optimal for their environmental structure. This use of greater-than-optimal recency was especially peculiar because participants were demonstrably capable of using the optimal recency values of at least some of the environments. In the current study we investigate this behavior more closely, exploring possible explanations for it.

We believe that greater-than-optimal recency can be explained by a declining ability to perceive and utilize autoregressions with older timesteps, tendencies to reduce use of one timestep cue when another cue is used more (cue competition), or a combination of both. We investigate these explanations in a task where outcomes are autoregressed on one or more previous timesteps. We will use the term *validity* to refer to the actual regression coefficients of previous outcomes on the current outcome and the term *utilization* to refer to regression coefficients of previous outcomes on people's predictions of the current outcome. Hence, validity is the actual autoregression of task outcomes and utilization is the subjective autoregression used by participants. These terms originate from Brunswik (1955) lens theory and have been frequently used in the function and multiple cue learning literature (e.g., Speekenbrink & Shanks, 2010)

The first phenomenon we are interested in studying is the declining ability to identify regression to previous outcomes with their increasing temporal distance from the current prediction. This could produce patterns of declining utilization (i.e., subjective regression weight) of older outcomes. For instance, if the outcome at every timestep t regresses on

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

timesteps $t-1$ and $t-2$ equally, participants might display greater utilization of timestep $t-1$, due to reduced ability to identify regression to timestep $t-2$. This declining ability to identify regressions to older timesteps might be why participants frequently use recency as a decision strategy even in environments where it is not optimal. For instance, in 2AFC tasks where the probabilities of the presented alternatives are unequal and stable (not autoregressed) across trials, the optimal strategy is to utilize all previous trials equally, thereby always choosing the more probable outcome (probability maximizing). However, even in such environments, participants utilize recent trials with greater weighting, choosing alternatives that were successful in a recent window of trials (Luthra & Todd, 2019). We investigate declining utilization with temporal distance of timesteps in Experiment 1—across three conditions, outcomes on each timestep regress on the outcomes one, two, or three timesteps back; we analyze whether utilization reduces as the temporal distance from the current timestep increases.

The second phenomena of interest is cue competition among previous timesteps—that is, reducing the utilization of one timestep when the utilization of another timestep is increased. Cue competition was initially reported in the multiple cue learning literature where participants must learn to predict one variable from multiple other variables (Kruschke & Johansen, 1999). Here, researchers often find that the increased utilization of one cue detracts from the utilization of other cues (for instance, arising due to attention limitations, Kruschke & Johansen, 1999). In more extreme versions of cue competition, participants have been found to make decisions based on a single best cue, ignoring all others (e.g., take-the-best heuristic; Gigerenzer & Goldstein, 1999). In the current study, we determine whether such cue competition effects extend to temporal cues (timesteps in a sequence). There is some evidence in support of this—-for instance, a win-stay-lose-shift strategy (making decisions based on a single previous outcome, even if a larger number of them are predictive; Worthy & Maddox, 2014) could result from temporal cue competition. Further, studies frequently find that people tend to perceive nonexistent patterns in random temporal sequences (Hyman & Jenkin, 1956)—this possibly occurs because people give undue importance to few temporal cues (leading to exaggeration of patterns), instead of weighing the larger sequence equally (which will enable accurate perception of randomness). We investigate cue competition in Experiment 2—here each outcome regresses on two previous outcomes, from one and two timesteps back. By varying the validity of these two previous timesteps across three conditions, we observe how participant utilization of one timestep changes with the increased validity of another.

To study this, we used a task with three discrete mutually exclusive outcomes (a rabbit appears from one of three holes). However, unlike a typical multi-alternative forced choice task, participants in our task did not predict only the actual outcome; rather they expressed their perceived probabilities of the three outcomes (by placing a dog distanced from the three holes in accordance to the predicted underlying probabilities). Predictions were made in this format (i.e, as probabilities instead of outcomes) to help us analyze how participants combined multiple timestep cues to make predictions. As mentioned earlier, in Experiment 2, the outcome on each trial t regressed on two previous trials, $t-1$ and $t-2$. If, for instance, the validity of $t-1$ was 0.4 and the validity of $t-2$ was 0.2, under the forced choice task, optimal responding would entail always choosing the outcome that appeared on trial $t-1$ and completely ignoring trial $t-2$—that is, probability maximizing. If participants used this strategy, it would impede us from investigating their abilities to combine multiple cues. Hence, we created a novel task where incentivized optimal responding would entail accurately expressing the underlying probabilities, as we describe further in Methods below.

