

Reviews

The value of choice facilitates subsequent memory across development

Perri L. Katzman^a, Catherine A. Hartley^{a,b,*}^a New York University, Department of Psychology, 6 Washington Place, New York, NY 10003, United States^b New York University, Center for Neural Science, 4 Washington Place, New York, NY 10003, United States

ARTICLE INFO

Keywords:

Memory
Development
Agency
Learning

ABSTRACT

Both children and adults are more likely to remember information when they have control over their learning environment. Despite many demonstrations of this effect in the literature, it is still unclear how and why people are more likely to remember information that is obtained through their own actions rather than passively received. One possibility is that individuals are biased to remember the outcomes of their choices because doing so may often be beneficial. Having agency, or the ability to exert control, is valuable when individuals can act in an instrumental manner to achieve their goals. Preferentially encoding information encountered in such contexts may confer an advantage when making similar decisions in the future. However, it has not been directly examined whether modulating the value, or utility, of agency affects its mnemonic benefit. Additionally, the developmental trajectory of how the utility of agency affects memory is unclear. The current study examines whether the mnemonic benefit of agency is modulated by the utility of choice and whether this effect varies as a function of age. We tested 96 participants, ages 8 to 25, in a paradigm in which agency and its utility were separately manipulated at encoding. In contrast to previous studies, we did not find that simply having the ability to make a choice enhanced memory. Rather, when the utility of agency varied within the task, we only observed an agency-related memory benefit when the ability to choose had the greatest utility. This pattern was age-invariant, suggesting that this effect on memory is present in middle childhood and persists through adulthood.

1. Introduction

Why are we more likely to remember a route taken when we are driving rather than sitting in the passenger seat? Both the driver and passenger view the same visual input through the windshield and can make choices about what streets to take. However, only the driver can turn the steering wheel and control where the car will go. Agency, or the ability to exert control by causally influencing choice outcomes, has been shown to facilitate memory in adults (Gureckis & Markant, 2012; Voss, Galvan, & Gonsalves, 2011; Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011) and in children (Ruggeri, Markant, Gureckis, Bretzke, & Xu, 2019). Even relatively inconsequential control, like determining the timing of the onset of stimuli (Markant, DuBrow, Davachi, & Gureckis, 2014) or merely having the opportunity to make a choice (Murty, DuBrow, & Davachi, 2015) enhances subsequent memory for the resulting outcomes. Although a large corpus of work has demonstrated that agency facilitates memory across development, it is still unclear why individuals are more likely to remember the consequences of their own actions.

Agents may prioritize memory for controllable events because it is

adaptive to do so. This interpretation is in line with a rational analysis of cognition, which posits that individuals are more likely to remember information and events that will be useful in the future (Anderson & Milson, 1989; Steyvers & Griffiths, 2008). Given the costs associated with memory encoding, storage, and retrieval (Dukas, 1999), prioritizing information in memory according to its future utility optimizes use of one's cognitive resources (Anderson & Milson, 1989; Chater & Oaksford, 1999). Remembering the consequences of previous choices is adaptive because it can help an individual determine which action might lead to a desired outcome when faced with a similar decision later on. Thus, prioritization of memory for choice outcomes is a particularly effective way to aid future decisions in real-world environments in which reward statistics are learnable, rewards are contingent upon actions, and these relations exhibit some degree of temporal consistency. In such contexts, agency can be considered to have value, or utility, because individuals can make choices to bring about desired outcomes. While past studies in both children and adults have demonstrated that agency alone can facilitate memory, it is currently unclear whether memory is sensitive to the degree to which agency is useful, and how sensitivity to the utility of agency may develop with

* Corresponding author at: New York University, Department of Psychology, 6 Washington Place, New York, NY 10003, United States.

E-mail address: cah369@nyu.edu (C.A. Hartley).

age.

If the mnemonic benefit of agency is modulated by its utility, the effect may exhibit multiple possible developmental trajectories. One possibility is that the ability to prioritize memory based on its potential future value requires cognitive control processes that continue to develop into adulthood. The ability to adaptively modulate cognitive effort based on potential reward improves through late adolescence (Castel et al., 2011; Davidow, Insel, & Somerville, 2018; Hantén et al., 2007; Insel, Kastman, Glenn, & Somerville, 2017; Störmer, Eppinger, & Li, 2014). If prioritizing the encoding of choice outcomes is an effortful goal-directed process, then the ability to preferentially allocate this cognitive resource to contexts in which agency has utility would likely continue to develop into adulthood. Alternatively, agency has been proposed to function as an intrinsic reinforcer (Bobadilla-Suarez, Sunstein, & Sharot, 2017; Leotti & Delgado, 2011), which, in turn, may facilitate memory even in the absence of cognitive effort (Murty et al., 2015). As children exhibit behavioral and neural sensitivity to both intrinsic and extrinsic reinforcers (Galvan et al., 2006; Satterthwaite et al., 2012), such an account would suggest that their memory might be just as sensitive as adults' to contexts in which agency can be leveraged to achieve rewarding outcomes, and thus we would not expect any age-related effects. Consistent with such early emergence of agency effects on memory, children show better memory when they have control over their learning environment (Feldman & Acredolo, 1979; McComas, Dulberg, & Latter, 1997; Partridge, McGovern, Yung, & Kidd, 2015; Ruggeri et al., 2019; Sim, Tanner, Alpert, & Xu, 2015), with one study observing an effect size comparable to that of adults by approximately age eight (Ruggeri et al., 2019). However, neither the previous studies in children nor those with adult participants have manipulated the utility of agency in their paradigms.

The current study addresses two aims: to determine, first, whether the mnemonic benefit of agency is modulated by its utility and, second, whether this effect varies as a function of age. We designed a novel child-friendly paradigm in which agency and the utility of agency varied across three choice contexts. Participants, ages 8 to 25, were instructed to maximize the amount of reward they received during the task, and told they would earn bonus money based on the amount of reward they earned. On half of the trials, they had agency and could make a choice between two options, and on the remaining trials they lacked agency and one option was pre-selected. After selecting one of the options, participants saw a trial-unique item, learned whether or not it was rewarding, and were instructed to remember it for the next day's task. The utility of agency varied according to the amount of reward participants could earn through their choices. If memory for an event is enhanced by the utility of agency at encoding, we would expect that the memory benefit for items in the agency condition would be greater in the contexts for which having the ability to choose can result in greater reward. Alternatively, if simply having the ability to make a choice were beneficial for memory, then we would not expect any differences across our conditions. Our results support the first hypothesis. We found that participants were more likely to remember the items from the agency trials, but only in the context for which the ability to choose had the greatest utility. The magnitude of this effect did not change with age, suggesting that sensitivity to the utility of agency, and its resulting mnemonic benefit, is evident from middle childhood through adulthood.

