Modeling Dynamic Control in Normal Aging

Brian D. Glass (b.glass@qmul.ac.uk)
Biological and Experimental Psychology Centre, School of Biological and Chemical Sciences, Queen Mary College, University of London, Mile End, London E1 4NS, UK.

Magda Osman (m.osman@qmul.ac.uk)
Biological and Experimental Psychology Centre, School of Biological and Chemical Sciences, Queen Mary College, University of London, Mile End, London E1 4NS, UK.

Abstract

Complex and dynamic decision making environments are common throughout life, but little is known about how normal aging influences performance on these types of scenarios. To determine performance differences associated with normal aging, we test older and younger adults in a dynamic control task. The task involves the control of a single output variable via multiple and uncertain input controls. A computational model is developed to determine the behavioral characteristics associated with normal aging in a dynamic control task. Older adults exhibit a positivity effect, congruent with previous research. Model based analysis demonstrates a unique performance signature profile associated with normal aging.

Keywords: dynamic decision making; learning; normal aging; computational modeling

Introduction

Normal human aging is associated with cognitive changes that lead to differences in the way older adults approach and perform in decision making tasks. Specifically, older adults appear to suffer from executive control deficits (Braver, et al., 2001; Kray, Li, & Lindenberger, 2002; Ortega, et al., 2012). However, emerging evidence suggests that older adults can utilize compensatory strategies to return performance to or beyond baseline levels (Glass, et al., 2012; Huang, et al., 2012; Worthy, et al., 2011).

While previous research has focused on classic paradigms such as category learning, task switching, and single-response choice procedures, little is known about normal aging in dynamic control tasks for which the participant controls multiple input variables in an integrative and uncertain task environment. Such complex dynamic environments are analogous to many real-life situations. For example, we make several distinct health choices on a daily basis which influence our overall health and wellbeing in uncertain ways. These types of environments are often noisy and the specific influence of the various choices is often unclear or unspecified.

The present research contrasts older adult and younger adult performance in a dynamic control task designed to simulate such real-life dynamic decision making environments (Osman & Speekenbrink, 2011). A novel computational modeling technique is developed to assess individual behavioral characteristics and strategies in the dynamic control task.

Method

Procedure

In the present dynamic control task, the participant attempts to control a single outcome value towards a goal. To do so, on each trial the participant chooses values for three separate cues. These cue values are then combined via the dynamic control equation (Equation 1) then summed with the outcome value plus some normally distributed random noise (standard deviation = 8). In this way, the participant's cue selections guide the outcome value. The outcome value is initialized at 178 with a goal value of 62 and a “safe range” (±10 around the goal value)

$$\boldsymbol{y(t) = y(t-1) + 0.65x_1 - 0.65x_2 + e}$$

Equation 1.

where $y(t)$ is the outcome on trial t, x_1 is the positive cue, x_2 is the negative cue, and e is an error term randomly sampled from a normal distribution with a mean of 0 and SD of 8.

The dynamic control equation was designed such that one cue has a positive impact on the outcome value, one cue has a negative impact, and a third cue has no impact. The impact of the cue is not labeled or available to the participant, thus the participant must learn to control the outcome value based solely on the resulting movement of the outcome value on each trial. After each trial, the cue input values are reset to 0. The participant can then freely select input values for each of the three cues before confirming the choices.

A critical feature of this control task is that the outcome value can swing below the target, meaning the participant must dynamically adapt in order to maximize performance. After an initial learning phase, the participants completed 2 Test blocks of 20 trials each. The first Test block was a Congruent Test in which the starting value and goal criterion were equivalent to the learning phase. The second Test block was a Transfer Test with a different starting value and goal value than the earlier phases. At the beginning of each block, the control task was reset to the initial state.

Participants

27 younger participants aged 18 to 25 (M = 22.3, SD = 5.4) and 15 older participants aged 61 to 75 (M = 67.92, SD = 5.03) participated in the dynamic control task. The younger participants were recruited from the Queen Mary, University of London undergraduate community and paid £6 ($9.50). The older participants were recruited via the National Hospital for Neurology and Neurosurgery. The older adults were recruited as a healthy control group via the National Hospital for Neurology and Neurosurgery. To qualify for the healthy adult participation pool, the older adults completed the. Beck Depression Inventory-II (BDI-II; Beck, et al., 1996) and Mini-Mental State Examination (MMSE) (Folstein, et al., 1975). All scores fell within the normal cutoff range for both the MMSE (greater than 27) and BDI-II (less than 18). None of the HCs had a history of neurological or physical or psychiatric illness, head injury or drug or alcohol abuse.

