Supplemental Methods and Results

S1 Main Task

S1.1 Example regression design matrix

To clarify how predictors were coded in the multiple regression predicting Test Phase choices on the second day of the study, we provide an example design matrix for a single trial.

Crew Member, Trial	I_{-1}	R_{-1}	R_{-2}	R_{-3}	EI_{-1}	ER_{-1}	EC_{-1}
White Beard, 4	1	0	1	1	0	0	0
Red Beard, 4	0	0	0	0	1	1	-0.33
Black Beard, 4	0	0	0	0	0	0	0.45

Table S1: This table depicts the rows of the design matrix that code for the fourth trial on day two of the study, given the following scenario. On the three preceding choice trials, White Beard was selected, and the first two selections were rewarded. After the third trial, a memory probe (an image of a banana, see Figure 1C) was presented and evoked a Day 1 trial on which Red Beard was chosen. This trial occurred on the cityscape island/context, where Red Beard was chosen on 18 trials and rewarded on 6, and Black Beard was chosen on 22 trials and rewarded on 16. The design matrix contains three rows per trial, each reflecting the contribution of the independent variables to the probability of picking a crew member (White Beard, Red Beard, Black Beard). Predictors encode choice-relevant information from past trials. The first column (I_{-1}) indicates whether the crew member of interest was selected on the most recent trial, and the second column (R_{-1}) indicates whether that choice was rewarded. The third and fourth columns (R_{-2}, R_{-3}) indicate whether the crew member was rewarded on the two preceding trials. If the most recent trial before the current choice was a recognition memory trial in which an "old" object was presented, the fifth through seventh columns contain choice and value information from the trial evoked associated the object probe: the identity of the chosen crew member (EI_{-1}), whether that choice was rewarded (ER_{-1}), and the rewards associated with that crew member across the evoked context (EC_{-1}). The context-level reward was computed as the number of rewarded choices minus unreward choices, divided by the total number of times the crew member was chosen within the context; if the crew member was never chosen, this value was set to 0. The dependent variable is binary, encoding whether the crew member of interest was selected on the current trial (trial 4 in this example). Regression coefficients therefore reflect the contribution of each predictor to the probability of selecting a given crew member. Predictors not shown are age and its interactions with each predictor shown above.

S1.2 Memory performance

Both recognition memory and source memory in the main task were poor (probe recognition: mean $d' = 0.83$, SD = 0.76; source memory: mean proportion correct = 0.18, SD = 0.07) but significantly better than chance (probe recognition: one-sample t-test against 0, $t(105) = 11.25$, $p < .001$; source memory: one-sample t-test against 0.17, $t(105) = 2.57$, p = .012). Neither recognition memory or source memory varied significantly with age (Fig. S1; probe recognition: ρ=0.079, p=.42; source memory: ρ=0.14, p=.15). Recognition memory and source memory were marginally related to one another (ρ=0.23, p=.064) as well as both being related to a measure of memory precision (lure discrimination index i.e., LDI) taken from a separate task, the Mnemonic Similarity Task (recognition memory & LDI: ρ=0.33, p=.006; source memory & LDI: ρ=0.27, p=.024; S3).

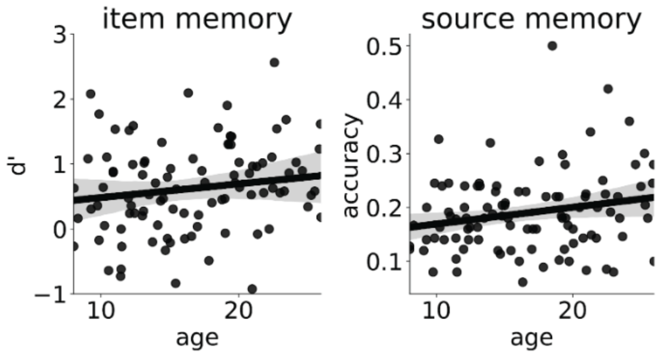

Figure S1: **Item and source memory do not improve with age. A. Item memory.** Recognition memory of probe objects, as measured by d', did not change with age ($\rho(104)$ =0.079 , p=.42). **B. Source memory.** Similarly, accuracy on source memory trials also did not increase with age ($\rho(104)$ =0.14, p=.15).