2. Methods

2.1. Participants

One hundred ten participants ages 8 to 25 were recruited from the New York University and New York City communities for a behavioral study. Informed consent was obtained in a manner approved by the New York University Committee on Activities Involving Human Subjects. Adult participants provided informed consent; participants

under eighteen years of age assented to participate, and a parent or legal guardian provided consent. Participants were paid a baseline of \$30 for participation, plus a \$5 bonus for performance. Fourteen participants were excluded from analyses: five dropped out before completing the second day, eight were excluded due to an error in the task, and one for self-reporting not paying attention during the learning phase. Of the remaining 96 participants included in the analyses, there were 31 children (15 female, ages 8–12, $M = 10.6$, $SD = 1.6$), 31 adolescents (17 female, ages 13–17, $M = 15.6$, $SD = 1.5$), and 34 adults (16 female, ages 18–25, $M = 22.2$, $SD = 2.5$). Participants' self-identified race and ethnicity were as follows: 10% Black/African American, 45% Caucasian/White, 23% Asian, and 22% Mixed Race. Twenty-five percent identified as Hispanic/Latinx.

2.2. Materials

Memoranda comprised 240 color photos of household objects on white backgrounds. Stimuli were randomly selected from this collection for each participant such that 120 images were used in the learning phase and 60 were used as lures for the memory phase. The three choice contexts in the learning phase were represented as galaxies. Three photos of different galaxies, and the names "THETA", "GAMMA", and "OMEGA" were each randomly assigned to a choice context for every participant. Each galaxy was randomly assigned two planets from a collection of 6 cartoon planet images in different colors. The experimental script was coded using Psychtoolbox v3.0.15 with MATLAB v9.5.

2.3. Experimental design

The study took place over two consecutive days. All subjects completed the same version of the tasks with a child-friendly backstory of searching for treasure in space:

Aliens came to Earth in search of treasure. Though the objects they found seemed ordinary to us, the aliens thought that some were treasure and some were trash. On their way back to their home planet they lost their objects in three different galaxies. Each galaxy has two planets; planets may have mostly treasure, mostly trash, or have an equal amount of treasure and trash.

2.3.1. Learning phase

In each trial of the task (Fig. 1), participants first saw which of three galaxies they would be visiting (Context Cue), followed by whether the planet would be chosen by the Captain or the Autopilot (Agency Cue). On Captain trials, participants had agency and could choose which planet to visit. They were instructed to respond via keypress and could press 1 to visit the left planet or 2 to visit the right. On Autopilot trials, participants lacked agency. One of the planets was pre-selected and they had to press the corresponding key to visit. In both the Captain and Autopilot conditions, participants had no time limit to respond. Once they made a keypress, participants saw a trial-unique object and learned if the aliens thought it was trash or treasure (Outcome). Participants were forewarned to pay attention to and remember the objects for the next day's task. A 1-second fixation cross was displayed between each screen. Participants completed 120 trials total; 40 in each galaxy, with 20 Captain and 20 Autopilot trials each. Trial order was pseudo-randomized to ensure an equal number of each trial type throughout the task.

While agency was manipulated on a trial level, the utility of agency differed across the three galaxies. In the Non-Learnable galaxy, both planets had 50% reward probability (i.e., treasure outcome on half of the trials). Autopilot sampled both planets equally in a pseudo-randomized order. Due to the reward probabilities, both Captain and Autopilot would find treasure on roughly half of the trials regardless of the planet selected. Agency had no utility in this galaxy because

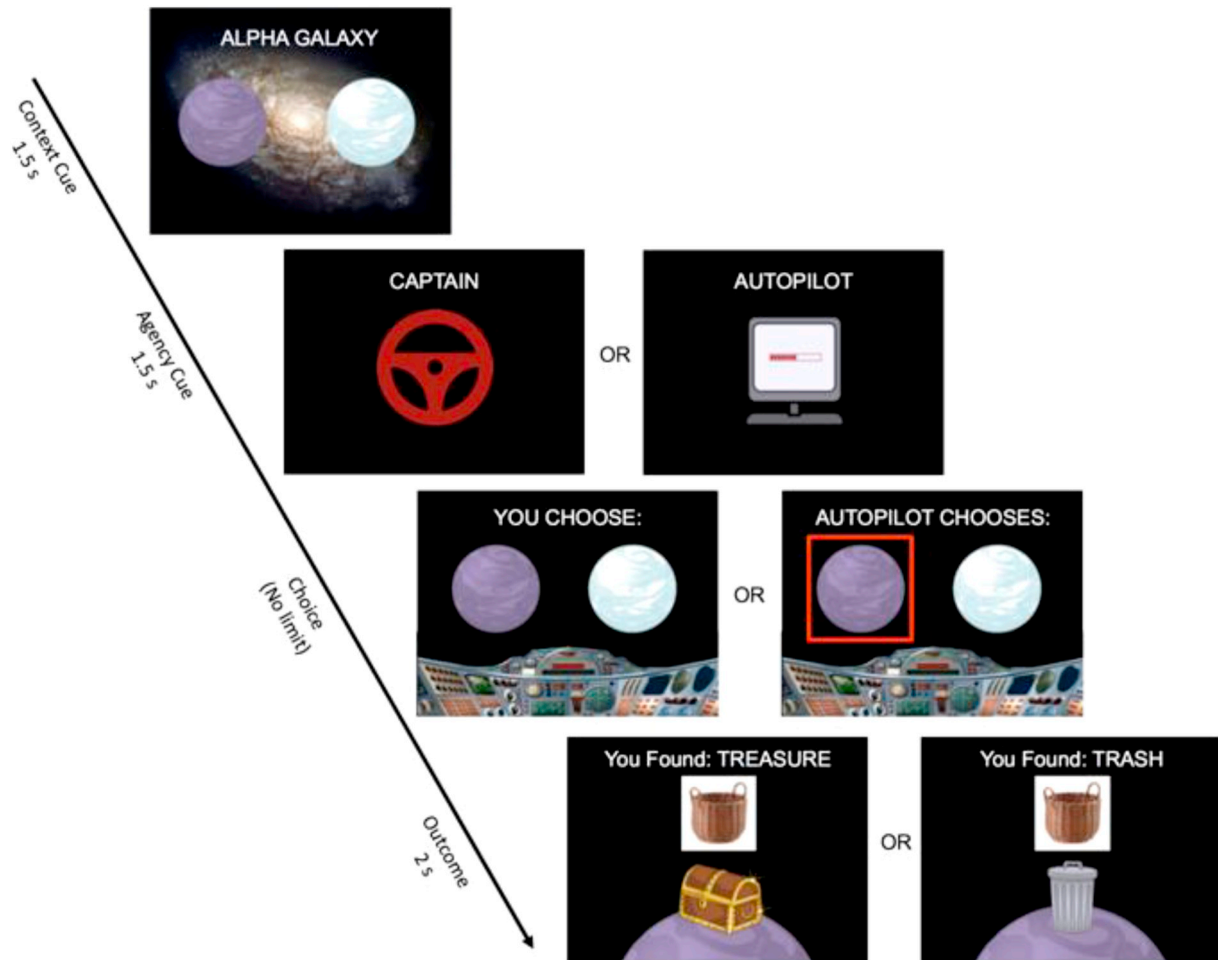


Fig. 1. Learning phase design. During the instructions and practice trials, participants viewed this galaxy name (ALPHA) and these two color planets. For the actual task, the galaxies had different names and planets were six different colors.