Computational Model

A computational model of behavior in the dynamic control task was constructed to determine behavioral characteristics of individual participants. The model is based on memory trace reinforcement learning. After each trial, a reinforcement history for each of the three cues is updated according to whether the cue choices resulted in the discrepancy between achieved outcome value and goal value increasing or decreasing. On the following trial, the reinforcement history becomes the basis for a probabilistic action selection function using Luce's choice. Previous research has found that participants often vary the value of more than one cue on each trial. Thus, the model includes an inter-cue gating mechanism which allows each cue value selection to take into account the action selection probabilities of the other two cues.

The resulting model features four free parameters: an exploitation parameter governing the action selection function, an inter-cue gating parameter, and two memory-updating reinforcement strengths (one for successful trials, and one for unsuccessful trials). To evaluate the model, the model's probability of selecting the human participant's cue choice are combined across all trials and all three cues into a single model fit value. The model is fit to an individual participant's responses by an optimization procedure that determines the parameter values which maximize the fit value.

Memory-Updating Reinforcement Strengths

After each trial, the computational model determines whether the cue values it selected resulted in the outcome value moving towards or away from the goal. For each cue, a Gaussian curve with a mean equal to the chosen cue is constructed (Equation 2).

$$P_{update}(v) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{v-v_p}{\sigma}\right)^2}$$

Equation 2.

where P_{update} (v) is the probability of selecting a value of v when the previous selected value was v_p.

This curve is then summed (successful trial) or subtracted (unsuccessful trial) to the cue's former reinforcement history. A free parameter (one for successful trials, one for unsuccessful trials) determines the relative weight of the updating summation. For example, if the memory-updating positive reinforcement strength is 0.8, then the reinforcement history is updated such that 80% of the new reinforcement history reflects the current cue value choice and 20% reflects the previous reinforcement history (Equation 3).

$$P_{History}(v) = [(1-\gamma_s)P(v)] + \left[sign(R) \cdot \gamma_s \cdot\ P_{update}(v)\right]$$

Equation 3.

where $P_{History}(v)$ is the cue selection probability history for cue value v, γ_s is the memory-updating reinforcement strength for feedback s (positive or negative), and R is the change in the outcome value's distance to the goal from the previous trial.

In summary, there are two memory-updating reinforcement strengths, one for positive outcomes and one for negative outcomes. Each strength represents the weight with which current choices impact choice history (see Figure 1).

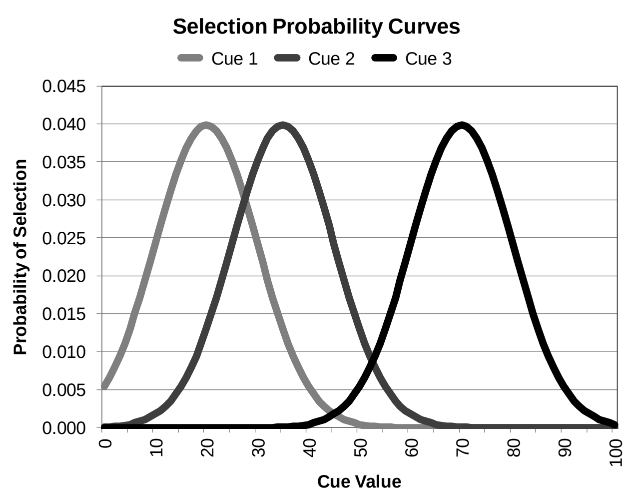

Figure 1: Sample probability density curves of selecting a given value for a given cue. Over the course of a block, the curves will alter in various ways depending on the model parameters, trial success, and uncertainty inherent in the outcome value.

