Published in final edited form as:

Neuroimage. 2016 April 15; 130: 13–23. doi:10.1016/j.neuroimage.2015.12.004.

Neural Correlates of State-Based Decision-Making in Younger and Older Adults

Darrell A. Worthy¹, Tyler Davis², Marissa A. Gorlick³, Jessica A. Cooper³, Akram Bakkour⁴, Jeanette A. Mumford⁵, Russell A. Poldrack⁶, and W. Todd Maddox³

1

2

3

4

5

6

Abstract

Older and younger adults performed a state-based decision-making task while undergoing functional MRI (fMRI). We proposed that younger adults would be more prone to base their decisions on expected value comparisons, but that older adults would be more *reactive* decision-makers who would act in response to recent changes in rewards or states, rather than on a comparison of expected values. To test this we regressed BOLD activation on two measures from a sophisticated reinforcement learning (RL) model. A value-based regressor was computed by subtracting the immediate value of the selected alternative from its long-term value. The other regressor was a state-change uncertainty signal that served as a proxy for whether the participant's state improved or declined, relative to the previous trial. Younger adults' activation was modulated by the value-based regressor in ventral striatal and medial PFC regions implicated in reinforcement learning. Older adults' activation was modulated by state-change uncertainty signals in right dorsolateral PFC, and activation in this region was associated with improved performance in the task. This suggests that older adults may depart from standard expected-value based strategies and recruit lateral PFC regions to engage in reactive decision-making strategies.

Keywords

Aging; Decision-Making; fMRI; reinforcement learning; computational modeling

Introduction

Reward-based decision-making involves selecting from among multiple alternatives in order to maximize rewards and/or minimize losses. Expected Value theory is the dominant theory

regarding how individuals make decisions in reward-based decision-making tasks. This theory proposes that individuals learn, from feedback, to predict which options will be more rewarding than others, and that their behavior is guided by such a comparison (Edwards, 1954; Rangel, Camerer, & Montague, 2008; Samanez-Larkin & Knutson, 2015). This representation and comparison of expected values of alternative choices may be particularly compromised by the normal aging process.

Recently there has been a surge of work aimed at examining how the neurobiological, cognitive, and social changes associated with the normal aging process affect reward-based decision-making (Lighthall, Huettel, & Cabeza, 2014; Maddox, Gorlick, & Worthy, 2015; Mata, Josef, Samanez-Larkin, & Hertwig, 2011; M. Mather, 2005; Mara Mather & Carstensen, 2005; Samanez-Larkin & Knutson, 2015; Worthy, Gorlick, Pacheco, Schnyer, & Maddox, 2011). Older adults experience age-related declines in the integrity of the mesolimbic dopamine system, and its associated neural structures, that is critical for tracking and representing expected reward values (Bäckman, Nyberg, Lindenberger, Li, & Farde, 2006; Chowdhury et al., 2013; S.-C. Li, Lindenberger, & Sikström, 2001). Another study demonstrated that decline in the white matter integrity of glutamatergic pathways may also impact reward-based decision-making (Samanez-Larkin, Levens, Perry, Dougherty, & Knutson, 2012). However, despite these declines older adults often show decision-making behavior that is equally as good as or even better than younger adults (Worthy et al., 2011; Worthy, Otto, Doll, Byrne, & Maddox, 2015). One way in which older adults might compensate for decline in neural structures that mediate expected-value based decision-making is to use alternative, heuristic-based strategies that are heavily reliant on recent outcomes. Several studies have demonstrated older adults' abilities to be adaptive decision-makers who utilize alternative strategies to achieve equivalent performance to that of younger adults (Castel, Rossi, & McGillivray, 2012; Mata, Josef, & Lemaire, 2015; Worthy, Cooper, Byrne, Gorlick, & Maddox, 2014; Worthy et al., 2011; Worthy & Maddox, 2012).

While it is likely that both older and younger adults make decisions based on the subjective value of each alternative, these subjective values may be based on different information for younger and older adults. Younger adults' subjective values may closely correspond to the long-term expected value of each option. Older adults, however, may base their subjective values of each alternative on whether the rewards they have received for each action were improvements or declines to rewards received on earlier trials. For example, older adults are more likely to switch to a different option following steep declines in reward than younger adults (Worthy et al., 2015), and are more likely to utilize heuristics such as 'win-stay-lose-shift' which are solely based on whether the last outcome was an improvement or a decline in reward relative to the previous outcome (Worthy & Maddox, 2012). Older adults are also more likely to display a 'model-free' pattern of decision-making in the two-step task where they tend to switch following unrewarded trials and stay following rewarded trials regardless of whether transitions to second stage states were common or rare. Younger adults, however, are more likely to make decisions based on model-based expected values of each option (Eppinger, Walter, Heekeren, & Li, 2013). We propose that this tendency for older adults to use heuristics or to be more adaptive in the types of strategies they use leads them to be more 'reactive' decision-makers who act in response to recent outcomes rather than based on a comparison of expected value (Worthy et al., 2015).

Further evidence supporting this assertion comes from a recent study by Vink and colleagues that found increased activation in reward related brain regions in younger adults during reward anticipation, but greater activation in older adults during reward receipt (Matthijs Vink, Kleerekooper, van den Wildenberg, & Kahn, 2015). Several other studies have found similar results (Dreher, Meyer-Lindenberg, Kohn, & Berman, 2008; Samanez-Larkin et al., 2007; Samanez-Larkin, Kuhn, Yoo, & Knutson, 2010; Schott et al., 2007; Spaniol, Bowen, Wegier, & Grady, 2015). The aging process appears to reduce the tendency for activation in reward-related regions to shift from reward receipt to reward anticipation over the course of learning, with older adults responding more to reward receipt throughout the course of the task than younger adults (Matthijs Vink et al., 2015). Rather than basing decisions on a comparison of the relative expected values of each action, older adults may operate in a more reactive manner and base their decisions on the influence that an action has had on recent changes in state or received rewards. Thus for older adults, the subjective value of each alternative may be based on whether selecting each alternative led to improvements or declines in states or rewards, rather than on expected values derived from reinforcement-learning models. While we acknowledge that our theory that older adults are more reactive decision-makers than younger adults is relatively new, it is nevertheless consistent with recent work demonstrating an enhanced reliance on recent outcomes for older adults during decision-making situations.

In the current experiment, older and younger adults performed a state-based dynamic decision-making task while undergoing MRI. The task required participants to learn how actions led to either improvement or declines in their future state, which ultimately affected their long-term cumulative reward. Previous work in our labs suggests that older adults can perform as well as or better than younger adults in this task, although they may rely on more reactive, heuristic-based strategies than on comparing expected values (Worthy et al., 2014; Worthy et al., 2011). To test our theory that older and younger adults would base their decisions on different types of information we regressed the blood-oxygen-level dependent (BOLD) signal on estimates from a state-based reinforcement-learning model that we have used in prior behavioral work (Worthy et al., 2014).