participants could not use their choices to achieve their goal of finding more treasure. The other two galaxies were “learnable”; both had one planet with 80% reward probability, and the other with 20%. If participants were able to learn that one planet had a higher reward probability, they could use that knowledge to choose optimally. In the Learnable-Yoked galaxy, Autopilot’s behavior was yoked to the participant’s previous choices in that galaxy. More specifically, every four Autopilot choices were the same as the participant’s previous four choices in a random order. For example, if the participant’s first four choices were “left”, “left”, “right”, “left”, then Autopilot’s next four choices may be “left”, “right”, “left”, “left.” The first four Autopilot trials sampled equally from both sides in a random order. As a result, having agency meant the participant could earn more treasure overall, but would not find substantially more treasure on the Captain trials compared to the Autopilot trials. In the Learnable-Dissociable galaxy, Autopilot selected both planets equally. Consequently, participants who learned to select the high-reward planet could earn more treasure on the Captain trials than on the Autopilot trials. Agency had the most utility in this galaxy.

2.3.2. Memory phase

Participants returned after 24 h to complete the memory test. On each trial, participants saw an object and judged whether they thought it was definitely new, maybe new, maybe old, or definitely old. We refer to this response as item recognition. When participants endorsed an object as old, they were then asked whether the object had been definitely or maybe treasure or trash. We refer to this response as source

memory discrimination. All decisions were made without a time limit. Participants made these judgments of all 120 items from the Learning Phase and 60 novel lure items, pseudo-randomized to ensure an equal ratio of old and new items for each fifth of the task. We included fewer lures than old objects in order to reduce the duration of the task, and previous studies have demonstrated that this does not affect hit rate (e.g., Cox & Dobbins, 2011).

2.3.3. Post-task questions

After participants completed the memory test, they answered a few questions about their experience during the Learning Phase. They ranked their preference for each galaxy, estimated the probability of finding treasure on each of the six planets (from Not Likely to Very Likely on a 7-point Likert scale), and rated how much control they felt they had over their ability to find treasure when they visited each galaxy as a Captain (from No Control to Full Control on a 7-point scale). Once they answered these questions, all participants were told they earned a \$5 bonus.

2.4. Analysis approach

Participant data and analysis code can be accessed on the Open Science Framework: <https://osf.io/n4kam/>. We used R version 3.6.0 (R Core Team, 2013) to conduct all data processing and statistical analyses. Linear models were fit using the *lme4* package (v1.1-21; Bates, Mächler, Bolker, & Walker, 2015). For all mixed-effects models, we included subject-level random intercepts (Baayen, Davidson, & Bates,

2008; Jaeger, 2008), unless otherwise noted. In constructing our models, we began with the maximal model in order to minimize Type I error (Barr, Levy, Scheepers, & Tily, 2013). We set the number of model iterations to 10,000 and used the ‘bobyqa’ optimizer. If the model failed to converge, we first removed correlations between random intercepts and random slopes, then interactions between random slopes, and finally removed random slopes themselves, until the reduced model converged. When comparing results between the three choice contexts, the Non-Learnable galaxy was treated as a baseline condition to which the Learnable-Dissociable and Learnable-Yoked galaxies were compared. We used the *aod* package (v1.3.1; Lesnoff & Lancelot, 2012) to compute the Wald X^2 test for statistically significant effects from linear models, and the *BayesFactor* package (v0.9.12-4.2; Morey & Rouder, 2018) to compute Bayes factor for all *t*-tests and ANOVAs.

A primary aim of this study was to determine whether there were age-related changes in how the utility of agency influenced memory. To determine whether any potential age effects followed a linear or quadratic pattern, we began each analysis by comparing whether the inclusion of a squared age regressor provided a significantly better fit to the data, as indexed by ANOVA tests and through comparison of Bayesian information criterion (BIC) scores (Raftery, 1995; Wagenmakers, 2007). Age was always z-scored when included as a regressor, and to maintain the relationship between age and squared age the latter was calculated as the square of z-scored age.

3. Results

3.1. Learning results

3.1.1. Learning phase accuracy

In this experiment, agency was only consequential if participants could learn to select the high-reward planets. To examine performance across the three galaxies, we fit a generalized linear mixed-effects logistic regression model to participants' choices (“correct” high-reward planet versus “incorrect” low-reward planet) on Captain trials during the learning phase. In the Non-Learnable galaxy, the “correct” choice was arbitrarily coded as choosing left. Galaxy, participant age, z-scored trial number, and all interactions were included as model predictors, with a random intercept and random slopes for Galaxy and trial number for each participant. An ANOVA comparing the linear and quadratic age models found the latter did not fit significantly better ($X^2 = 3.47$, $df = 6$, $p = .747$; $BIC_{\text{Linear}} = 6806.6$, $BIC_{\text{Quadratic}} = 6855.1$). In the linear age model, participants made more correct choices in the Learnable-Yoked galaxy ($X^2 = 37.7$, $df = 1$, $p < .001$) and the Learnable-Dissociable galaxy ($X^2 = 26.9$, $df = 1$, $p < .001$) compared to the Non-Learnable galaxy. There was a significant interaction with trial number, indicating that performance in both the Learnable-Yoked ($X^2 = 21.5$, $df = 1$, $p < .001$) and Learnable-Dissociable ($X^2 = 20.3$, $df = 1$, $p < .001$) galaxies improved over the course of the task compared to the Non-Learnable galaxy (Fig. 2; see Supplementary Fig. 1 for data plotted by age group). In order to compare performance in the two learnable galaxies, we re-ran the analysis with the Learnable-Yoked galaxy as the baseline. We found no significant difference between the Learnable-Dissociable and Learnable-Yoked galaxies ($X^2 = 0.08$, $df = 1$, $p = .78$). There were no significant effects of age, indicating that participants across our age range may learn to select the more rewarding planets equally well.

3.1.2. Learning phase reaction times

In addition to examining how choice accuracy improves throughout the learning phase, we also tested whether reaction times (RTs) varied by age and galaxy condition. We first log-transformed the RTs to make the data normally distributed. We then fit a linear mixed-effects model with galaxy, age, and z-scored trial number as predictors and a random slope for each participant. A model that included an additional squared age regressor did not provide a significantly better fit ($X^2 = 7.49$,

$df = 6$, $p = .285$; $BIC_{\text{Linear}} = 17,840$, $BIC_{\text{Quadratic}} = 17,884$). In the linear age model we found a significant negative interaction between trial number and the Learnable-Dissociable galaxy compared to the Non-Learnable galaxy ($X^2 = 4.8$, $df = 1$, $p = .028$). As the task progressed, participants responded relatively more quickly in the Learnable-Dissociable galaxy and relatively more slowly in the Non-Learnable galaxy. We also found a negative interaction between trial number and age ($X^2 = 5.8$, $df = 1$, $p = .016$), such that over the course of the task, children get slower, adolescents stay roughly consistent, and adults get faster. To assess any reaction time differences between the two learnable galaxies, we re-ran the linear mixed-effects model with the Learnable-Yoked galaxy as a baseline. We found no significant effects for the Learnable-Dissociable galaxy relative to the Learnable-Yoked galaxy.