Methods

Participants

We had 159 and 175 participants in Experiments 1 and 2 respectively, recruited on MTurk. In both experiments participants completed only one of three conditions.

Task

Participants played our Catch-the-Rabbit game, repeatedly placing their dog somewhere in a triangular carrot-patch to prevent a rabbit from stealing their carrots (see Figure 1). On each trial (of 150 total trials), a rabbit appeared on the computer screen from one of three equally spaced holes. Prior to the appearance of the rabbit, participants placed their dog in between the three holes to chase away the rabbit. On each trial, the rabbit could steal a maximum of 50 carrots and the closer the participants placed their dog to the hole at which the rabbit appears, the more carrots they would save on that trial. However, the number of carrots saved did not increase linearly with closeness of the dog to the correct hole—if that reward mapping were used, probability maximizing would be optimal. Rather, the rewards increased as a log function of closeness to the correct hole. This reward structure ensured that the participants saved the most carrots in the long term if they placed the dog in accordance to the underlying probabilities of the rabbit next appearing at the three holes. The cross-marks for dog placement allow participants to express their predicted probabilities for the rabbit's next location at 0.1 intervals. For all cross-marked positions, the total probabilities of the three holes added to 1. For instance, the position of the dog in Figure 1 would be optimal if the rabbit had a 0.7 probability of appearing from Hole A, a 0.2 probability of appearing from Hole B, and a 0.1 probability of appearing from Hole C. See Appendix for more detail on reward structure.

Participants were informed that they would earn the highest reward (i.e., save the most carrots) overall if they correctly matched the underlying trial-by-trial likelihoods of rabbit positions with their dog placements. Participants played demo

Figure 1: Snapshot of the Catch-the-Rabbit game. Participants place the dog on one of the cross-marked positions in accordance to their judged probabilities of the rabbit's appearance from the three holes. Following placement, the rabbit appears and participants are informed about the number of carrots saved in the trial.

trials to help them understand the reward structure better and match the underlying probabilities.

We ran an initial experiment (Experiment 0; 49 participants) to test whether our novel task worked to get participants to predict the underlying probabilities through their placement of the dog. The underlying probabilities of the three holes was fixed at 0.6, 0.3, and 0.1 (randomly assigned across Holes A, B, and C) throughout the game. Participants successfully learned and expressed these underlying probabilities early in the game—on average they placed the dog at 0.57, 0.28, and 0.15 closeness to the three holes (Figure 2). These results assured us that our task can successfully enable participants to express their probabilities of outcomes.

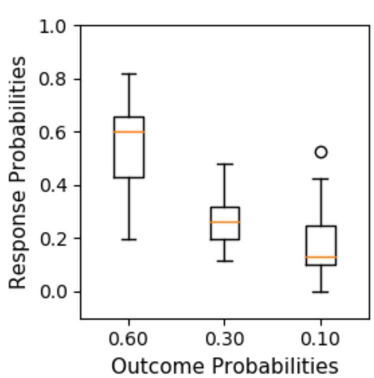

Figure 2: Distribution of participant's responses in Experiment 0. Participants appear to match the outcome probabilities.

Autoregressive Sequences

The sequences were created using the following regression equation:

$$P(outcome_t) = \sum_{n=1}^{3} v_n \cdot outcome_{t-n} + c \tag{1}$$

where $P(outcome_t)$ is a vector of the probabilities of the three outcomes at timestep t, $outcome_{t-n}$ is a vector of the outcomes at timestep $t-n$ (e.g., [1, 0, 0] in Figure 1), v_n is the corresponding validity (regression coefficient) of timestep $t-n$, and c is a constant probability assigned to the three holes on every trial. Coefficients v_n and constant c were varied between timesteps and conditions. The stimuli sequences were created beforehand and parameter recovery was conducted to ensure that the intended coefficients could be correctly recovered from the stimuli sequences.