Inter-cue Parameter

Before the final probabilistic selection of the cue value occurs, for each of the three cues, the reinforcement history of the two other cues are taken into consideration. The level of this consideration is controlled by an inter-cue parameter. This parameter determines the strength at which the reinforcement history of other two cues will influence the action selection of the cue at hand. This is done using a gating equation which weights the alternate cues using the inter-cue parameter (Equation 4).

$$P_{Intercue}(v_{c_A}) = \left[\left(1-\frac{2\beta}{3}\right)P_{History}(v_{c_A})\right] + \left[\frac{\beta}{3}\cdot P_{History}(v_{c_B})\right] + \left[\frac{\beta}{3}\cdot P_{History}(v_{c_C})\right]$$

Equation 4.

where $P_{Intercue}(v_{cA})$ is the probability of selecting value v for cue c_A (e.g., cue 1), β is the inter-cue parameter, and c_A and c_B are the other two cues (e.g., cue 2 and 3). As the inter-cue parameter approaches 1, the computational model is more likely to pick similar cue values for all three cue inputs. As the inter-cue parameter approaches 0, the model is less likely to select an action for one cue based on the reinforcement history of the other two. In this way, the computational model can vary the strength in which cue values vary together in the action selection state of the decision process.

Exploration Parameter

On each trial, the computational model evaluates the reinforcement history of each cue to generate the probability of selecting each of the 100 cue value options. From these options, a single value is chosen using the Softmax decision rule (Equation 5). The equation's exploitation parameter, K, determines the level of determinism in the choice process (Daw & Doya, 2006). As K approaches ∞, the process is more likely to choose the most probable option. As K approaches 0, the equation is more likely to pick a less probable option.

$$P_{Final}(v_i) = \frac{e^{[P_{Intercue}(v_i)\cdot K]}}{\sum_{j=0}^{100} e^{[P_{Intercue}(v_j)\cdot K]}}$$

Equation 5.

where P_{Final} (v_i) is the final probability of selecting cue value v_i, K is the exploitation parameter, and v_j are all the cue values from 0 to 100 for given cue.

Results

Task Analysis

By considering the optimal cue actions that will maximize the outcome value's movement toward the target, the optimal selections can be computed for each trial (Equation 5). The difference between the optimal selections and the actual chosen selections results in an optimality score for each participant. Figures 2 and 3 shows that the Younger group tended to select more optimal responses in both Test blocks, although the difference was not statistically significant.