3.2. Memory results

3.2.1. Recognition memory

One primary aim was to determine whether the utility of agency modulated subsequent memory. To test this, we first calculated each participant's hit rate (proportion of items *correctly* endorsed as “definitely” or “maybe” old in the memory phase) separately for items from Captain and Autopilot trials in each of the three galaxies. We subtracted each participant's false alarm rate (proportion of items *incorrectly* endorsed as “definitely” or “maybe” old in the memory phase) from those six hit rates to calculate corrected hit rates (cHR). We then fit a linear mixed-effects model with agency condition (Captain or Autopilot), galaxy (Non-Learnable, Learnable-Yoked, or Learnable-Dissociable), and age as regressors to predict cHR, with a random intercept and random slope for agency for each participant. A separate model that included an additional squared age regressor did not provide a significantly better fit ($X^2 = 4.33$, $df = 6$, $p = .632$; $BIC_{\text{Linear}} = -639.4$, $BIC_{\text{Quadratic}} = -605.6$). In the linear age model, we found a significant agency by galaxy interaction such that cHR was higher for the Captain (vs. Autopilot) condition in the Learnable-Dissociable galaxy compared to the Non-Learnable galaxy ($X^2 = 4.3$, $df = 1$, $p = .038$). No other main effects or interactions reached significance. A post-hoc paired *t*-test between the Captain and Autopilot conditions in the Learnable-Dissociable galaxy confirmed a significant effect of agency ($t(95) = 4.97$, $p < .001$; Bayes Factor (BF) = 5128.2).

The reason agency has the greatest utility in the Learnable-Dissociable galaxy is because participants could earn more treasure on the Captain trials than the Autopilot trials. Since people are more likely to remember high-reward information (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Castel et al., 2011), we wanted to rule out the possibility that the effect of agency was driven by the difference in treasure earned. For this analysis, we further subdivided the Captain and Autopilot trials in each galaxy by whether the item had been treasure or trash (Fig. 3). We then calculated cHR for each sub-category and fit a linear mixed-effects model with agency condition, galaxy, age, and reward (treasure or trash) as regressors to predict cHR, and included a random intercept and random slopes for agency and reward for each participant. For this analysis we again collapsed across “definitely” and “maybe” responses (see Supplemental information for exploratory analyses of confidence ratings). We found a main effect of reward such that participants had greater cHR for treasure than trash items overall ($X^2 = 5.2$, $df = 1$, $p = .023$). Critically, the interaction between galaxy and agency remained with greater cHR for Captain than Autopilot trials in the Learnable-Dissociable galaxy compared to the Non-Learnable galaxy ($X^2 = 4.9$, $df = 1$, $p = .026$; see Supplementary Fig. 2 for interaction plot). No other effects in this model reached significance (see Supplementary Fig. 3 for data plotted by age group). Post-hoc paired *t*-tests between the Captain and Autopilot cHR in the Learnable-Dissociable galaxy revealed a significant agency memory benefit for both the treasure ($t(95) = 2.49$, $p = .015$; BF = 2.09) and trash ($t(95) = 4.39$, $p < .001$; BF = 588.0) items. Thus, the positive

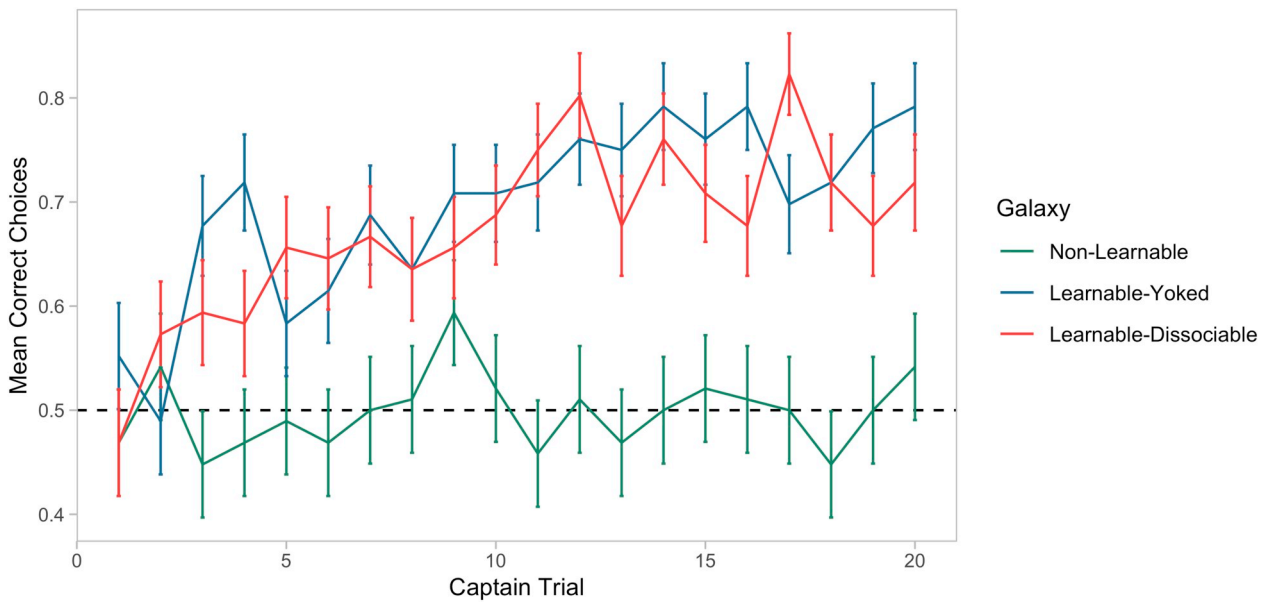


Fig. 2. Learning phase performance, averaged across all participants for each Captain trial for each galaxy. Chance performance denoted by dashed line. Error bars denote standard error.

effect of agency on memory in the Learnable-Dissociable galaxy cannot be attributed to there being more treasure for Captain trials than Autopilot trials. To ensure our effects were not influenced by participants performing at chance level, we re-ran the analysis excluding individuals for whom overall CHR was at or below zero and all significant effects still held (see Supplemental information). Additionally, the pattern of results is similar when d' is used as the dependent measure rather than CHR (see Supplemental information). To assess whether the effect of agency differed between the two learnable galaxies, we re-ran the linear mixed-effects model analysis with the Learnable-Yoked galaxy as the baseline. We found a significant agency by galaxy interaction such that memory was stronger for Captain (vs. Autopilot) trials in the Learnable-Dissociable galaxy relative to the Learnable-Yoked galaxy ($X^2 = 6.8$, $df = 1$, $p = .009$). This indicates that participants had better memory for information encoded with agency, but only in the context for which

agency had the greatest utility.