In Experiment 1, outcomes regressed with 0.6 coefficient on one of three previous outcomes. Therefore, across Conditions 1 to 3, v_1, v_2, or v_3 was set to 0.6 respectively while the other coefficients were fixed at 0. Hence, on each trial, there was a 0.4 probability of the outcome being randomly chosen from the three holes, adding to each hole a constant probability c of 0.133 ($\frac{0.4}{3}$) across all trials.

In Experiment 2, across all conditions, outcomes regressed on two previous timesteps with differing coefficients. In Condition 1, v_1 was 0.4, v_2 was 0.2 and c was 0.133; in Condition 2, v_1 was 0.2, v_2 was 0.4 and c was 0.133; and in Condition 3, both v_1 and v_2 were 0.4 and c was 0.066. These combinations allowed us to study whether and how utilization of one timstep changes with the validity (and utilization) of another timestep, as expected through cue competition. In Experiment 2, across all conditions, v_3 was fixed at 0—we believe that two previous timesteps should sufficiently answer our questions regarding cue competition. Here input timesteps in the regression equation (i.e., $t-1$ and $t-2$) were correlated on every trial, which could impede accurate detection of regression coefficients. Parameter recovery was conducted to ensure that only appropriate sequences where multicollinearity did not interfere with coefficient detection were used in the experiment.

Experiment 1

Because the outcomes in this experiment regressed on only one of three previous outcomes, the regression equation simplified to the following across the three conditions with n being varied between 1 and 3:

$$P(outcome_t) = 0.6 \cdot outcome_{t-n} + 0.133 \tag{2}$$

The goal of this experiment was to determine whether and how participant utilization of previous timesteps reduced as distance to the current timestep increased. We hypothesize that utilization of the valid timestep would be highest in Condition 1 (where n is set to 1), lower in Condition 2 (n = 2), and lowest in Condition 3 (n = 3).

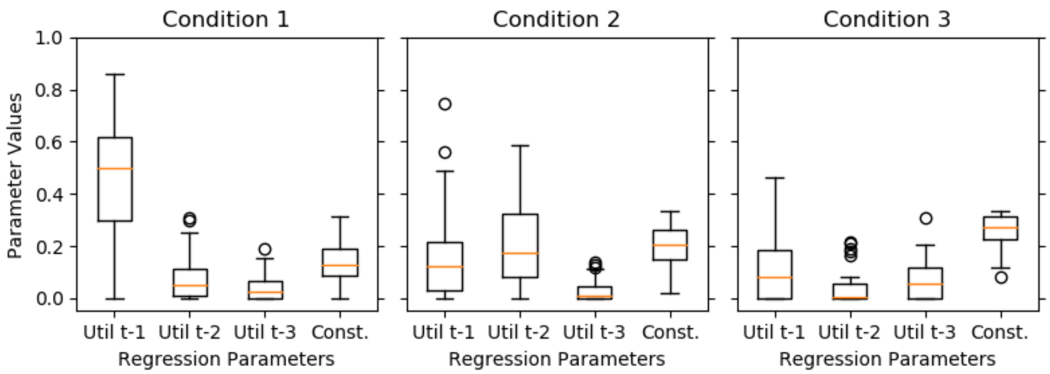

Figure 3: Distribution of participant's regression parameters for three conditions of Experiment 1. We find that utilization of valid timesteps is greatest for Condition 1 (where $t-1$ is the valid timestep), lower for Condition 2 ($t-2$ is the valid timestep), and lowest for Condition 3 ($t-3$ is the valid timestep).