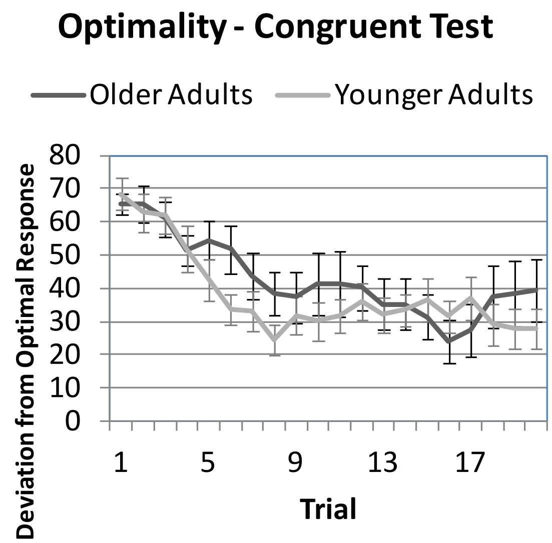

Figure 2: Optimality scores for Congruent Test block

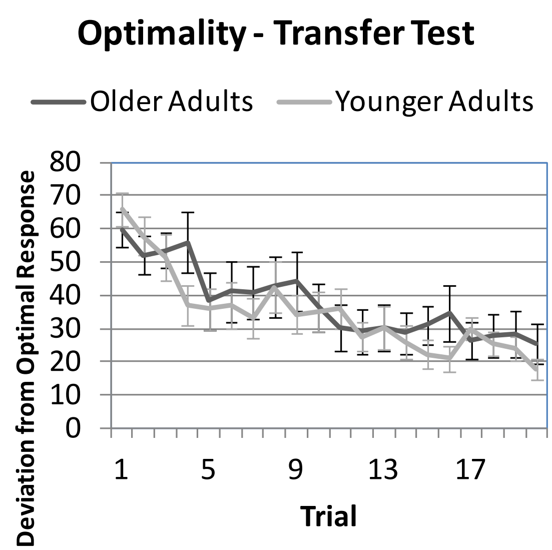

Figure 3: Optimality scores for Transfer Test block

At first blush, it may seem that the Older group performed similarly to the Younger group. However, further analysis of the strategies used by both groups demonstrates critical differences in the way the Older adults completed the dynamic control task. The strategy analysis considered four different types of cue changes: varying none, varying one cue, varying two cues, and varying all three cues. Figures 4 and 5 illustrate the cue varying strategies for both groups on both the Congruent Test and Transfer Test. A 2 (Older, Younger) x 2 (Congruent, Transfer) x 4 (Strategy Type) repeated measures ANOVA reveals an Age by Block by Strategy interaction, $F(3, 120) = 2.95$, $p < 0.05$, $\eta = 0.07$. There was also a main effect of strategy, $F(3, 120) = 24.42$, $p < 0.001$, $\eta = 0.38$. No other main effects of interactions were statistically significant.

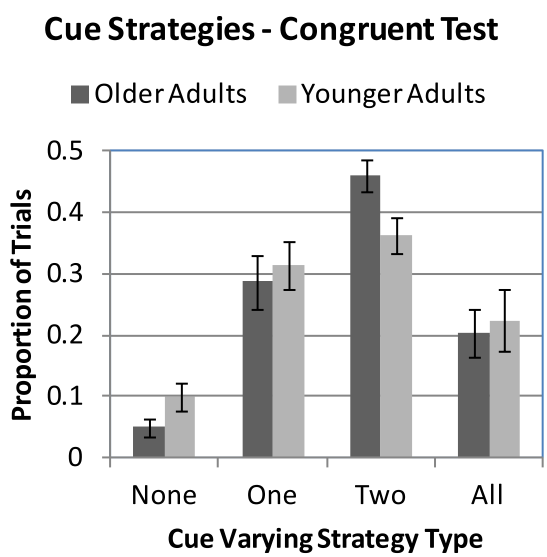

Figure 4. Cue varying strategies for Congruent Test

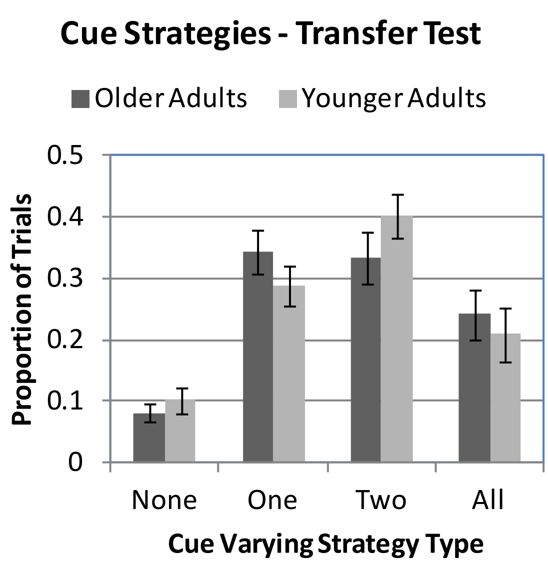

Figure 5. Cue varying strategies for Transfer Test

Not only did the Younger and Older groups differ in their cue varying strategies, they also differed in the values selected for the cues. Figures 6 and 7 report the mean cue values (between 0 and 100) selected for each of the three Cue Types. A 2 (Older, Younger) x 2 (Congruent, Transfer) x 4 (Strategy Type) repeated measures ANOVA revealed a main effect of Cue Type, $F(2, 80) = 5.11$, $p < 0.01$, $\eta = 0.11$, as well as an interaction of Age and Cue Type, $F(2, 80) = 3.51$, $p < 0.05$, $\eta = 0.08$. This suggests that the Older group tended to select higher values for the Positive and Null cues