3.2.2. Source memory

We fit a generalized linear mixed-effects logistic regression to model source memory discrimination accuracy using reward, agency condition, galaxy, and age as predictors, with random intercepts for each memoranda item and each subject, and a subject-level random slope for reward. We only included studied items that were correctly remembered. According to an ANOVA analysis comparing the linear and quadratic age models, the inclusion of the squared age term provided a significantly better fit ($X^2 = 21.2$, $df = 12$, $p = .048$). In the quadratic age model there was a significant main effect of squared age ($X^2 = 4.3$, $df = 1$, $p = .038$), such that performance was slightly better in adolescents, relative to children and adults. Additionally, there was an interaction between reward and squared age ($X^2 = 9.8$, $df = 1$,

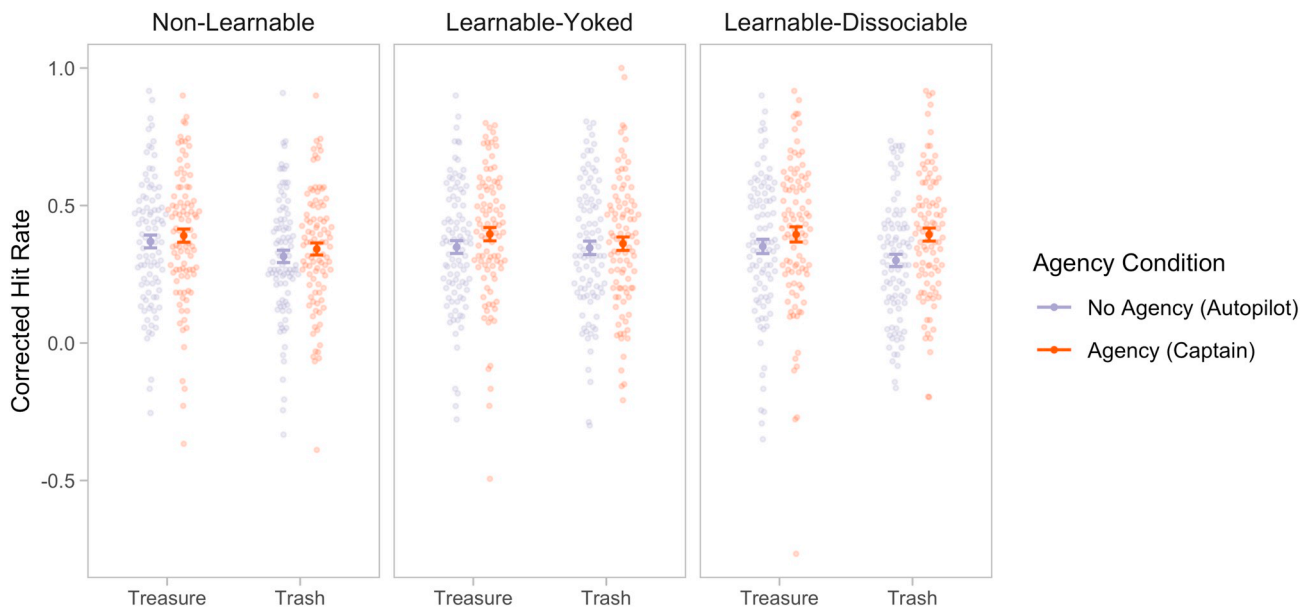


Fig. 3. Corrected hit rates for trash and treasure items in each galaxy. A corrected hit rate of 0 indicates an equal proportion of hits and false alarms for that condition. Positive values indicate better memory performance. Smaller dots represent individual subjects' corrected hit rate for the corresponding trial types, with the means represented as larger dots, and error bars reflecting standard error.

$p = .002$) such that children and adults had substantially better source memory accuracy for trash items than for treasure items, while adolescents had relatively equal source memory accuracy for both categories (see Supplementary Fig. 4). No other effects reached significance; there was no evidence to suggest that agency influenced source memory accuracy. However, the BIC value for the linear age model (8920.3) was lower than that of the quadratic model (9005.2), indicating the former provides a better fit. All statistically significant effects are reported here, though these findings should be interpreted in a degree of caution. In the linear age model, we found a significant negative effect of reward ($X^2 = 10.6$, $df = 1$, $p = .001$), reflecting overall worse source memory accuracy for treasure items, and a negative interaction between age and the Learnable-Dissociable galaxy compared to the Non-Learnable galaxy ($X^2 = 4.0$, $df = 1$, $p = .044$), such that older participants had increasingly worse source memory accuracy for the Learnable-Dissociable galaxy relative to the Non-Learnable galaxy. Additionally, there was a triple interaction between reward, age, and the Learnable-Dissociable galaxy compared to the Non-Learnable galaxy ($X^2 = 5.3$, $df = 1$, $p = .022$). In the Non-Learnable galaxy, source memory accuracy for treasure items and trash items both decreased with age. However, in the Learnable-Dissociable galaxy, source memory accuracy for trash more steeply decreased with age, while accuracy for treasure stayed constant.

To investigate the potential interaction between reward and squared age from the quadratic age model, we fit a linear mixed-effects model with age, squared age, and galaxy as regressors and a random intercept for each participant to predict decision criterion, which reflects response bias. We found a significant effect of squared age ($X^2 = 11.1$, $df = 1$, $p < .001$), which indicated that children and adults were biased to endorse remembered items as trash instead of treasure, whereas adolescents did not exhibit a response bias. No other effects reached significance.

3.3. Post-task questions

3.3.1. Estimating reward probability

After the memory task, participants answered questions probing their explicit attitudes about the planets and galaxies from the learning phase. We asked participants “How likely is it to find treasure on this planet?” on a scale from 1 (“Not likely”) to 7 (“Very likely”). Participants accurately rated the reward probabilities of the high-reward planets as greater than those of the low-reward planets in the Learnable-Yoked ($M_{\text{high}} = 5.08$, $SD_{\text{high}} = 1.40$, $M_{\text{low}} = 3.36$, $SD_{\text{low}} = 1.44$; paired t -test: $t(95) = 8.19$, $p < .001$; $BF = 6,727,117,845$) and Learnable-Dissociable ($M_{\text{high}} = 5.27$, $SD_{\text{high}} = 1.66$, $M_{\text{low}} = 3.40$, $SD_{\text{low}} = 1.70$; paired t -test: $t(95) = 6.40$, $p < .001$; $BF = 1,871,391$) galaxies. They rated both planets in the Non-Learnable galaxy similarly ($M_{\text{left}} = 4.23$, $SD_{\text{left}} = 1.51$, $M_{\text{right}} = 4.29$, $SD_{\text{right}} = 1.57$; paired t -test: $t(95) = -0.237$, $p = .81$; $BF = 0.116$). These results suggest participants had accurate explicit judgments of the relative reward probabilities in the three choice contexts.