S1.3 Explicit source memory's relationship with context sensitivity

Our proposed mechanism relies on contextual information to trigger the retrieval of multiple contextually linked episodes. A natural hypothesis is that the contextual information retrieved at decision time should predict choice. Consistent with this, Bornstein & Norman (2017) showed that covert neural measures of context reinstatement mediate the probed-context effect. In contrast, our study relies on behavioral data and therefore lacks a trial-by-trial measure of context reinstatement at decision time. Instead, we must rely on participants' explicit source memory judgments, collected at the end of the task, as a proxy for retrieved context. Using these judgments, we conducted two analyses to test whether source memory modulated sensitivity to the evoked context. First, we examined whether source memory accuracy modulated the influence of the probed context. To do so, we allowed the probed-context regressor to interact with trial-level source memory accuracy (1 = correct, 0 = incorrect). Second, rather than assuming that participants retrieved the true context associated with each probe, we used their source memory responses to infer which context they retrieved. Then, we tested whether this inferred-context reward regressor better predicted choice than the reward computed under the assumption of accurate retrieval.

In the first analysis, all primary effects reported in the main text remained (Table S2). However, we found no significant interaction between source memory accuracy and the probed-context effect. In the second analysis, the main effects again remained robust, including the effect of the inferred-context reward regressor and its interaction with age. However, although the inferred-context regressor was significant, its effect was slightly weaker than that of the original context-reward regressor. In addition, the model fit was slightly worse when using the inferred-context regressor (AIC = 62841.3) than when using the original context-reward regressor (AIC = 62735.1).

We found no association between explicit source memory judgments and the probed-context effect. The absence of a relationship could reflect that context retrieval is an implicit process, one that our explicit source memory measure is a poor proxy for.

S1.4 Distinguishing context-guided episodic sampling from context-specific Q-learning

To explain participants' behavior in the Test Phase, we propose a context-guided episodic sampling mechanism. A potential alternative to our account is that participants do not sample multiple contextually linked episodes at decision time. Instead, they may be averaging the outcomes from each action in a given context during the Encoding Phase, storing these values, and retrieving them when a memory probe reinstates the context during the Test Phase. In reinforcement learning, this strategy corresponds to a context-specific Q-value. Context-guided episodic sampling and context-specific Q-value learning differ along two dimensions. First, they differ in how actions are constructed. In our variant of episodic sampling, past experience is uniformly averaged across a context. In Q-learning, experience is recency-weighted during averaging. Second, they differ in when action values are constructed. In episodic sampling, values are constructed at decision time, while in Q-learning, values are updated incrementally during encoding. Because the two strategies construct values differently, we can compute the values each would assign actions based on a participant's Day 1 experience and test which better predicts their choices on Day 2. Because the two strategies differ in when values are constructed, they make different predictions about sensitivity to temporal context. In episodic sampling, the temporal position of the probed trial within its context should modulate the strength of the probed-context effect. This is because trials temporally proximal to the probed trial

term	estimate	SE	p-value
Intercept	-0.65	0.01	$p < .001$
last choice	0.99	0.015	$p < .001$
reward_{-1}	0.71	0.048	$p < .001$
reward_{-2}	0.26	0.026	$p < .001$
reward_{-3}	0.18	0.024	$p < .001$
choice probed trial	0.082	0.021	$p = .0001$
reward probed trial	0.04	0.021	$p = .05$
probed context	0.13	0.062	$p = .034$
source accuracy	-0.018	0.06	$p = .77$
age	0.015	0.011	$p = .17$
probed context x source accuracy	0.029	0.016	$p = .86$
last choice x age	0.28	0.015	$p < .001$
reward_{-1} x age	0.19	0.049	$p = .039$
reward_{-2} x age	0.0037	0.026	$p = .89$
reward_{-3} x age	0.018	0.024	$p = .46$
choice probed trial x age	-0.045	0.022	$p = .037$
reward probed trial x age	-0.0095	.021	$p = .65$
probed context x age	0.099	0.050	$p = .045$
source accuracy x age	0.013	0.058	$p = .82$
probed context x source accuracy x age	-0.011	0.12	$p = .93$

Table S2: Estimated coefficients from a model predicting Test Phase choices. Regressors include those from the main text along with trial-level source memory accuracy. Source memory accuracy was allowed to interact with the probed-context regressor and its interaction with age.

are more likely to be retrieved; thus, when the probed trial occurs closer to the reversal point, the probed-context effect should be stronger, because trials from the other side of the reversal are more likely to be sampled. This temporal context effect would not be predicted by a context-specific Q-value mechanism.