As further detailed below, one regressor we computed was a *relative value* component defined as the difference between the state-based and reward-based value of the action that was selected on each trial. This served as a proxy for the relative long-term value of each option compared to the immediate expected value. The second regressor we computed was a *state-change uncertainty signal* that represented whether the prior action led to an improvement or a decline in each participant's state and how uncertain or unexpected the change in state was (detailed below). These state-change uncertainty signals should be very useful in allowing participants to learn which actions lead to improvements or declines in future states. Based on our theory that older adults are more reactive decision-makers who base their actions on recent changes in rewards or states rather than on expected value comparisons we predicted that BOLD activation in older adults would be more related to the state-change uncertainty signals compared to younger adults, who would show greater activation related to the relative long-term value of each option. Based on recent work highlighting enhanced frontal compensation in older adults we predicted that areas of the

dorsolateral prefrontal cortex (DLPFC) would show enhanced activation related to state-change uncertainty signals, and furthermore, that this compensatory activation would be related to improved performance in the task (Cabeza, Anderson, Locantore, & McIntosh, 2002; Park & Reuter-Lorenz, 2009; Reuter-Lorenz & Cappell, 2008).

To test this prediction we performed a region of interest (ROI) analysis in this region as well as whole-brain analyses. A goal of the analysis was to determine whether any state-change uncertainty-weighted activation in this region was tied to enhanced performance in older adults. The link between activation and improved performance has been proposed as a way to test whether activation can be considered compensatory (Lighthall et al., 2014). We also predicted that younger adults' activation would be more strongly tied to the expected value regressor compared to older adults, particularly in ventral striatal and medial PFC regions commonly implicated in value-based decision making (Hare, O'Doherty, Camerer, Schultz, & Rangel, 2008; Rangel et al., 2008; Samanez-Larkin, Worthy, Mata, McClure, & Knutson, 2014).

Materials and Methods

Participants

Participants from the Austin community and students of the University of Texas at Austin were recruited from alumni mailings, fliers, and newspaper ads. Eighteen healthy younger adults (mean age 23.61 years, range 18-31; 10 F; mean years of education = 15.64) and eighteen healthy older adults (mean age 67 years, range 61-79; 8 F; mean years of education = 18.37) were included in this study and compensated with \$10/hr for their participation. Five additional subjects were recruited but were excluded from analysis (for non-completion of study (3), experimenter error (2), and structural abnormalities (1)). All volunteers gave informed written consent according to procedures approved by the University of Texas at Austin Internal Review Board. All volunteers were right-handed native English speakers.

Before completing the study all participants were screened for conditions that prevent them from being in an MRI environment. Participants were also screened for neurological disorders, drugs known to influence blood flow and/or cognition. Older adults were administered a battery of neuropsychological tests (assessing attention, verbal memory, visual memory, speed, and executive function) in order to determine whether they were functioning within the normal range for their age.¹

Decision Making Task

Participants performed four 75-trial runs of a two-option state-based dynamic decision making task where current rewards depended on past choices (history-dependent; Otto, Markman, Gureckis, & Love, 2010 Figure 1A; Worthy et al., 2011). On any given experimental trial, the *reward-maximizing* option gives higher rewards than the *state-maximizing* option. However, the state participants are in improves as the state-maximizing

¹Other similar models have recursively derived long-term values of each action over infinite horizons using the Bellman Equation. Here we simply estimated the maximum rewards that could be obtained from transitions to each possible future state on the next trial. This provided an expected reward value that was directly comparable to the reward-based value provided by the SARSA learner.

option is chosen more frequently (move to the right on the x-axes in Figure 1A), whereas the state declines as the reward-maximizing option is chosen more frequently (move to the left on the x-axes in Figure 1A). The number of state-maximizing options selected was initialized to 5 on the first trial of each run. This guarantees that all participants start the task at the same point on the reward functions, and does not bias participants toward one or the other end states (0 or 10 state-maximizing option selections during the past 10 trials). Thus, reward values for both options depend on how often the state-maximizing option has been chosen over a window of the last 10 trials. Each experimental trial lasted for four seconds where participants could select one of the two options for 2 seconds followed by a 2 seconds of feedback. After every 25 experimental trials there was a 16-second fixation trial.

A hypothetical scenario was presented where participants selected one of two oxygen extraction systems on Mars (Figure 1B). Participants were told they would repeatedly test the two systems and a bar, representing a small oxygen tank, would show the amount of oxygen they had extracted on the current trial. A line on the larger cumulative tank corresponded to their goal – collecting enough oxygen collected to complete their mission. In order to reach this goal, participants needed to select the state-maximizing option 80% of the time across four runs. They were told nothing about the underlying reward structure, but were required to learn from experience.

Imaging Acquisition

T1-weighted scans were collected from each participant using 3 Tesla GE Signa EXCITE MRI equipped with an eight channel phased array head coil (*GE Medical Systems, Milwaukee Wisconsin*) at the University of Texas at Austin Imaging Research Center. Head motion was minimized with foam inserts. Stimuli were viewed using a back projection screen and a mirror mounted on top of the head coil. MR safe glasses were provided as needed to correct vision to normal. Responses were collected with two four-button MR compatible optical transmission device, one held in each hand.

Functional EPI images were collected using a parallel imaging approach with GRAPPA reconstruction utilizing whole head coverage with slice orientation to reduce artifact (approximately 10% grade off the AC–PC plane, TR = 2s, TE = 30ms, 32 axial slices oriented for best whole head coverage, acquisition voxel size = $3.125 \times 3.125 \times 3 \text{ mm}^3$ with a 0.3 mm inter-slice gap). These parameters were implemented to optimize coverage of the ventral prefrontal cortex without sacrificing whole-brain acquisition. The first four EPI volumes were discarded to allow scans to reach equilibrium. One high-resolution T1-weighted SPGR structural image (TR, 2.3; TE, 2.0; FOV, 256; matrix, 192×192 ; sagittal plane; slice thickness, 1 mm; 172 slices) was collected on each participant for the anatomical co-registration with functional imaging datasets.

Data Analysis

Computational Modeling—We fit participants' behavioral data with a HYBRID state-based RL model, a win-stay-lose-shift (WSLS) model, and a baseline model that assumed stochastic responding. The RL model assumed that participants observed the hidden state (s) on each trial and valued options based on both the probability of reaching a given state on

the next trial (s') by selecting action a , and on the rewards experienced in each state. This model is similar to other models that have assumed that subjects use state-based information to determine behavior (Glascher et al., 2010; Eppinger et al., 2013), and identical to one we used in a previous behavioral study (Worthy et al., 2014).