3.3.2. Subjective control

To assess subjective sense of control for each context, we asked participants “When you were CAPTAIN, how much control did you have on finding treasure in this galaxy?” on a scale from 1 (“No Control”) to 7 (“Full Control”). We conducted a one-way repeated measures ANOVA to test for overall differences in controllability ratings across the galaxies. There was a marginally significant effect of galaxy ($F(2,190) = 2.93$, $p = .056$; $BF = 0.60$, proportional error estimate = $\pm 0.85\%$). Post-hoc paired t -tests revealed a significant difference in controllability ratings between the Learnable-Dissociable ($M = 5.02$, $SD = 1.56$) and Non-Learnable ($M = 4.53$, $SD = 1.55$; $t(95) = 2.46$, $p = .0158$; $BF = 1.95$). There was no significant difference in controllability rating between the Learnable-Yoked ($M = 4.66$,

$SD = 1.54$) and Non-Learnable galaxies ($t(95) = 0.586$, $p = .559$; $BF = 0.133$), or between the Learnable-Dissociable and Learnable-Yoked galaxies ($t(95) = 1.67$, $p = .097$; $BF = 0.433$). This suggests that participants were explicitly aware that the utility of their agency varied across choice contexts.

Because participants could detect the relative utility of their actions in the three galaxies, we next wanted to examine whether subjective assessments of controllability could predict the degree to which agency benefitted memory. To test this we first calculated participants’ “agency benefit” in each galaxy as the corrected hit rate for Captain trials minus that of Autopilot trials. We then fit a linear mixed-effects model with age and z-scored controllability as regressors and a random intercept for each subject to predict agency benefit. We found that perceived controllability did not predict agency benefit ($X^2 = 2.1$, $df = 1$, $p = .152$). That is, individual variability in the degree to which agency benefits memory cannot be predicted from subjective levels of control.

3.3.3. Context preference

Participants ranked the three galaxies in the order they were preferred (from 1 as most preferred to 3 as least). We ran a linear model predicting this preference score with age and galaxy as predictors and found that the Learnable-Dissociable galaxy ($M = 1.78$, $SD = 0.80$; $X^2 = 16.6$, $df = 1$, $p < .001$) and Learnable-Yoked galaxy ($M = 1.97$, $SD = 0.84$; $X^2 = 6.0$, $df = 1$, $p = .014$) were both preferred to the Non-Learnable galaxy ($M = 2.25$, $SD = 0.76$). There was no effect of age. Though the Learnable-Dissociable galaxy was ranked numerically as the most preferred, it did not significantly differ from the Learnable-Yoked (paired t -test: $t(95) = -1.26$, $p = .209$; $BF = 0.244$). To directly test the relationship between preference score and controllability, we ran a linear mixed-effects model with preference score and age as regressors and subject-level random intercepts to predict controllability and found that preference score significantly predicts controllability ($X^2 = 44.5$, $df = 1$, $p < .001$) where the galaxies that were more preferred were ranked as more controllable. These results suggest that participants prefer contexts in which they have a greater subjective sense of control.

4. Discussion

In this study, we examined whether the utility of agency during learning influenced subsequent memory and whether this effect varied across development. Participants ages 8 to 25 completed a two-day memory paradigm in which the utility of agency, reflected in the degree to which making choices is consequential for earning reward, varied across three choice contexts during encoding. Participants had greater memory accuracy for the items encoded with agency in the context for which agency had the greatest utility, suggesting that the ability to choose is beneficial for memory insofar as it is useful. This effect did not vary across our age range. Collectively, these findings suggest that, by middle childhood, individuals are sensitive to the degree to which their agency has utility, which is reflected in memory performance.

Our results suggest that merely making a choice, in and of itself, may be insufficient to enhance memory in some situations. If memory were prioritized for all choice outcomes in our study, then we would have expected to observe better recollection for information encoded with agency across all three encoding contexts. Instead, we found that the ability to choose facilitated memory only in the context in which agency had the greatest utility. This result stands in apparent contrast to previous findings that simply the act of choosing improves subsequent memory. Murty and colleagues (2015) conducted a similar study in which participants either had the opportunity to choose between one of two occluder screens or were instructed to select a highlighted option. After participants made their response, a trial-unique item was revealed from behind the selected occluder and participants were instructed to remember it for a memory test the next day. Unlike our paradigm, there were no reward values associated with any outcomes and participants did not have any explicit goal other than

completing the task. Murty et al. (2015) found that participants had better memory for the trial-unique items from the agency trials compared to the no-agency trials. Though we did not find any main effect of agency, a parsimonious account of the seemingly disparate findings of these two studies may be that memory is prioritized for information encoded with the greatest relative degree of control. Inferences of control have been proposed to take into account multiple properties that can be conceptualized as conferring incrementally greater levels of control (Huys & Dayan, 2009; Ly, Wang, Bhanji, & Delgado, 2019). At the most basic level, the opportunity to choose an action confers a minimal degree of control, with contingent relationships between actions and distinct outcomes, and the ability to obtain higher levels of reward through contingent action each increasing control estimates. As our ability to encode information in memory is constrained, priority may be given to information encountered under conditions of relatively higher control. In studies contrasting choice and no-choice conditions, this might promote encoding of outcomes obtained through choice alone (e.g., Murty et al., 2015), whereas, in our study, outcomes experienced in the condition that presented choice opportunity, contingent outcomes, and the ability to obtain greater reward through action was prioritized over outcomes in contexts affording a lesser degree of control.

Perceived control has been proposed to function as an intrinsic reward (Bobadilla-Suarez et al., 2017; Deci & Ryan, 2000; Leotti & Delgado, 2011), which suggests a potential mechanism underlying the effects on memory observed in our study. Both humans and non-human animals are sensitive to the controllability of their environments (Moscarello & Hartley, 2017) and prefer having agency (Bown, Read, & Summers, 2003; Catania & Sagvolden, 1980; Cordova & Lepper, 1996; Leotti, Iyengar, & Ochsner, 2010; Suzuki, 1997; Suzuki, 1999). People will pay a premium to have control even doing so is suboptimal (Bobadilla-Suarez et al., 2017; Owens, Grossman, & Fackler, 2014; Wang & Delgado, 2019). Consistent with this evidence that control is subjectively rewarding, striatal brain regions typically engaged during the anticipation and receipt of reward are also sensitive to conditions of agency (Fujiwara et al., 2013; Leotti & Delgado, 2011; Murty et al., 2015; Wang & Delgado, 2019), and rewards that are obtained contingently elicit greater striatal activity (Tricomi, Delgado, & Fiez, 2004) and function as more robust reinforcers (Cockburn, Collins, & Frank, 2014) than those that are passively received. Across a variety of experimental paradigms, memory is enhanced by the anticipation and receipt of rewards (Shohamy & Adcock, 2010), and the magnitude of striatal activation during encoding is associated with better subsequent memory (Adcock et al., 2006; Murty et al., 2015). Collectively, these findings suggest that both the intrinsic reward value of agency, along with the greater relative reward that could be obtained through action, might contribute to the observed memory benefit in the study condition in which agency had the greatest utility. However, future neuroimaging studies are needed to determine whether there is empirical support for this hypothesized mechanism.