Results and Discussion

We fit a regression model to participant data that was similar to that used for generating the data:

$$response\,weights_t = \sum_{n=1}^{3} u_n \cdot outcome_{t-n} + c; \qquad (3)$$

$$response\ prob_t = \frac{response\,weights_t}{\sum_{w=1}^{3} response\,weights_t^w} \qquad (4)$$

where u_n are participants' utilizations (subjective regression coefficients) and c is the constant probability participants assigned to the three holes across all trials. These parameters provide us with $response\,weights_t$ which is a vector representing the weight given to the three outcomes on timestep t. Luce's choice rule was used on these weights to provide $response\ prob_t$, the probabilities assigned by participants to the three outcomes.

Therefore, four parameters were fit to participants—u_1, u_2, u_3, and c. The regression model was fit using MLE and we used a multinomial distribution to determine the probability of participant data given the estimated response probabilities.

Figure 3 shows the distribution of regression parameters for the three conditions. We find that utilization of the valid timestep (u_1 for Condition 1, u_2 for Condition 2, and u_3 for Condition 3) varied across conditions—as anticipated, utilization of the regressing timestep reduced in conditions where it was more distant. This difference in utilization of the valid timestep across conditions was significant ($p < .001$). The reduced utilization of distant timesteps could be a result of difficulty in detecting regressions due to declining memory and/or failures in recollecting older outcomes during prediction.

We also find that in all three conditions, participants utilize recent timesteps more even if they have 0 validity—they tend to detect spurious regressions to recent timesteps more than to older ones (e.g., in Condition 2, u_1 is significantly greater than u_3, though both have 0 validity). In fact, in Condition 3, participants utilize timestep $t-1$ similarly as $t-3$ even though v_1 is 0 and v_3 is 0.6. This pattern suggests that participants are perhaps more sensitive to similarities to previous timesteps than to dissimilarities—they appear to "reward" timesteps (increase their strength) for their ability to predict outcomes more than they tend to "punish" them (decrease their strength) for their inability to predict. If they rewarded and punished them equally, we would expect that u_1 would be lower than u_3 in Condition 2—because of superior memory for recent timesteps, participants would punish u_1 more stringently for its inability to predict outcomes. We verify this account through our decision model in the Modelling section.

Experiment 2

Here, each trial regressed on two previous timesteps. In Condition 1, v_1 was 0.4, v_2 was 0.2 and c was 0.133; in Condition 2, v_1 was 0.2, v_2 was 0.4 and c was 0.133; and in Condition 3, both v_1 and v_2 were 0.4 and c was 0.066.

We anticipate that due to cue competition, the increased validity (and utilization) of one timestep will accompany the reduced utilization of another timestep. Hence, we expect that utilization of $t-1$ should be significantly lower in Condition 3 as compared to Condition 1 although v_1 is the same in both conditions and similarly, utilization of $t-2$ should be significantly lower in Condition 3 as compared to Condition 2.

Results and Discussion

We used the model described in Experiment 1 to estimate participant utilization, but truncated to only include regression coefficients u_1 and u_2 and constant c. We did not fit u_3 because we found that participants displayed $\sim$0 utilization of $t-3$ and it appeared to be irrelevant to our analyses on cue competition.

Figure 4 reports the distribution of u_1, u_2, and c for the three conditions. Utilization u_2 is slightly, but significantly, lower in Condition 3 than in Condition 2 ($p < .05$), consistent with cue competition. However, there is no difference in