The model tracks the *reward-based* expected values for each action in each state ($Q_{RB}(s,a)$) using a SARSA learner. On each trial the model computes the reward prediction error (RPE):

$$\delta_{RPE} = r - Q_{MF}(s, a) \quad (1)$$

The prediction error is then used to update the expected value for the current state action pair:

$$Q_{RB}(s, a) = Q_{RB}(s, a) + \alpha \delta_{RPE} \quad (2)$$

Here α is a free parameter that represents the learning rate for state-action pairs on each trial. The model also has the ability to allow reward information gained for actions in a specific state to be generalized across states. For each state other than the state on the current trial (denoted as s^*) the Q_{MF} value for the same action selected on the current trial is updated:

$$Q_{MF}(s^*, a) = Q_{MF}(s^*, a) + \theta (r - Q_{MF}(s^*, a)) \quad (3)$$

Here θ represents the degree to which the rewards received on each trial are generalized to the same action in different states.

To compute the *state-based* expected value of each option the model first tracks the state-transition probabilities associated with each action in each state. Following each trial in state s and arriving in state s' after having taken action a the model computes a state prediction error (SPE):

$$\delta_{SPE} = 1 - T(s, a, s') \quad (4)$$

Next, the model updates the state transition probability:

$$T(s, a, s') = T(s, a, s') + \eta \delta_{SPE} \quad (5)$$

Here η is a free parameter that controls the learning rate for the state-transition probabilities. The state-transition probabilities for all other states not arrived at (denoted as s'') are reduced according to:

$$T(s, a, s'') = T(s, a, s'') \cdot (1 - \eta) \quad (6)$$

This ensures that all transition probabilities at a given state sum to 1.

After updating state-transition probabilities and expected reward value information the model computes a state-based value for each action in each state ($Q_{SB}(s,a)$) using the state-transition probabilities and the maximum estimated model-free reward value in each possible future state to determine the future value of each action¹. In this task there are three possible states that participants will transition to on the next trial (s') following action a on the current trial (stay in the same state or move up or down one state). We estimated the Q_{SB} value for each action by the following equation:

$$Q_{SB}(s,a) = \sum_{s' \in \{s-1, s, s+1\}} T(s,a,s') * \max[Q_{MB}(s',a')] \quad (7)$$

This function multiplies the probability of transitioning to each possible state on the next trial, having taken action a in trial t , by the maximum expected reward in state s' for either action.

The model then determines a net value for each action ($Q_{Net}(s,a)$) by taking a weighted average of the model-based and model-free expected values:

$$Q_{Net}(s,a) = \omega \cdot Q_{SB}(s,a) + (1 - \omega) \cdot Q_{RB}(s,a) \quad (8)$$

Where ω is a free parameter that determines the degree to which choices are based on the state-based versus reward-based components of the model.

Finally, the probability of selecting each action is determined using the Softmax rule:

$$P(a,t) = \frac{e^{\beta \cdot [Q_{Net}(s,a) + \pi \cdot \text{rep}(a)]}}{\sum_{j=1}^n e^{\beta \cdot [Q_{Net}(s,j) + \pi \cdot \text{rep}(j)]}} \quad (9)$$

The state-change uncertainty signal that we used in our fMRI analysis was computed for each trial by first comparing the current state to the state on the previous trial:

$$SC(t) = s_t - s_{t-1} \quad (10)$$

This was simply equal to -1, 0, or 1, depending on whether there was a decline, no change, or an improvement in state after the previous action. Next, state changes were multiplied by the state-based prediction error from Equation 1 above to compute the state-change uncertainty signal (SCU):

$$SCU(t) = SC(t) \cdot \delta_{SPE} \quad (11)$$

This signal represents both the direction of the state change experienced from trial $t-1$ to trial t as well as the uncertainty associated with the state-change. Greater magnitude uncertainty signals should be more informative in updating state-transition probabilities. Note that this metric is similar to the state prediction error used by Glascher and colleagues (Gläscher,

Daw, Dayan, & O'Doherty, 2010), but it is a signed version of this metric where the change in state is multiplied by the state prediction error to represent not only uncertainty regarding the state transition, but whether the state transition was positive, neutral, or negative.

Finally, to compare the relative long-term versus immediate value of each action we simply computed the difference between the state-based and reward-based values for the chosen option on each trial.

$$\text{Diff}(Q) = Q_{\text{SB}}(s, a) - Q_{\text{RB}}(s, a) \quad (12)$$

This metric was chosen because it allowed us to account for the difference between the long-term value provided by the state-based Q-value and the immediate reward value provided by the reward-based Q-value. It is also a metric that is not affected by the state participants are in. Using *Qnet* values, for example, would be problematic because those values will be larger when participants are in a higher state and smaller when participants are in a lower state. This would reduce such a regressor's ability to track differences in value between the two options on any trial, regardless of the state participants are in.

The model was fit to the data individually for each participant. State-change uncertainty signals (*SCU*) and relative expected values (*Diff(Q)*) were computed for each subject by using the average best-fitting parameter values across all subjects. In addition, we also fit a variant of the model that did not include a perseveration term (*rep(a)*). This model was fit in order to compare its fits with the full model and with the win-stay-lose-shift model which makes predictions regarding perseveration and switching, but does not incorporate information about expected value.

The WSLS model is identical to the model we used in previous papers from our labs (Worthy et al., 2014; Worthy & Maddox, 2012). The model has two free parameters. The first parameter represents the probability of staying with the same option on the next trial if the reward received on the current trial is equal to or greater than the reward received on the previous trial:

$$P(a_i, t | \text{choice}_{t-1} = a_i \& r(t-1) \geq r(t-2)) = P(\text{stay} | \text{win}). \quad (13)$$

In Equation 1 *r* represents the reward received on a given trial. The probability of switching to another option following a win trial is $1 - P(\text{stay} | \text{win})$.

The second parameter represents the probability of shifting to the other option on the next trial if the reward received on the current trial is less than the reward received on the previous trial:

$$P(a_j, t | \text{choice}_{t-1} = a_i \& r(t-1) < r(t-2)) = P(\text{shift} | \text{loss}). \quad (14)$$

MRI Analysis—FMRI data processing was carried out using FEAT (FMRI Expert Analysis Tool) Version 6.00, part of FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). BET (brain extraction tool) was used to extract the brain from the image of the skull and surrounding tissue.

Motion correction was performed with MCFLIRT, using a normalized correlation ratio cost function and linear interpolation. Registration was conducted through a 2-step procedure, whereby the mean EPI image was first registered to the MPAGE structural image, and then into standard space [Montreal Neurological Institute (MNI); avg152 template], using a correlation ratio cost function and 12-parameter affine transformation (Jenkinson & Smith, 2001). This process was performed separately for each imaging run. Statistical analyses were performed in native space, with the statistical maps normalized to standard space prior to higher-level analysis. Whole-brain statistical analysis was performed using a three-level multi-stage approach to implement a mixed-effects model treating participants as a random effect.