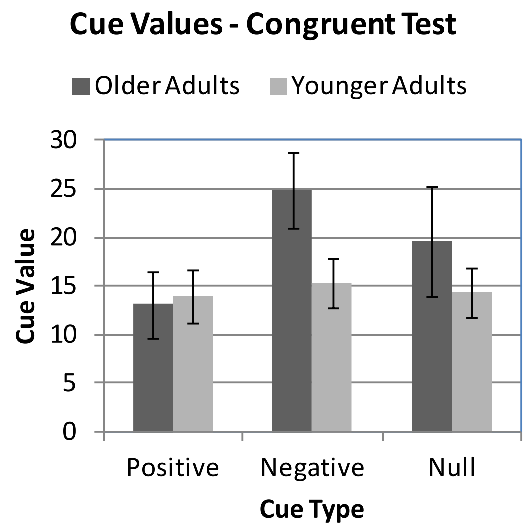

Figure 6. Cue values selected for Congruent Test

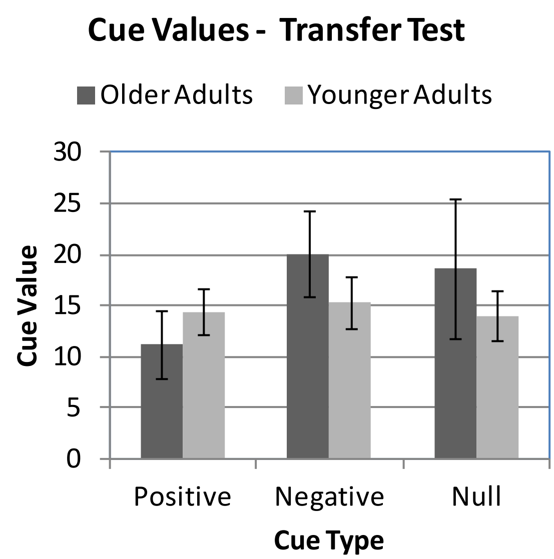

Figure 7. Cue values selected for Transfer Test

Taken together, analyses of surface level behavior suggest the Older group differed from the Younger group in completing the dynamic control task. However, the nature of the underlying cognitive processes which led to these patterns of behavior remains elusive using basic task analysis. In order to distill psychologically relevant characteristics of the processes involved in the dynamic decision making task performance, a computational reinforcement learning model of the dynamic control task was fit to individual participant data.

Model Based Analysis

Task behavior was fit to the computational model using an optimization procedure that attempted to minimize the difference between observed trial-by-trial cue value selections and the expected cue value selections as determined by the model. This was done by considering the probabilities given to the various cue values for each cue on a given trial. The optimization procedure attempted to determine best fitting free parameters (exploitation parameter, inter-cue parameter, positive and negative reinforcement sensitivity parameter) that maximized the probability that the model would select the same cue values as the human participant on a given trial.