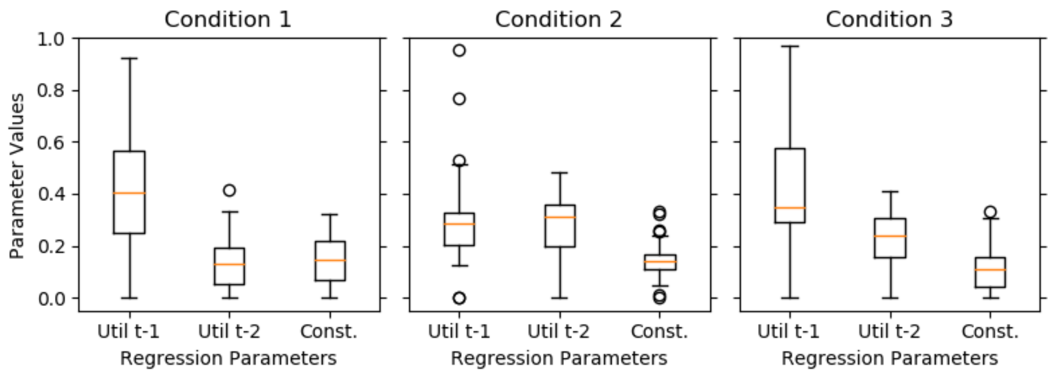

Figure 4: Distribution of participant's regression parameters for three conditions of Experiment 2. We find that utilization of timestep $t-2$ is significantly higher in Condition 2 compared to Condition 3 (possibly due to lower competition from timestep $t-1$). However, there is no significant difference in utilization of timestep $t-1$ between Condition 1 and 3.

u_1 between Conditions 1 and 3. This might be because older outcomes overall have less utilization, hence reducing competition from $t-2$ to utilization of $t-1$ in Condition 3. These results provide us with an inconclusive answer regarding cue competition—we attempt to understand cue competition further in the trial-by-trial decision model described next.

As in Experiment 1, we find an overall reduced utilization of u_2 as compared to u_1. In Condition 3, u_2 is significantly lower than u_1 although both have similar regression coefficients ($p < .01$). Further, in Condition 2, distribution of u_1 is similar to u_2, even though regression on $t-2$ is higher.

Modelling

We used a trial-by-trial decision model to formalize the mechanisms participants used as they performed the task. By comparing several versions of the model (using BIC) we attempted to determine decision mechanisms (i.e., utilization decline and cue competition) that were essential to participants' prediction behaviors. In this section, we report model comparison results using combined data from Experiments 1 and 2 since we expect participants to perform the same trial-by-trial decision processes across both experiments.

Participants are modeled as using outcomes from previous timesteps as cues for predicting the current timestep. Using a variation of the delta rule (Busemeyer & Stout, 2002), they estimate the weight $w_{t,t-n}$ they should give to timestep $t-n$ to predict timestep t as:

$$w_{t,t-n} = (1-\delta)\cdot w_{t-1,t-n} + \delta \cdot M_{t-n} \cdot (1-\alpha)^{n-1} \quad (5)$$

$$where\, M_{t-n} = \begin{cases} 1\, if\, O_t = O_{t-n}, \\ 0\, if\, O_t \neq O_{t-n} \end{cases} \quad (6)$$

Above, δ is the learning rate, with higher values producing quicker learning; M_{t-n} is 1 if O_t (the outcome from timestep t) and O_{t-n} (the outcome from timestep $t-n$, whose weight $w_{t,t-n}$ are being currently evaluated) are the same; and α is a decline rate multiplied to outcomes. High values of α impede learning of regression to older timesteps, thereby implementing a form of utilization decline.

In this version of the delta rule, we only reward timestep cues for accurate prediction and we do not punish them for incorrect prediction. We also fit data with a version of the punishing model (described in Kelley & Busemeyer, 2008) and obtained higher BIC values (see Appendix). This support for a model without punishment fits with our findings on the pattern of spurious regressions observed in Experiment 1.

We implement a specific form of cue competition (Kruschke & Johansen, 1999) as follows: Once weights are learned at timestep t, they compete for attention during prediction through a softmax function. These new competitive weights, $cw_{t,t-n}$ are calculated as:

$$cw_{t,t-n} = \frac{e^{\theta \cdot w_{t,t-n}}}{\sum_{m=1}^{3} e^{\theta \cdot w_{t,t-m}}} \quad (7)$$

where θ is the cue competition rate. In simulations we find that θ values close to 3.5 lead competitive weights to be roughly similar to initial input weights (i.e., no competition), higher values produce increases in already high weight values (i.e., positive competition), and lower values produce more equivalent weight values (i.e., negative competition).