The first-level regressors of interest were convolved using trial onset times with a canonical (double-gamma) hemodynamic response function as well as their temporal derivative. Additional nuisance variables included a stimulus presentation regressor that was parametrically modulated by response time (RT) and single volume regressors to “scrub” volumes exceeding a framewise displacement threshold of 0.9 mm (Siegel et al., 2014). A quasi-Poisson regression revealed no differences between groups in terms of number of scrubbed volumes ($t < 1$; median younger adults = 2; median older adults = 1). Time-series statistical analysis was carried out using generalized least squares in FILM (FMRIB's Improved Linear Model) with local autocorrelation correction (Woolrich, Ripley, Brady, & Smith, 2001) after pre-statistics processing including spatial smoothing using a full-width-half-maximum Gaussian kernel of 6mm to reduce noise and high pass temporal filtering (Gaussian-weighted least-squares straight line fitting, with $\sigma=100.0s$).

We first examined differences in task-related activation between younger and older adults, as well as task-related activation that was modulated by response time. In addition, we also conducted an analysis where state-change uncertainty signals (SCUs) and the difference between the state-based and reward-based expected value of the chosen option ($Diff(Q)$) from the HYBRID RL model presented above were mean-centered and included as parametric modulators of the BOLD response on a trial-by-trial basis in the first-level analysis. The purpose of mean centering is so the interpretation of the unmodulated regressors remains, “the mean activation during the stimulus” as opposed to “mean activation during the stimulus when parametric modulators have a value of 0”, which is less interpretable. Mean centering does not actually impact the interpretation of the modulated regressors

First-level, within-run parameter estimates were averaged across the four functional runs in a second-level, within-participant fixed effects model. Quality control checks at the second-level found ($n=4$, 2 Younger) participants with scan artifacts where run 2 needed to be excluded and one older adult where run 1 needed to be excluded. The contrast maps from the second-level were combined into a third, group-level mixed model that included

participant as a random effect. The third-level mixed model employed a two-stage Bayesian estimation for the variance components, which included full Markov-chain Monte Carlo estimation for all near threshold voxels (FSL's FLAME 1&2). Older adults, younger adults, and older adults versus younger adult analyses were conducted. Final whole-brain statistical maps were corrected for cluster extent at the $p < 0.05$ level using FSL's Gaussian random field theory-based correction and a cluster-forming threshold (primary uncorrected threshold) of $z > 1.96$. The minimum cluster size to achieve $p < 0.05$ for Gaussian random field theory corrections depends upon both the primary uncorrected threshold and the observed spatial smoothness of the data for a given third-level model. The observed spatial smoothness for each Group-level contrast are as follows (FWHM x,y,z, in voxels): SCUs = 4.72, 4.82, 4.90; $\text{Diff}(Q)$ = 4.90, 4.91, 5.12; Reaction Time = 5.26, 5.36, 5.41; Stimulus Onset = 5.64, 5.73, 5.90. The resulting minimum cluster sizes were: SCUs = 861; $\text{Diff}(Q)$ = 932; Reaction Time = 1105; Stimulus Onset = 1314.

In addition to the whole brain results, we conducted an ROI analysis on the right dorsolateral prefrontal cortex using an anatomically defined mask of the middle frontal gyrus from the Harvard-Oxford Atlas. The purpose of the ROI analysis was to examine whether individual differences in task performance correlated with the degree to which subjects' DLPFCs tracked the model-based SCU regressor. We also conducted similar analysis in the ventral striatum for both the SCU and DiffEV regressors.

Results

Behavioral Results

Figure 2 shows the average proportion of trials that participants selected the optimal option (the state-maximizing option) in each of the four 75-trial runs. From a mixed (run X age group) ANOVA we observed a significant effect of run, $F(3,105)=7.52$, $p<.001$, partial $\eta^2=.177$. Participants learned to select the optimal choice more as the task progressed. There was no effect of age, and no age X run interaction (both $F<1$). Thus, older adults selected the optimal option at the same rate as younger adults. We also examined average response times. Younger adults ($M=481\text{ms}$, $SD=77\text{ms}$) responded significantly faster than older adults ($M=608\text{ms}$, $SD=148\text{ms}$), $t(35)=-3.25$, $p<.01$.

Modeling Results

We fit each model to each individual data set by maximizing log-likelihood, and used Akaike's Information Criterion (AIC; (Akaike, 1974)) to compare the fit of each model. Table 1 lists the average AIC values and best-fitting parameter values for younger and older adults for each model. The WSLS and HYBRID-RL models fit the data roughly equally well, with the RL model fitting younger adults' data slightly better and the WSLS model fitting older adults' data slightly better. The inclusion of the perseveration term markedly improved the fit of the RL model as can be seen by comparing the AIC values of the model with and without the perseveration term. Accounting for participants' tendencies to perseverate can drastically improve fits of RL models and is the main reason why the WSLS model can fit data from a variety of tasks as well as RL models (Worthy, Hawthorne, & Otto, 2013; Worthy, Pang, & Byrne, 2013).

Following previous work from our lab we examined the degree to which the WSLs model fit the data better than the RL model that did not include a perseveration term by subtracting the AIC of the WSLs model from the AIC of the RL model with no perseveration term. This is analogous to the comparison we performed in our 2012 paper where we found greater reliance on the heuristic-based WSLs strategy in older adults, and greater reliance on an expected value based RL strategy in younger adults (Worthy & Maddox, 2012). Although the difference was not statistically significant, the relative fit values for the WSLs model were larger for older adults ($M=63.96$, $SD=74.48$) than for younger adults ($M=55.69$, $SD=79.38$). This is consistent with other studies where we have found that younger adults are better fit by models that rely exclusively on expected value, while older adults are better fit by models that rely exclusively on whether the most recent outcome was an improvement or a decline in reward.

Imaging Results

Task-Related Activation—We first examined age differences in task-related activation. The top of Table 2 lists regions where activation was significantly greater for younger adults than for older adults. Younger adults showed greater activation in regions of the ventral striatum, including caudate, putamen, and accumbens, and in the anterior cingulate and frontal pole while older adults show greater deactivation of the medial orbitofrontal cortex. Younger adults also showed greater activation in occipital regions.

RT-Modulated Activation—We performed similar analyses for task related activation modulated by RTs. We did not find greater activation for older adults in any brain regions. Table 2 lists the regions where RT-modulated activation was greater for older adults than for younger adults. There was greater activation for younger adults in temporal and occipital regions as well as in the thalamus and superior frontal gyrus.