We did not observe any age-related differences in how the utility of agency influenced subsequent memory. We hypothesized that distinct potential developmental patterns might emerge in our study depending on the neurocognitive mechanisms underpinning the influence of the utility of agency on memory. One possibility was that prioritizing memory according to the utility of agency requires cognitive control processes that continue to develop with age (Davidow et al., 2018; Insel et al., 2017; Störmer et al., 2014). Past work has demonstrated that intentional prioritization of memory for information explicitly labeled as “high-value” is evident in middle childhood, but markedly improves into adulthood (Castel et al., 2011; Hantén et al., 2007). Though we cannot make strong inferences on the basis of the present null result, the absence of any age-related increases in the mnemonic enhancement in the present study suggests that memory prioritization for outcomes of consequential choices may not rely on such “top-down” cognitive control processes. Instead, the observed age-invariant effect supports

the idea that the sensitivity of memory to the utility of agency might arise from engagement of neural reward circuitry, which exhibits functional sensitivity to both extrinsic reinforcement (Galvan et al., 2006; Van Den Bos, Cohen, Kahnt, & Crone, 2011) and intrinsic reward (Satterthwaite et al., 2012) in childhood. In our study, participants across our age range reported comparably strong preferences for the condition in which agency had high utility. Though direct empirical evidence is still needed, it is possible that such subjectively rewarding conditions may engage reward circuitry to an equivalent degree in children as well as adults, facilitating subsequent memory.

This study has some limitations. Because we did not observe any significant effects of age, it is unclear how early in development the utility of agency begins to influence memory. A recent study by Ruggeri et al. (2019) suggests that active control of learning benefits memory to a small degree in five year olds, and that this effect increases with age until reaching adult-like levels by eight years old. Though that study did not examine whether children are sensitive to the utility of agency, it is possible that the same developmental changes underlying age-related increases in the memory benefit from mere choice would also apply. As such, the degree to which memory is sensitive to the utility of memory may improve into middle childhood. Another limitation is that we are unable to make clear inferences about the temporal dynamics of the influence of agency on memory. We tested memory after a 24-hour delay, but it is unclear whether the mnemonic benefit of agency would be present in an immediate test, or whether it requires time-dependent consolidation, as has been observed in several past studies of reward facilitation of memory (Fischer & Born, 2009; Murayama & Kitagami, 2014; Patil, Murty, Dunsmoor, Phelps, & Davachi, 2017). Further research is needed to investigate how early this effect is present, as well as how long this memory benefit persists. Finally, we did not have an explicit measure of attention during the task. As attention has a profound influence on encoding success (Bouret & Sara, 2004; Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014; Muzzio, Kentros, & Kandel, 2009), it is possible that our results could be at least partially explained by differing levels of attention during the task. Participants may have been attending more to the controllable contexts, to the trials in which they had the opportunity to choose, or to the rewarding outcomes. Differential engagement of attention across these conditions could lead participants to preferentially encode the outcomes of specific trials. Further work is needed to investigate the role of attention in the mnemonic benefit of agency.

4.1. Conclusions

The present study demonstrates that the utility of choice modulates subsequent memory. Our findings suggest that individuals are sensitive not only to their ability to make choices, but to the degree to which those choices are consequential, and this in turn influences how memory is prioritized for information obtained through control. Moreover, we found that this effect was evident and did not vary in magnitude from middle childhood through young adulthood, suggesting that the underlying cognitive and neural mechanisms underlying this memory benefit are functional by middle childhood.

Funding

This work was supported by a Klingenstein-Simons Fellowship Award, National Science Foundation Career Grant 1654393, a Jacobs Foundation Early Career Fellowship, and the NYU Vulnerable Brain Project (C.A.H.) and an NSF GRFP (P.L.K.).

Declaration of competing interest

The authors report no conflicts of interest.

Acknowledgements

We would like to thank our many participants and their families for their time, Lena Church for her assistance in data collection, our colleagues in the lab and reviewers for their constructive feedback on this manuscript, and Macrina Cooper-White for her support.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2020.104239>.