Finally, $cw_{t,t-n}$ is used to make predictions for outcome probabilities on trial t:

$$pred\, weights = \sum_{n=1}^{3} cw_{t,t-n} \cdot O_{t-n} \cdot (1-\alpha)^{n-1}; \quad (8)$$

$$pred\, prob[i] = \frac{e^{\psi \cdot pred\, weights[i]}}{\sum_{j=1}^{3} e^{\psi \cdot pred\, weights[j]}} \quad (9)$$

Here, *pred weights* are the weights assigned to the three outcomes and are obtained by adding weighted timestep cues. Predicted probabilities of three outcomes, *predprob*, is obtained through a softmax version of Luce's choice rule where ψ represents amount of exploration by participants when selecting probabilities of the three outcomes. This is the same equation used in Eqn. 7; however, it is applied to predictions instead of regression coefficients $cw_{t,t-n}$. Although exploration of outcomes is not relevant directly to the topic of the current study, it is useful to include in the model since it allows for individual differences in exploratory responding.

Decline parameter α is used for suppressing older outcomes both in Eqn. 5 (at time of learning regression) and in Eqn. 8 (at time of prediction). Including α in Eqn. 5 leads it to interact with cue competition in Eqn. 7 such that older timesteps might receive a smaller boost from cue competition because of initial suppression. We also fit a model where α was only used during prediction in Eqn. 8 and not in Eqn. 5—here we obtained only a very small increase in BIC (1.45). Overall, the current model and experiment design are not sufficient to determine whether suppression is occurring during learning and/or prediction.

The full model described above has four parameters—learning rate (δ), decline rate (α), competition rate (θ), and exploration (ψ). We compared four versions of the model—the full model, one without utilization decline, one without cue competition, and one base model without both cue competition and decline. Comparing BICs for these four models allows us to estimate the importance of utilization decline and cue competition in predicting participant behavior.

Table 1 shows a comparison of the four models. The full model has lowest BIC values. However, the model without cue competition has only a slightly higher BIC. Further, we found that the average estimated cue competition rate $cw_{t,t-n}$ in Eqn. 7 was 3.9, producing very little competition. This corresponds with results from Experiment 2—cue competition could be a less essential aspect of participant behavior.

Table 1: BIC values of model variations.

Model Variation	Mean BIC Values
Full model	1217.53
Model without utilization decline	1263.65
Model without cue competition	1225.23
Base model (without utilization decline and cue competition)	1274.36

Discussion

The goal of the current study was to investigate how people use previous states of a system to predict a new state. Researchers have frequently studied people's abilities to predict systems using the states of other systems through the multiple cue learning paradigm—however, in our study, previous states of the same system serve as cues. Studies have also used statistical learning paradigms where participants learn fixed sequences of outcomes. Our study adds to the literature by using probabilistic sequences where best performance can be obtained only through identifying regression to previous timesteps. Many real world systems behave in similar ways—their current state is probabilistically dependent on their previous ones. We focused on two aspects of people's behavior in making predictions in such systems—decline in cue utilization with temporal distance and cue competition between predictive temporal cues.

Our studies found that decline in past cue utilization with temporal distance was central to people's behavior. Participants' subjective regression coefficients were significantly lower for older timesteps in Experiment 1 (Figure 3). Further, BIC values were considerably higher when utilization decline was removed from the decision model. Reduction in utilization could occur due to forgetting during learning and/or prediction. In previous work, we found that such recency-based decline was not correlated with working memory capacity (Luthra & Todd, 2019), suggesting that individual variation in it might not be mediated by memory limits.

Experiment 2 showed some evidence of cue competition in participant behavior—utilization of timestep $t-2$ reduced significantly when utilization of timestep $t-1$ was increased. However, we did not obtain a similar decrease in utilization of timestep $t-1$ with changes in utilization of $t-2$ (Figure 4). These mixed results could be because the difference in validity of the two previous timesteps in Conditions 1 and 2 was low (only 0.2). Further, in our work the predictive timesteps were correlated (due to temporal autoregression). This inescapable real-world tradeoff impacts the ability to investigate cue competition, contrasting with previous studies that found clear cue competition effects when using uncorrelated cues (Busemeyer et al., 1993).