State-Change Uncertainty-Related Activation—Next we examined activation that was parametrically modulated by the *SCU* measure from the HYBRID RL model, while adjusting for $Diff(Q)$. We first examined activation relative to baseline in both younger and older adults (Figure 3). In younger adults *SCUs* were correlated with activation in a few regions of the ventral striatum, as well as the amygdala. In older adults there were a variety of regions where activation was parametrically modulated by *SCUs*, most notably in R. DLPFC, as well as ventral striatal regions and the frontal pole, in addition to occipital regions.

When contrasting *SCU*-related activation between older and younger adults, older adults showed greater activation than younger adults in the R. DLPFC and precentral motor regions, and in the frontal pole and occipital regions. There were no regions where *SCU*-related activation was greater for younger adults than for older adults. The greater *SCU*-related activation we observed in older adults in the R.DLPFC was a key prediction we made prior to the study. This region has been associated with compensatory activation in older adults in prior studies and may be implicated in comparing one's current state to the state from the last trial. While this activation differs from a common pattern of frontal compensation where older adults recruit bilateral regions in the same area that younger

adults recruit unilaterally, another test of compensatory activation is whether increased activation in a brain region is associated with better performance (Lighthall et al., 2014).

To test this we performed a region of interest (ROI) analysis in R. DLPFC and examined the association between *SCU*-related activation and the proportion of optimal choices in the task. This ROI was determined a priori, and the mean contrast of the *SCU* parameter estimate across R. DLPFC was used as our measure of activation. Figure 4 plots this relationship for both younger and older adults. The correlation between *SCU*-related activation in R.DLPFC and performance was significant in older adults ($r = 0.655$, $p < .01$), but not in younger adults ($r = 0.17$, $p = .49$), however the was not significant across the whole sample, $t(31) = 1.64$, $p = .11$, and the interaction did not reach significance, $t(1.69)$, $p = .10$.

Although exploratory, we also conducted an ROI analysis in the ventral striatum to examine whether there was a correlation between *SCU*-related activation and performance in either younger or older adults. These relationships are plotted in Figure 5. Despite no age differences in *SCU* related activation in this region we found a positive correlation between performance and *SCU*-related activation in this region within younger adults ($r = 0.53$, $p < .05$), but not for older adults ($r = 0.07$, $p = .8$). The slope was not significant across the whole sample, $t(31) = 1.64$, $p = .11$ and the interaction did not reach significance, $t(29) = 1.75$, $p = .09$.

Expected Value-Related Activation—We also examined activation that was parametrically modulated by *Diff(Q)*, defined as the difference between the state-based (long-term) and reward-based (immediate) value of each option, while controlling for *SCU* (Figure 6). For the comparison to baseline there were a number of regions where activation was modulated by *Diff(Q)* within younger adults, including areas of the ventral striatum, medial and lateral PFC, as well as occipital regions. In contrast, for older adults we only found a region in the R. precentral gyrus where activation was modulated by *Diff(Q)*, and there were no regions where *Diff(Q)*-related activation was greater for older adults compared to younger adults. Younger adults showed greater *Diff(Q)*-related activation than older adults in the left putamen, as well as in temporal and occipital regions.

We also conducted an exploratory ROI analysis in the ventral striatum to examine whether *Diff(Q)* related activation was correlated with performance. These plots are shown in Figure 7. Interestingly we found negative associations between *Diff(Q)* related activation and performance for both younger ($r = -0.62$, $p < .01$) and older adults ($r = -0.66$, $p < .01$). While this relationship was unexpected and these analyses were exploratory in nature, one possibility is that the enhanced *Diff(Q)* related activation found in low performing older and younger adults is due to enhanced activation related to overriding the tendency to maximize immediate reward in favor of the larger delayed rewards.

Discussion

We examined BOLD activation in older and younger adults while they were performing a state-based decision-making task that required attention to both the value of each option and to how selecting each option caused changes in future states. While younger and older adults performed the task equally well our model-based fMRI analyses suggest that these two

groups differed in the information that was associated with activation in specific brain regions during the task. Within older adults, activation in several regions was parametrically modulated by state-change uncertainty signals that we derived from a HYBRID RL model. These signals reflect whether participants experienced an improvement or a decline in their state, relative to the previous trial, as well as how uncertain that state-change was. Older adults recruited R. DLPFC in addition to striatal regions in order to track these state changes. State-change related activation was significantly greater for older adults compared to younger adults in R. DLPFC, and our ROI analysis suggests that greater state-change weighted activation in this region was associated with superior performance in the task. These results are in line with the frontal compensation hypothesis of aging (Cabeza et al., 2002; Lighthall et al., 2014; Park & Reuter-Lorenz, 2009), although it should be noted that our results differ from prior work in that we did not observe bilateral recruitment of DLPFC. However, many prior studies that have found evidence of age-related frontal compensation in DLPFC have used basic memory tasks, and the pattern of age-related compensation may differ for complex decision-making tasks.

We also observed a positive association within younger adults between performance and state-change related activation in the ventral striatum. This is interesting given the lack of a relationship between performance and R. DLPFC activation in younger adults. One possibility is that is heightened activation in the ventral striatum for high-performing younger adults represents processing of the state-based prediction error, which aids in learning of the state-based structure of the task and in the development of state-based expected values. Activity in the ventral striatum has been associated with both reward-based and state-based prediction errors in prior work (Daw, Gershman, Seymour, Dayan, & Dolan, 2011), and our findings are in line with these results.

BOLD activation in older adults did not appear to be heavily modulated by a regressor that estimated the difference between the state-based and the reward-based expected value for the option selected on each trial. In contrast, we observed expected-value-related activation in younger adults in several regions that have consistently been implicated in reward-based decision-making and reinforcement learning (Rangel et al., 2008; Samanez-Larkin et al., 2014). However, it's important to point out that many of the regions where we found greater expected-value related activation in younger adults were in parietal and posterior brain regions not typically implicated in reward-based decision-making. Future work should address whether these regions are commonly recruited in decision-making tasks, in addition to frontostriatal regions often implicated in reward-based decision-making.

Overall, the results support our hypothesis that older adults are more reactive decision-makers who base their decisions on responses to recent events, rather than on a prospective comparison of expected values associated with each option at the time of choice. This assertion is supported by recent work in our labs that suggests that older adults are highly responsive to recent events, and by the current results that suggests that older adults are responsive to the most recent changes in their state (Worthy & Maddox, 2012; Worthy et al., 2015). It is also supported by other neuroimaging studies that show greater activation during reward receipt for older adults and greater activation at the time of choice for younger adults (Matthijs Vink et al., 2015).