References

- Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*(3), 507–517.
- Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, *96*(4), 703–719.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), <https://doi.org/10.1016/j.jml.2012.11.001>.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48.
- Bobadilla-Suarez, S., Sunstein, C. R., & Sharot, T. (2017). The intrinsic value of choice: The propensity to under-delegate in the face of potential gains and losses. *Journal of Risk and Uncertainty*, *54*(3), 187–202.
- Bouret, S., & Sara, S. J. (2004). Reward expectation, orientation of attention and locus coeruleus-medial frontal cortex interplay during learning. *European Journal of Neuroscience*, *20*(3), 791–802.
- Bown, N. J., Read, D., & Summers, B. (2003). The lure of choice. *Journal of Behavioral Decision Making*, *16*(4), 297–308.
- Castel, A. D., Humphreys, K. L., Lee, S. S., Galván, A., Balota, D. A., & McCabe, D. P. (2011). The development of memory efficiency and value-directed remembering across the life span: A cross-sectional study of memory and selectivity. *Developmental Psychology*, *47*(6), 1553–1564.
- Catania, A. C., & Sagvolden, T. (1980). Preference for free choice over forced choice in pigeons. *Journal of the Experimental Analysis of Behavior*, *34*(1), 77–86.
- Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, *3*(2), 57–65.
- Cockburn, J., Collins, A. G. E., & Frank, M. J. (2014). A reinforcement learning mechanism responsible for the valuation of free choice. *Neuron*, *83*(3), 551–557.
- Cordova, D. I., & Lepper, M. R. (1996). Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. *Journal of Educational Psychology*, *88*(4), 715–730.
- Cox, J. C., & Dobbins, I. G. (2011). The striking similarities between standard, distractor-free, and target-free recognition. *Memory & Cognition*, *39*(6), 925–940.
- Davidow, J. Y., Insel, C., & Somerville, L. H. (2018). Adolescent development of value-guided goal pursuit. *Trends in Cognitive Sciences*, *22*(8), 725–736.
- Deci, E. L., & Ryan, R. M. (2000). The “what” and “why” of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry*, *11*(4), 227–268.
- Dukas, R. (1999). Costs of memory: Ideas and predictions. *Journal of Theoretical Biology*, *197*(1), 41–50.
- Feldman, A., & Acredolo, L. (1979). The effect of active versus passive exploration on memory for spatial location in children. *Child Development*, *50*(3), 698–704.
- Fischer, S., & Born, J. (2009). Anticipated reward enhances offline learning during sleep. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(6), 1586–1593.
- Fujiwara, J., Usui, N., Park, S. Q., Williams, T., Iijima, T., Taira, M., ... Tobler, P. N. (2013). Value of freedom to choose encoded by the human brain. *Journal of Neurophysiology*, *110*(8), 1915–1929.
- Galvan, A., Hare, T. A., Parra, C. E., Penn, J., Voss, H., Glover, G., & Casey, B. J. (2006). Earlier development of the accumbens relative to orbitofrontal cortex might underlie risk-taking behavior in adolescents. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *26*(25), 6885–6892.
- Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, reward, and information seeking. *Journal of Neuroscience*, *34*(46), 15497–15504.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, *7*(5), 464–481.
- Hantén, G., Li, X., Chapman, S. B., Swank, P., Gamino, J., Roberson, G., & Levin, H. S. (2007). Development of verbal selective learning. *Developmental Neuropsychology*, *32*(1), 585–596.
- Huys, Q. J. M., & Dayan, P. (2009). A Bayesian formulation of behavioral control. *Cognition*, *113*(3), 314–328.
- Insel, C., Kastman, E. K., Glenn, C. R., & Somerville, L. H. (2017). Development of corticostriatal connectivity constrains goal-directed behavior during adolescence. *Nature Communications*, *8*(1), 1605.
- Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, *59*(4), 434–446.
- Leotti, L. A., & Delgado, M. R. (2011). The inherent reward of choice. *Psychological Science*, *22*(10), 1310–1318.
- Leotti, L. A., Iyengar, S. S., & Ochsner, K. N. (2010). Born to choose: The origins and value of the need for control. *Trends in Cognitive Sciences*, *14*(10), 457–463.
- Lesnoff, M., & Lancelot, R. (2012). *aod: Analysis of overdispersed data. R package version, 1, 3–1*.
- Ly, V., Wang, K. S., Bhanji, J., & Delgado, M. R. (2019). A reward-based framework of perceived control. *Frontiers in Neuroscience*, *13*. <https://doi.org/10.3389/fnins.2019.00065>.
- Markant, D., DuBrow, S., Davachi, L., & Gureckis, T. M. (2014). Deconstructing the effect of self-directed study on episodic memory. *Memory & Cognition*, *42*(8), 1211–1224.
- McComas, J., Dulberg, C., & Latter, J. (1997). Children's memory for locations visited: Importance of movement and choice. *Journal of Motor Behavior*, *29*(3), 223–229.
- Moscarello, J. M., & Hartley, C. A. (2017). Agency and the calibration of motivated behavior. *Trends in Cognitive Sciences*, *21*(10), 725–735.
- Morey, R. D., & Rouder, J. N. (2018). *BayesFactor: Computation of Bayes factors for common designs. R package version 0.9.12.4-2*.
- Murayama, K., & Kitagami, S. (2014). Consolidation power of extrinsic rewards: Reward cues enhance long-term memory for irrelevant past events. *Journal of Experimental Psychology: General*, *143*(1), 15–20.
- Murty, V. P., DuBrow, S., & Davachi, L. (2015). The simple act of choosing influences declarative memory. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(16), 6255–6264.
- Muzzio, I. A., Kentros, C., & Kandel, E. (2009). What is remembered? Role of attention on the encoding and retrieval of hippocampal representations. *The Journal of Physiology*, *587*(12), 2837–2854.
- Owens, D., Grossman, Z., & Fackler, R. (2014). The control premium: A preference for payoff autonomy. *American Economic Journal: Microeconomics*, *6*(4), 138–161.
- Partridge, E., McGovern, M. G., Yung, A., & Kidd, C. (2015). Young children's self-directed information gathering on touchscreens. Proceedings of the 37th annual conference of the Cognitive Science Society, Austin, TX. Cognitive Science Society.
- Patil, A., Murty, V. P., Dunsmoor, J. E., Phelps, E. A., & Davachi, L. (2017). Reward retroactively enhances memory consolidation for related items. *Learning & Memory*, *24*, 65–69. <https://doi.org/10.1101/lm.042978.116>.
- R Core Team. (2013). R: A language and environment for statistical computing. Retrieved from <ftp://ftp.uvigo.es/CRAN/web/packages/dplR/vignettes/intro-dplR.pdf>.
- Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology*, *25*, 111–164.
- Ruggeri, A., Markant, D. B., Gureckis, T. M., Bretzke, M., & Xu, F. (2019). Memory enhancements from active control of learning emerge across development. *Cognition*, *186*, 82–94.
- Satterthwaite, T. D., Ruparel, K., Loughhead, J., Elliott, M. A., Gerraty, R. T., Calkins, M. E., ... Wolf, D. H. (2012). Being right is its own reward: Load and performance related ventral striatum activation to correct responses during a working memory task in youth. *NeuroImage*, *61*(3), 723–729.
- Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, *14*(10), 464–472.
- Sim, Z. L., Tanner, M. M., Alpert, N. Y., & Xu, F. (2015). *Children learn better when they select their own data.* (CogSci).
- Steyvers, M., & Griffiths, T. L. (2008). *Rational analysis as a link between human memory and information retrieval.* The Probabilistic Mind: Prospects for Bayesian Cognitive Science 329–349.
- Störmer, V., Eppinger, B., & Li, S.-C. (2014). Reward speeds up and increases consistency of visual selective attention: A lifespan comparison. *Cognitive, Affective, & Behavioral Neuroscience*, *14*(2), 659–671.
- Suzuki, S. (1997). Effects of number of alternatives on choice in humans. *Behavioural Processes*, *39*(2), 205–214.
- Suzuki, S. (1999). Selection of forced- and free-choice by monkeys (Macaca fascicularis). *Perceptual and Motor Skills*, *88*(1), 242–250.
- Tricomi, E. M., Delgado, M. R., & Fiez, J. A. (2004). Modulation of caudate activity by action contingency. *Neuron*, *41*(2), 281–292.
- Van Den Bos, W., Cohen, M. X., Kahnt, T., & Crone, E. A. (2011). Striatum-medial prefrontal cortex connectivity predicts developmental changes in reinforcement learning. *Cerebral Cortex*, 1–9.
- Voss, J. L., Galvan, A., & Gonsalves, B. D. (2011). Cortical regions recruited for complex active-learning strategies and action planning exhibit rapid reactivation during memory retrieval. *Neuropsychologia*, *49*(14), 3956–3966.
- Voss, J. L., Gonsalves, B. D., Federmeier, K. D., Tranel, D., & Cohen, N. J. (2011). Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nature Neuroscience*, *14*(1), 115–120.
- Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, *14*(5), 779–804.
- Wang, K. S., & Delgado, M. R. (2019). Corticostriatal circuits encode the subjective value of perceived control. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhz045>.