Additionally, our work suggests that participants tend to reward cues for their ability to predict with a stronger weight than they punish them for their inability to predict. This is displayed by the greater utilization of recent timesteps even when they have 0 regression to the current timestep (Figure 3). This finding helps explain why people frequently assume positive autoregression (clumpiness) even when sequences are random (Scheibehenne et al., 2011)—they are more sensitive to spurious similarities with recent timesteps than with older timesteps. This tendency of primarily learning through positive examples has been displayed in other domains of human behavior—for instance, children successfully learn language by exposure to correct usage (Denis, 2001). Assuming positive autoregressions in the environment is possibly a cognitive adaptation to the real world structures that organisms encounter, which often display such characteristics.

The current work focuses on a unidirectional influence of the environment on human behavior—we study how people's predictions change with changes in their autoregressive environment. Future work should explore dynamic bi-directional interactions between the two, for instance, by simultaneously modelling how people's predictions of autoregressions get reflected in the environment structure through the choices they make. Identifying stable states of such systems can help explain and predict autoregressive structures of cultural environments (e.g., fast fashion trends; Luthra & Todd, 2021a).

Data, Code, and Appendix

Data, code (for models and plotting), and appendix are publicly available in the GitHub repository mahiluthra/autoregression-decisions.

Acknowledgments

We thank the Adaptive Behavior and Cognition Lab (West) and the Percepts and Concepts Labs for useful feedback on the paper. This research was supported in part by the John Templeton Foundation grant, "What drives human cognitive evolution".

References

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*(3), 193.

Busemeyer, J. R., Myung, I. J., & McDaniel, M. A. (1993). Cue competition effects: Empirical tests of adaptive network learning models. *Psychological Science*, *4*(3), 190–195.

Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, *14*(3), 253.

da Silva, C. F., Victorino, C. G., Caticha, N., & Baldo, M. V. C. (2017). Exploration and recency as the main proximate causes of probability matching: A reinforcement learning analysis. *Scientific Reports*, *7*.

Denis, F. (2001). Learning regular languages from simple positive examples. *Machine Learning*, *44*(1), 37–66.

Gigerenzer, G., & Goldstein, D. G. (1999). Betting on one good reason: The take the best heuristics. In G. Gigerenzer, P. M. Todd, & ABC Research Group (Eds.), *Simple heuristics that makes us smart* (pp. 75–95). Oxford University Press.

Hyman, R., & Jenkin, N. S. (1956). Involvement and set as determinants of behavioral stereotypy. *Psychological Reports*, *2*(3), 131–146.

Kelley, H., & Busemeyer, J. (2008). A comparison of models for learning how to dynamically integrate multiple cues in order to forecast continuous criteria. *Journal of Mathematical Psychology*, *52*(4), 218–240.

Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(5), 1083.

Luthra, M., & Todd, P. M. (2019). Role of working memory on strategy use in the probability learning task. In A. Goel, C. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st annual conference of the cognitive science society* (pp. 721–728). Cognitive Science Society.

Luthra, M., & Todd, P. M. (2021a). Bursty spatial and temporal activity resulting from social search. *ALIFE Workshop on Agent-Based Modelling of Human Behavior*.

Luthra, M., & Todd, P. M. (2021b). People adjust recency adaptively to environment structure. In T. Fitch, C. Lamm, H. Leder, & K. Teßmar-Raible (Eds.), *Proceedings of the 43rd annual conference of the cognitive science society* (pp. 569–575). Vienna: Cognitive Science Society.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928.

Salisu, A. A., Gupta, R., & Ogbonna, A. E. (2022). A moving average heterogeneous autoregressive model for forecasting the realized volatility of the US stock market: Evidence from over a century of data. *International Journal of Finance & Economics*, *27*(1), 384–400.

Scheibehenne, B., Wilke, A., & Todd, P. M. (2011). Expectations of clumpy resources influence predictions of sequential events. *Evolution and Human Behavior*, *32*(5), 326–333.

Speekenbrink, M., & Shanks, D. R. (2010). Learning in a changing environment. *Journal of Experimental Psychology: General*, *139*(2), 266.

Worthy, D. A., & Maddox, W. T. (2014). A comparison model of reinforcement-learning and win-stay-lose-shift decision-making processes: A tribute to W. K. Estes. *Journal of Mathematical Psychology*, *59*, 41–49.