Our exploratory ROI analysis of expected value related activation in the ventral striatum revealed surprising negative associations between activation and performance for both younger and older adults. While this relationship was unexpected and these analyses were exploratory in nature, one possibility is that the enhanced expected value related activation found in low performing older and younger adults is due to enhanced activation due to conflict over the immediate and delayed reward values associated with each option. The greater expected value related activation for poorer performing younger and older adults could be due to greater conflict in forgoing the larger immediate rewards provided by the reward-based option in favor of the larger delayed rewards provided by the state-based option. While speculative this type of conflict between immediate and delayed reward could have also caused the greater expected value related activation in younger adults compared to older adults, as older adults have shown reduced delay discounting compared to younger adults (Green, Fry, & Myerson, 1994).

Alternatively, our relative expected value metric was generally larger for the state-maximizing option than for the reward maximizing option. Thus, values were generally higher on this metric when individuals selected the state-maximizing option and lower on trials where individuals selected the reward-maximizing option. The inverse relationship between ventral striatum activation related to this regressor and performance could be due to higher performing participants having clearer knowledge that selecting the reward-maximizing option led to larger immediate reward, while selecting the state maximizing option led to smaller reward. Thus greater activation in high-performing individuals on trials where they selected the reward-maximizing option could be due to anticipation of the larger immediate rewards provided by this option (M. Vink, Pas, Bijleveld, Custers, & Gladwin, 2013). However, we urge caution in interpreting these correlations as we did not have specific a priori predictions about the relationship between expected value related activation in the ventral striatum and performance, like we did for *SCU* related activation in DLPFC for older adults.

It's important to note that activation for younger and older adults was modulated more for certain model-based regressors than for others. The use of multiple measures from models may be particularly necessary in aging studies. Given well-documented differences in strategy use it is important to ensure that an appropriately large model-space is utilized in order to account for both older and younger adult behavior. It's also important to note that we observed differences in age-related brain activity and strategy use despite no differences in performance in the task. Other researchers have found similar results and have pointed to older adults compensating for age-related declines in fluid intelligence with enhanced crystallized intelligence learned over the lifespan (Y. Li, Baldassi, Johnson, & Weber, 2013). Our results may reflect a similar situation where older adults compensate for declines in the ability to use expected value representations to proactively guide decision-making, and instead rely on reactive decision-making strategies where the relative value of recent outcomes strongly affects the subjective value of each option on future trials.

The notion that older adults depart from standard reinforcement-learning based strategies of updating and comparing expected values is in line with other recent work that demonstrates that neural responses to reward prediction errors decline with healthy aging (Chowdhury et

al., 2013; Eppinger, Schuck, Nystrom, & Cohen, 2013; Samanez-Larkin et al., 2014). The decline in neural responses to reward prediction errors is likely due to age-related decline in the integrity of the mesolimbic dopaminergic reward system (S.-C. Li et al., 2001). These age-related neural declines likely compromise older adults' ability to update expected value representations in response to prediction errors, which makes reinforcement-learning based strategies very inefficient. As a result, we propose that older adults engage in more parsimonious, reactive-based strategies akin to win-stay-lose-shift (Cooper, Worthy, Gorlick, & Maddox, 2013; Worthy, Hawthorne, et al., 2013; Worthy & Maddox, 2012). These strategies utilize recent changes in rewards or states as guides for how to alter or persist in the same decision-making behavior across trials.

Future work should further test our assertion that healthy aging causes qualitative shifts in the strategies individuals use to make decisions. Older adults may be particularly prone to poorer decision-making in situations where actual expected values must be utilized and compared, and where reactive, heuristic-based strategies are not viable. One such situation may be when additional decision alternatives are added (Frey, Mata, & Hertwig, 2015; Worthy et al., 2014). Older adults perform more poorly when additional choices are added, and this could be because heuristic-based strategies such as win-stay-lose-shift are less effective as more alternatives are added. Older adults may also underperform when asked to provide expected value predictions or when asked to “pass” or “play” a pre-selected choice (Cauffman et al., 2010). Those types of situations would require accurate expected value representations and would not allow older adults to rely on recency-biased heuristics. Understanding the qualitative differences in the types of decision-making strategies older and younger adults utilize is critical for developing ways to assist older adults in important personal, work-related and financial decisions, and for determining when decision-making may be compromised by age. Such an understanding is becoming increasingly important as older adults continue to make up a larger segment of the population in developed nations, and continue to make many important decisions on a daily basis.

Limitations

There are a few limitations to note for this study. First, we did not explicitly assess possible differences in IQ and we only conducted cognitive testing for our older adult sample. We also did not specifically screen for mental or physical illnesses other than ones that would preclude participation in fMRI studies. This study also focused only on older and younger adults and did not include data from participants in the midsection of the adult lifespan. This makes it difficult to determine precisely when the changes we observed occurred, particularly whether the shift from the use of proactive to reactive decision-making strategies steadily occurs over the course of the adult's lifespan or whether the shift occurs abruptly near the end of middle age. Park and colleagues have presented a compelling data set that suggests that visuospatial and verbal memory abilities steadily decline over the course of the lifespan (Park et al., 2002). Thus, one possibility is that the ability to utilize past information to update expected value representations that are utilized at the time of choice steadily declines over the course of the adult lifespan. Future studies should include broader samples of participants across the lifespan.

Conclusion

We used a combination of behavioral, computational modeling, and fMRI methods to test our hypothesis that older adults are more reactive decision makers than younger adults, whose behavior is better explained by expected value theory. BOLD activation in older adults more closely aligned with a regressor that tracked recent changes in state, while activation in younger adults was aligned with a regressor that tracked the difference between long-term and immediate expected value. Future work should continue to examine whether older and younger adults exhibit qualitative differences in how they utilize information about rewards and states to make decisions. A strong characterization of these differences will be extremely useful in predicting and improving decision-making behavior across the lifespan.

Acknowledgments

This work was supported by NIA grant AG043425 to D.A.W. and W.T.M. We thank the MaddoxLab RAs for all data collection.

