NeuroImage 130 (2016) 13–23

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Neural correlates of state-based decision-making in younger and older adults

Darrell A. Worthy [a,*], Tyler Davis [b], Marissa A. Gorlick [c], Jessica A. Cooper [c], Akram Bakkour [d], Jeanette A. Mumford [e], Russell A. Poldrack [f], W. Todd Maddox [c]

[a] Texas A&M University, USA
[b] Texas Tech University, USA
[c] University of Texas at Austin, USA
[d] Columbia University, USA
[e] University of Wisconsin Madison, USA
[f] Stanford University, USA

ARTICLE INFO

Article history:
Received 18 August 2015
Accepted 5 December 2015
Available online 12 December 2015

Keywords:
Aging
Decision-making
fMRI
Reinforcement learning
Computational modeling

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.

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 et al., 2008; Samanez-Larkin and 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 et al., 2014; Maddox et al., 2015; Mata et al., 2011; Mather and Carstensen, 2005; Samanez-Larkin and Knutson, 2015; Worthy et al., 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 et al., 2006; Chowdhury et al., 2013; Li et al., 2001). Another study demonstrated that decline in the white matter integrity of glutamatergic pathways may also impact reward-based decision-making (Samanez-Larkin et al., 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, 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 et al., 2012; Mata et al., 2015; Worthy et al., 2011, 2014; Worthy and 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

* Corresponding author at: Department of Psychology, Texas A&M University, 4235 TAMU, College Station, TX 77843-4235, USA.
E-mail address: worthyda@tamu.edu (D.A. Worthy).

http://dx.doi.org/10.1016/j.neuroimage.2015.12.004
1053-8119/

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 and 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 et al., 2013b). 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 (Vink et al., 2015). Several other studies have found similar results (Dreher et al., 2008; Samanez-Larkin et al., 2007, 2010; Schott et al., 2007; Spaniol et al., 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 (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., 2011, 2014). 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 et al., 2002; Park and Reuter-Lorenz, 2009; Reuter-Lorenz and 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 et al., 2008; Rangel et al., 2008; Samanez-Larkin et al., 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/h 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.

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 Fig. 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 option is chosen more frequently (move to the right on the x-axis in Fig. 1A), whereas the state declines as the reward-maximizing option is chosen more frequently (move to the left on the x-axes in Fig. 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 4 s where participants could select one of the two options for 2 s followed by a 2 s of feedback. After every 25 experimental trials there was a 16-second fixation trial.