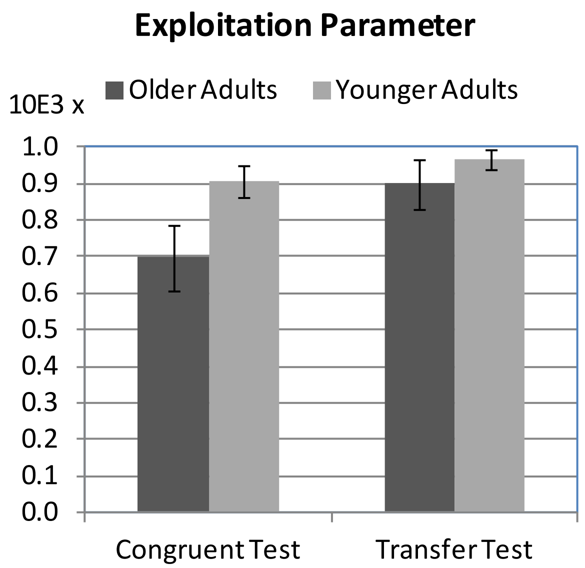

Figure 8. Exploitation Parameter

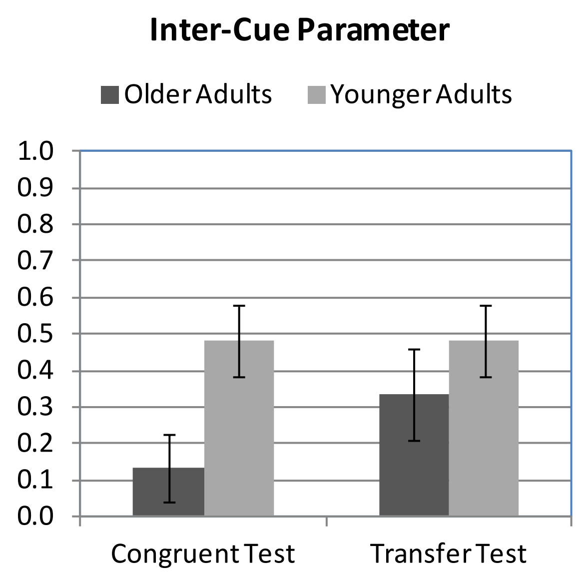

Figure 9. Inter-Cue Parameter

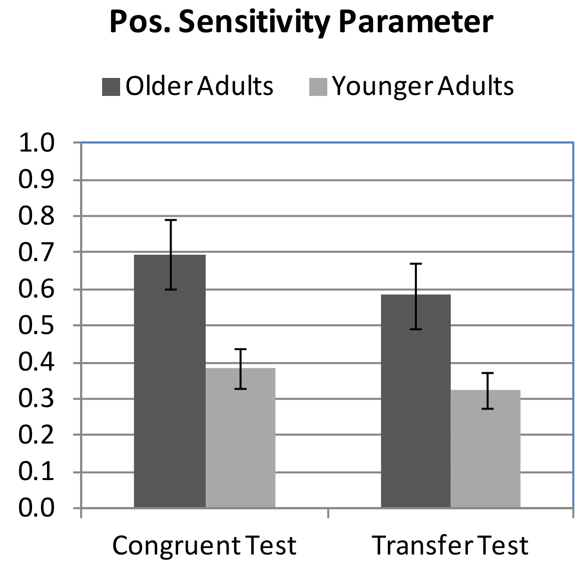

Figure 10. Positive Sensitivity Learning Parameter

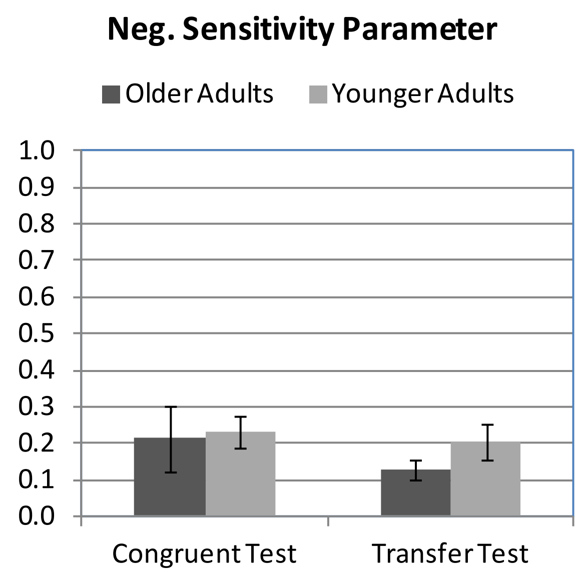

Figure 11. Negative Sensitivity Learning Parameter

Figures 8 through 11 reports the mean best fit parameter values for the Younger and Older groups. In the Congruent Test, those in the Older groups were best fit with a lower exploitation parameter ($t[40] = -2.37$, $p = 0.02$), a higher positive reinforcement strength parameter ($t[40] = 3.17$, $p = 0.003$), and a lower inter-cue parameter ($t[40] = -2.35$, $p = 0.02$). There was no significant difference in the negative reinforcement strength parameter between the two groups, $t(40) = -0.29$, $p = 0.85$. In the Transfer Test, the Older adults continued to be better fit by a higher positive reinforcement parameter than Younger adults, $t(40) = 2.74$, $p < 0.01$. In short, in the Congruent Test, the Older group's performance was better fit with parameters associated with higher exploration, higher positive feedback sensitivity, and lower inter-cue selection. In the Transfer Test, the Older group continued to be better fit by model parameters associated with higher positive feedback sensitivity.

Discussion

The present study examined the role of normal aging in a dynamic control task using a novel computational modeling technique. Standard behavioral analysis revealed older adults potentially utilized an alternative strategy in completing the dynamic control task than younger adults. A computational model of the task revealed specific behavioral characteristics associated with normal aging. In the Congruent block, older adults demonstrated more exploratory behavior, less inter-cue behavior, and more reliance on recent and positive success. On the Transfer block, older adults did not differ from younger adults in their exploratory and inter-cue behavior, but continued to demonstrate more reliance on recent and positive success.