References

- Akaike H. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*. 1974; 19(6):716–723.
- Bäckman L, Nyberg L, Lindenberger U, Li SC, Farde L. The correlative triad among aging, dopamine, and cognition: current status and future prospects. *Neuroscience & Biobehavioral Reviews*. 2006; 30(6):791–807. [PubMed: 16901542]
- Cabeza R, Anderson ND, Locantore JK, McIntosh AR. Aging gracefully: compensatory brain activity in high-performing older adults. *Neuroimage*. 2002; 17(3):1394–1402. [PubMed: 12414279]
- Castel AD, Rossi AD, McGillivray S. Beliefs about the “hot hand” in basketball across the adult life span. *Psychology and aging*. 2012; 27(3):601. [PubMed: 22288426]
- Cauffman E, Shulman EP, Steinberg L, Claus E, Banich MT, Graham S, Woolard J. Age differences in affective decision making as indexed by performance on the Iowa Gambling Task. *Developmental psychology*. 2010; 46(1):193. [PubMed: 20053017]
- Chowdhury R, Guitart-Masip M, Lambert C, Dayan P, Huys Q, Düzel E, Dolan RJ. Dopamine restores reward prediction errors in old age. *Nature neuroscience*. 2013; 16(5):648–653. [PubMed: 23525044]
- Cooper JA, Worthy DA, Gorlick MA, Maddox WT. Scaffolding Across the Lifespan in History-Dependent Decision Making. *Psychology and Aging*. 2013; 28(2):505–514. [PubMed: 23795765]
- Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ. Model-based influences on humans' choices and striatal prediction errors. *Neuron*. 2011; 69(6):1204–1215. [PubMed: 21435563]
- Dreher JC, Meyer-Lindenberg A, Kohn P, Berman KF. Age-related changes in midbrain dopaminergic regulation of the human reward system. *Proceedings of the National Academy of Sciences*. 2008; 105(39):15106–15111.
- Edwards W. The theory of decision making. *Psychological bulletin*. 1954; 51(4):380. [PubMed: 13177802]
- Eppinger B, Schuck NW, Nystrom LE, Cohen JD. Reduced striatal responses to reward prediction errors in older compared with younger adults. *The Journal of Neuroscience*. 2013; 33(24):9905–9912. [PubMed: 23761885]
- Eppinger B, Walter M, Heekeren HR, Li SC. Of goals and habits: age-related and individual differences in goal-directed decision-making. *Frontiers in neuroscience*. 2013; 7
- Frey R, Mata R, Hertwig R. The role of cognitive abilities in decisions from experience: Age differences emerge as a function of choice set size. *Cognition*. 2015; 142:60–80. [PubMed: 26022497]

- Gläscher J, Daw N, Dayan P, O'Doherty JP. States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*. 2010; 66(4):585–595. [PubMed: 20510862]
- Green L, Fry AF, Myerson J. Discounting of delayed rewards: A life-span comparison. *Psychological Science*. 1994; 5(1):33–36.
- Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A. Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *The Journal of Neuroscience*. 2008; 28(22):5623–5630. [PubMed: 18509023]
- Li SC, Lindenberger U, Sikström S. Aging cognition: from neuromodulation to representation. *Trends in cognitive sciences*. 2001; 5(11):479–486. [PubMed: 11684480]
- Li Y, Baldassi M, Johnson EJ, Weber EU. Complementary cognitive capabilities, economic decision making, and aging. *Psychology and aging*. 2013; 28(3):595. [PubMed: 24040999]
- Lighthall NR, Huettel SA, Cabeza R. Functional compensation in the ventromedial prefrontal cortex improves memory-dependent decisions in older adults. *The Journal of Neuroscience*. 2014; 34(47):15648–15657. [PubMed: 25411493]
- Maddox, WT.; Gorlick, MA.; Worthy, DA. Toward a three-factor motivation-learning framework in normal aging. In: Braver, T., editor. *Motivation and Cognitive Control*. United Kingdom: Taylor & Francis; 2015.
- Mata R, Josef AK, Lemaire P. Adaptive Decision Making and Aging. *Aging and Decision Making: Empirical and Applied Perspectives*. 2015; 105
- Mata R, Josef AK, Samanez-Larkin GR, Hertwig R. Age differences in risky choice: a meta-analysis. *Annals of the New York Academy of Sciences*. 2011; 1235(1):18–29. [PubMed: 22023565]
- Mather M. Aging and motivated cognition: The positivity effect in attention and memory. *Trends in cognitive sciences*. 2005; 9(10):496. [PubMed: 16154382]
- Mather M, Carstensen LL. Aging and motivated cognition: The positivity effect in attention and memory. *Trends in cognitive sciences*. 2005; 9(10):496–502. [PubMed: 16154382]
- Park DC, Lautenschlager G, Hedden T, Davidson NS, Smith AD, Smith PK. Models of visuospatial and verbal memory across the adult life span. *Psychology and aging*. 2002; 17(2):299. [PubMed: 12061414]
- Park DC, Reuter-Lorenz P. The adaptive brain: aging and neurocognitive scaffolding. *Annual Review of Psychology*. 2009; 60:173.
- Rangel A, Camerer C, Montague PR. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*. 2008; 9(7):545–556. [PubMed: 18545266]
- Reuter-Lorenz PA, Cappell KA. Neurocognitive aging and the compensation hypothesis. *Current directions in psychological science*. 2008; 17(3):177–182.
- Samanez-Larkin GR, Gibbs SEB, Khanna K, Nielsen L, Carstensen LL, Knutson B. Anticipation of monetary gain but not loss in healthy older adults. *Nature neuroscience*. 2007; 10(6):787–791. [PubMed: 17468751]
- Samanez-Larkin GR, Knutson B. Decision making in the ageing brain: changes in affective and motivational circuits. *Nature Reviews Neuroscience*. 2015
- Samanez-Larkin GR, Kuhnen CM, Yoo DJ, Knutson B. Variability in nucleus accumbens activity mediates age-related suboptimal financial risk taking. *The Journal of Neuroscience*. 2010; 30(4):1426–1434. [PubMed: 20107069]
- Samanez-Larkin GR, Levens SM, Perry LM, Dougherty RF, Knutson B. Frontostriatal white matter integrity mediates adult age differences in probabilistic reward learning. *The Journal of Neuroscience*. 2012; 32(15):5333–5337. [PubMed: 22496578]
- Samanez-Larkin GR, Worthy DA, Mata R, McClure SM, Knutson B. Adult age differences in frontostriatal representation of prediction error but not reward outcome. *Cognitive, Affective, & Behavioral Neuroscience*. 2014; 14(2):672–682.
- Schott BH, Niehaus L, Wittmann BC, Schütze H, Seidenbecher CI, Heinze HJ, Düzel E. Ageing and early-stage Parkinson's disease affect separable neural mechanisms of mesolimbic reward processing. *Brain*. 2007; 130(9):2412–2424. [PubMed: 17626038]
- Spaniol J, Bowen HJ, Wegier P, Grady C. Neural responses to monetary incentives in younger and older adults. *Brain research*. 2015; 1612:70–82. [PubMed: 25305570]