S1.4.1 Context-specific Q-values do not predict Test Phase choices

In the regression presented in the main text, the context-reward regressor reflects the mean value estimated by an episodic sampling process under minimal assumptions. Critically, each sampled episode contributes equally to this estimate. In our task, this property yields action values that are qualitatively distinct from those produced by a Q-learning mechanism that tracks a context-specific Q-value (e.g., an action value computed solely from experience within a given island). In Q-learning, outcomes contribute unequally to value estimates. Their impact depends on both the magnitude of the prediction error they elicit and their temporal recency. As a result, context-guided episodic sampling and context-specific Q-learning generate distinct predictions, allowing us to test which better accounts for participants' choices.

To generate a time series of Q-values, we fit participants' choices with a model that re-initializes action values at the beginning of each island/context. After each observed outcome, values are updated according to the reward prediction error scaled by a learning rate, α. We assume that participants select actions according to a mixture policy: with probability ϵ the agent chooses randomly, and with probability $1 - \epsilon$ the agent chooses according to a softmax rule with inverse temperature β. We treat α, β, and ϵ as free parameters and fit them to participants' choices from day 1 of the study. Median parameter estimates were $\alpha = 0.48$, $\beta = 5.10$, and $\epsilon = 0.07$. None of these parameters varied significantly with age (α: $\rho = -0.069$, $p = .48$; β: $\rho = 0.0081$, $p = .93$; ϵ: $\rho = -0.14$, $p = .15$).

For each participant, we extracted the Q-values at the end of each context and used those as a replacement for the episodic sampling context reward regressor we use in the main text. The effect of the replacement regressor was not statistically significant (b = -0.098, SE = 0.071, p = .17) nor was its interaction with age (b = 0.073, SE = 0.069, p = .29). When we included both the context-specific Q-value regressor and the episodic sampling context reward regressor in the same model, we found that the weight on the original context reward regressor remained essentially unchanged ($b = 0.15$, $SE = 0.045$, $p = .0012$). Meanwhile, the Q-value regressor showed only a trend-level effect and, critically, in the negative direction ($b = -0.10$, $SE = 0.057$, $p = .070$). These results argue against the alternative account that participants retrieved context-specific Q-values when presented with object

term	estimate	SE	p-value
Intercept	-0.65	0.011	$p < .001$
last choice	0.99	0.015	$p < .001$
reward$_{-1}$	0.71	0.057	$p < .001$
reward$_{-2}$	0.26	0.026	$p < .001$
reward$_{-3}$	0.18	0.024	$p < .001$
choice probed trial	0.082	0.022	$p = .0001$
reward probed trial	0.04	0.021	$p = .049$
probed inferred context	0.14	0.046	$p = .0021$
age	0.015	0.010	$p = .15$
last choice x age	0.28	0.015	$p < .001$
reward$_{-1}$ x age	0.10	0.047	$p = 0.34$
reward$_{-2}$ x age	0.0036	0.026	$p = .89$
reward$_{-3}$ x age	0.018	0.024	$p = .46$
choice probed trial x age	-0.045	0.021	$p = .034$
reward probed trial x age	-0.0096	.021	$p = .64$
probed inferred context x age	0.096	0.045	$p = .034$

Table S3: Estimated coefficients from a model predicting Test Phase. Rather than assuming retrieval of the true context associated with each probe, context was inferred from participants' source memory responses. The model includes the same predictors as in the main text, with the inferred-context reward regressor replacing the original context-reward regressor.

images from the previous day.