One possible interpretation of this pattern of results is that older adults were able to achieve the final performance profile of younger adults (as measured by deviation from optimal responses) by relying on compensatory mechanisms to engage the task. Specifically, in the congruent goal test, the older adults were more exploratory, relied less on the reinforcement history of alternative cues when determining cue values, and were more influenced by trials on which they received positive feedback. During the transfer goal test, the older adult's compensatory strategy gave way to a closer performance signature exhibited by younger adults. However, they remained more influenced by positive feedback. This interpretation is supported by previous research which has shown that older adults are able to achieve the performance levels of younger adults via a compensatory strategy (Glass, et al., 2012;Worthy & Maddox, 2012).

Another interpretation of the current results is that older adults approached the task by utilizing alternative mechanisms which may be enhanced in older adults. For example, older adults exhibit a positivity effect characterized by superior emotional processing of positively valenced content (Carstensen & Mikels, 2005). This could account for the older adults' higher learning rate sensitivity parameter for positive feedback, but not for negative feedback. Thus, when older adults encountered successful trials, their learning rate parameters increased such that prior knowledge was discounted. In this interpretation, older adults differed in their overall strategy due to specific enhancements associated with normal aging. This interpretation is supported by the positive learning rate sensitivity parameter remaining higher for older adults than younger adults in both the congruent and transfer tasks. Future research should determine whether the differences in strategies used by older adults to complete the dynamic control task are simply the result of slower overall learning rates, or due to differences in underlying cognitive mechanisms associated with normal aging. Future work should incorporate manipulations to test these interpretations, such as limiting feedback types to determine whether the aging positivity effect can account for performance differences.

Acknowledgments

The authors thank Agata Ryterska for help with data collection for the project.

References

Beck, A. T., Steer, R. A., Ball, R., & Ranieri, W. F. (1996). Comparison of Beck Depression Inventories-IA and-II in psychiatric outpatients. *Journal of Personality Assessment*, *67*, 588-597.

Braver, T. S., Barch, D. M., Keys, B. A., Cohen, J. D., Taylor, S. F., Carter, C. S., Kaye, J. A., Janowsky, J. S., Yesavage, J. A., Mumenthaler, M. S., Jagust, W. J., & Reed, B. R. (2001). Context processing in older adults: Evidence for a theory relating cognitive control to neurobiology in healthy aging. *Journal of Experimental Psychology: General*, *130*, 746-763.

Carstensen, L. L. & Mikels, J. A. (2005). At the intersection of emotion and cognition. *Current Directions in Psychological Science*, *14*, 117-121.

Daw, N. D. & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, *16*, 199-204.

Folstein, M. F., Folstein, S. E. & McHugh, P. R. (1975). Mini-Mental State: A practical method for grading the cognitive state of patients for the clinician. *Journal of Psychiatric Research*, *12*, 189-198.

Glass, B. D., Chotibut, T., Pacheco, J., Schnyer, D. M., & Maddox, W. T. (2012). Normal aging and the dissociable prototype learning systems. *Psychology and Aging*, *27*, 120-128.

Huang, C. M., Polk, T. A., Goh, J. O., Park, D. C. (2012). Both left and right posterior parietal activations contribute to compensatory processes in normal aging. *Neuropsychologia*, *50*, 55-66.

Kray, J., Li, K. Z. H., & Lindenberger, U. (2002). Age-specific changes in task-switching components: The role of task uncertainty. *Brain and Cognition*, *49*, 363-381.

Ortega, A., Gómez-Ariza, C. J., Román, P. E., & Bajo, M. T. (2012) Memory inhibition, aging and the executive deficit hypothesis. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *38*, 178-186.

Osman, M., & Speekenbrink, M. (2012). Prediction and control in a dynamic environment. *Frontiers in Psychology*, 3:68.

Worthy, D. A., Gorlick, M. A., Pacheco, J. L., Schnyer, D. M., & Maddox, W. T. (2011). With age comes wisdom: Decision making in younger and older adults. *Psychological Science*, *11*, 1375-1380.

Worthy, D.A. & Maddox, W. T. (2012). Age-based differences in strategy-use in choice tasks. *Frontiers in Neuroscience*, *5*, 1-10.