- Vink M, Kleerekooper I, van den Wildenberg WPM, Kahn RS. Impact of aging on frontostriatal reward processing. *Human brain mapping*. 2015; 36(6):2305–2317. [PubMed: 25704624]
- Vink M, Pas P, Bijleveld E, Custers R, Gladwin TE. Ventral striatum is related to within-subject learning performance. *Neuroscience*. 2013; 250:408–416. [PubMed: 23891754]
- Worthy DA, Cooper JA, Byrne KA, Gorlick MA, Maddox WT. State-based versus reward-based motivation in younger and older adults. *Cognitive, Affective, & Behavioral Neuroscience*. 2014; 14(4):1208–1220.
- Worthy DA, Gorlick MA, Pacheco JL, Schnyer DM, Maddox WT. With Age Comes Wisdom Decision Making in Younger and Older Adults. *Psychological Science*. 2011; 22(11):1375–1380. [PubMed: 21960248]
- Worthy DA, Hawthorne MJ, Otto AR. Heterogeneity of strategy use in the Iowa gambling task: A comparison of win-stay/lose-shift and reinforcement learning models. *Psychonomic bulletin & review*. 2013; 20:364–371. [PubMed: 23065763]
- Worthy DA, Maddox WT. Age-based differences in strategy use in choice tasks. *Frontiers in neuroscience*. 2012; 5
- Worthy DA, Otto AR, Doll BB, Byrne KA, Maddox WT. Older adults are highly responsive to recent events during decision-making. *Decision*. 2015; 2(1):27.
- Worthy DA, Pang B, Byrne KA. Decomposing the roles of perseveration and expected value representation in models of the Iowa gambling task. *Frontiers in psychology*. 2013; 4

Figure 1.

a.) Reward structure for the state-maximizing task. Rewards were a function of the number of times the state maximizing option had been selected over the previous ten trials. The reward maximizing reward led to larger immediate outcomes, but higher states and more cumulative reward could be achieved by selecting the state-maximizing option. **b.)** Sample screen shot of the experiment.

Proportion of Optimal Choices

Figure 2.
Proportion of state-maximizing choices during each 75-trial run.

Figure 3. State-change uncertainty (SCU) weighted activation for older adults (OA) and younger adults (YA).

Figure 4.

Scatterplot showing the relationship between the parameter for state-change uncertainty related activation in R. DLPFC and the proportion of state-maximizing choices within older adults. Greater state-change uncertainty related activation in this region was associated with better decision-making behavior in older adults.

Figure 5.

Scatterplot showing the relationship between the parameter for state-change uncertainty related activation in ventral striatum and the proportion of state-maximizing choices within older adults. Greater state-change uncertainty related activation in this region was associated with better decision-making behavior in younger adults.

Figure 6. Difference in expected value (*DiffQ*) weighted activation for older adults (OA) and younger adults (YA).

Figure 7.

Scatterplot showing the relationship between the parameter for difference in expected value related activation in ventral striatum and the proportion of state-maximizing choices within older adults. Greater difference in expected value related activation in this region was associated with poorer decision-making for both younger and older adults.

Table 1
Average AIC values and best-fitting parameters for each model

AIC Values		
	Younger	Older
HYBRID RL	267.48 (71.49)	299.56 (87.25)
HYBRID RL –No Perseveration	323.31 (98.75)	359.99 (83.61)
WSLS	267.62 (70.63)	296.03 (71.36)
Baseline	332.55 (84.02)	370.48 (75.31)
Parameter Estimates		
<u>HYBRID RL</u>		
	Younger	Older
State learning rate (η)	.26 (.37)	.41 (.40)
Reward learning rate (α)	.34 (.38)	.44 (.41)
Reward generalization rate (θ)	.10 (.25)	.07 (.23)
Model-based weight (ω)	.52 (.47)	.57 (.60)
Inverse temperature (β)	1.95 (2.14)	.87 (1.46)
Perseveration (π)	2.83 (6.39)	4.53 (5.58)
<u>HYBRID RL No Perseveration</u>		
	Younger	Older
State learning rate (η)	.29 (.42)	.45 (.45)
Reward learning rate (α)	.42 (.43)	.49 (.45)
Reward generalization rate (θ)	.05 (.08)	.05 (.08)
Model-based weight (ω)	.61 (.47)	.66 (.34)
Inverse temperature (β)	1.10 (1.65)	.46 (.71)
<u>WSLS</u>		
	Younger	Older
Win-stay probability	.81 (.19)	.82 (.09)
Lose-shift probability	.48 (.14)	.47 (.16)

Note: Standard deviations are listed in parentheses. Smaller AIC values indicate a better fit to the data.

Table 2
Brain activation patterns across older and younger adults

	Size (mm ³)	<i>x</i>	<i>y</i>	<i>z</i>
Task				
<i>YA > OA</i>				
Occipital Pole	2,628	18	-98	10
	10,150			
R. Thalamus		2	-2	4
L. Accumbens, L. Putamen		-12	8	-8
Anterior Cingulate		10	40	8
Frontal Pole, M. PFC		-2	56	6
R. Caudate		14	-2	18
RT				
<i>YA > OA</i>				
	35,324			
Temporal Pole		-28	10	-28
Temporal Fusiform Cortex		-30	-46	-22
Superior Frontal Gyrus		6	32	56
L. Thalamus		-8	-16	32
Inferior Temporal Gyrus		44	-16	-32
State-Change Uncertainty				
<i>YA</i>				
	1,660			
R. Caudate		34	0	18
<i>OA</i>				
F. Pole	1,926	-8	62	12
	29,234			
Lingual Gyrus		-2	-86	-8
Occipital Pole, Occipital Fusiform Gyrus		8	-88	40
Cuneal Cortex		4	-76	34
Caudate, Putamen, Thalamus		-10	-2	10
R. DLPFC		54	16	32
<i>OA > YA</i>				
	14,811			
M. Occipital Cortex, Lingual Gyrus		2	-86	-8
L. Occipital Cortex, Occipital Pole		6	-70	46
R. DLPFC, Precentral Gyrus, F. Pole		38	22	52
Difference in Expected Value				
<i>YA</i>				
R. Caudate, Thalamus, Posterior Cingulate	1,324	14	-18	24
L Lat. OFC, L Insula	1,724	-28	34	-2
R. Occipital Cortex, Angular Gyrus, Supramarginal Gyurs	2,018	40	-46	28

	Size (mm ³)	x	y	z
R. DLPFC, R. Lat. OFC	2,432	50	30	4
Anterior Cingulate Cortex, vmPFC	3,022	-2	54	18
L. Occipital Cortex, Angular Gyrus	3,419	-30	-70	28
Lingual Gyrus, Fusiform Gyrus	3,829	-2	-68	16
OA				
R. Precentral Gyrus	1,028	46	-6	58
YA>OA				
	8,350			
L. Lateral Occipital Cortex, Fusiform Gyrus		-30	-70	28
L. Putamen		-26	-24	6
L. Superior Temporal Gyrus		-60	-36	2
L Middle Temporal Gyrus				
Occipital Pole		-30	-70	28

Note: R=Right, L=Left, F=Frontal, Med.=Medial, Lat.=Lateral, YA= Younger adults, OA=Older adults