S1.4.2 Temporal context of the probed trial modulates the effect of distant experience on choice

To test whether the temporal context of the probe modulates the strength of probed-trial and probed-context effects, we added a regressor capturing the distance of the probed trial from the reversal in reward probabilities (which occurred at trial 11 in each context). We allowed this regressor to interact with the three regressors capturing distant experience: choice on the probed trial, reward on the probed trial, and how frequently an action was rewarded across the entire probed context. The primary effects reported in the main text remained robust. We continued to observe significant effects of recent rewards (Table S4; $b_{reward-1} = 0.67$, $SE = 0.014$, $p < .001$; $b_{reward-2} = 0.23$, $SE = 0.011$, $p < .001$; $b_{reward-3} = 0.15$, $SE = 0.012$, $p < .001$), choice on the probed trial ($b = 0.065$, $SE = 0.022$, $p = .0029$), and reward frequency across the probed context ($b = 0.11$, $SE = 0.050$, $p = .021$), as well as an interaction between age and the probed-context regressor ($b = 0.11$, $SE = 0.049$, $p = .028$). Turning to the probe-temporal-context regressor, we observed significant interactions between probe distance from reversal and both reward on the probed trial ($b = 0.022$, $SE = 0.0072$, $p = .0024$) and choice on the probed trial ($b = 0.017$, $SE = 0.0074$, $p = .018$). When the probed trial occurred earlier in the context (i.e., farther from the reversal), participants were more likely to repeat the option chosen on the probed episode—particularly when that episode had been rewarded—than when the probed trial occurred later (Fig. S2). In contrast, when the probed trial occurred close to the reversal, this pattern reversed. Participants were *less* likely to repeat an option if it had been rewarded on the probed trial relative to when it had not. This reversal is consistent with the probed episode triggering retrieval of temporally proximal episodes from the post-reversal period, when that option was no longer likely to yield reward. Finally, a three-way interaction between probe distance from reversal, choice on the probed trial, and age showed a trend-level effect ($b = 0.014$, $SE = 0.0074$, $p = .063$), suggesting that the probe-temporal-context effect was stronger in older participants. Overall, these results support our claim that decisions are shaped by the sampling of contextually linked episodes and suggest that this process strengthens with age.

S2 Demographics of participants who completed the Mnemonic Similarity Task and the Two Step Task

The subsample of participants who completed the Mnemonic Similarity Task and the Two-Step Task (N = 67, 39 females) was similar in age and gender composition to the full sample (N = 106, 54 females). In the subsample,

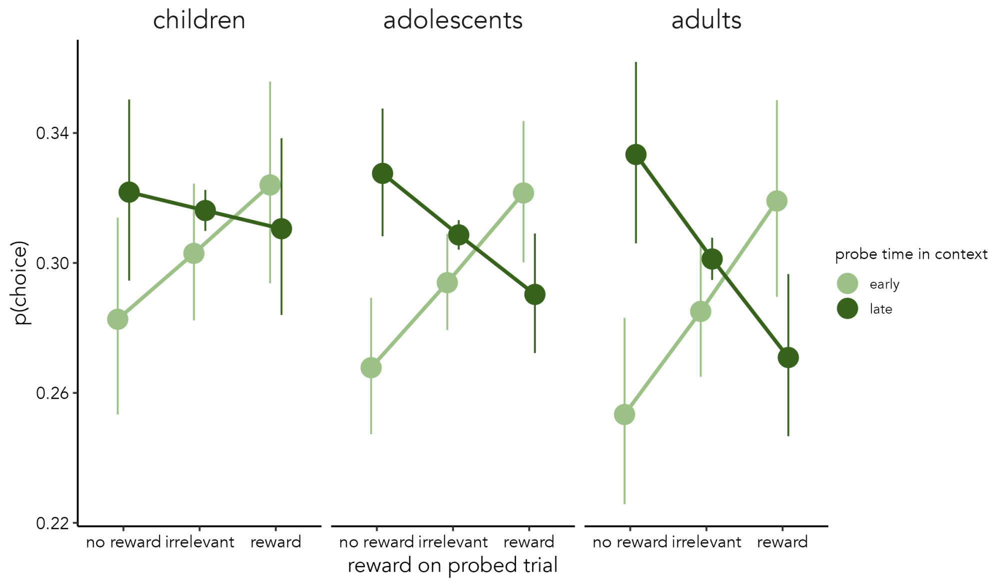

Figure S2: **Temporal context of memory probe modulates its effect on choice.** Markers show the estimated probability of choosing an option when it was rewarded on the evoked trial ("reward"), not rewarded ("no reward"), or when no trial was recently evoked ("irrelevant"). Marker color indicates when the probed trial occurred within its context. Error bars denote 95% confidence intervals. Age is plotted discretely (children: 8–12.99 years; adolescents: 13–17.99 years; adults: 18–25 years) as is time in context (early: near the beginning of the context; late near the reversal). However, both variables were treated as continuous in the statistical model.

mean age was 16.82 years (SD = 5.05; range = 8.03–25.97), compared to 16.96 years in the full sample (SD = 5.09; range = 8.03–25.97). The age distribution in the subsample was also relatively uniform.

S3 Mnemonic Similarity Task

We tested for age-related changes in memory precision as measured by the lure discrimination index (LDI). More specifically, LDI indexes an individual's ability to recognize an item as distinct from a highly similar, previously encountered item. LDI scores can range from -1 to 1 with positive values indicating successful discrimination, negative values indicating a response bias towards reporting that items are "similar", and 0 indicating chance performance. Participants across our age range demonstrated well-above-chance discrimination (mean LDI=0.35, SD=0.25, one sample t-test against 0 (t(66)=11.52, p¡.001). However, older participants tended to have higher LDI scores (Figure S4B; $\rho(65) = 0.25$, $p = .04$). As a baseline memory measurement, we also calculated participants' recognition memory (REC) scores for only the dissimilar items, with a range of -1 to 1 and positive values indicating better performance. Participants again performed well above chance (mean REC=0.77, SD=0.15; $t(66)$=40.75, $p < .001$), however, REC scores did not significantly increase with age (Figure S4C; $\rho(65) = 0.13$, $p = .29$).

S4 The Two Step Task

We aimed to replicate findings from prior work demonstrating age-related changes in model-based strategy use [Decker et al., 2016, Potter et al., 2017, Nussenbaum et al., 2020]. Across our sample, we observed both a main effect of reward and an interaction effect between reward and transition type, suggesting that participants used both model-free and model-based strategies throughout the task (b_{reward}=1.45, SE=0.17, $p < .001$; $b_{reward\,X\,transition}$=-1.52, SE=0.24, $p < .001$). Older participants were not significantly more likely to repeat their first stage choices regardless of the previous outcome (b_{age}=0.19, SE=0.12, $p = .11$). Importantly, we did not find an age by reward interaction effect ($b_{age\,X\,reward}$=0.27, SE=0.17, $p = 0.11$), suggesting that model-free learning did not vary with age. Finally, and of primary interest, we tested whether model-based strategy use changed with age. Contrary to previous findings, we did not observe significant age-related change, (Figure S5B; $b_{reward\,X\,transition\,X\,age}$=-0.30, SE=0.23, $p = .19$). This divergent finding may potentially be a result of our small sample size of 67. To determine

term	estimate	SE	p-value
Intercept	-0.66	0.021	$p < .001$
last choice	0.98	0.014	$p < .001$
reward$_{-1}$	0.67	0.014	$p < .001$
reward$_{-2}$	0.23	0.011	$p < .001$
reward$_{-3}$	0.15	0.012	$p < .001$
choice probed trial	0.065	0.022	$p = .0029$
reward probed trial	0.021	0.021	$p = .31$
probed context	0.11	0.050	$p = .021$
probe time in context	-0.0067	0.0037	$p = .075$
age	0.013	0.021	$p = .52$
choice probed trial x probe time in context	0.017	0.0074	$p = .018$
reward probed trial x probe time in context	0.022	0.0072	$p = .0024$
probed context x probe time in context	-0.016	0.016	$p = .32$
last choice x age	0.25	0.014	$p < .001$
reward$_{-1}$ x age	0.098	0.014	$p < .001$
reward$_{-2}$ x age	0.0041	0.011	$p = .70$
reward$_{-3}$ x age	0.017	0.012	$p = .14$
probe time in context x age	-0.00081	0.0037	$p = .83$
choice probed trial x age	-0.055	0.022	$p = .013$
reward probed trial x age	-0.015	.021	$p = .48$
probed context x age	0.11	0.049	$p = .028$
choice probed trial x probe time in context x age	0.014	0.0074	$p = .063$
reward probed trial x probe time in context x age	0.0093	0.0072	$p = .19$
probed context x probe time in context x age	-0.018	0.016	$p = .26$

Table S4: Estimated coefficients from a model predicting choices from the Test Phase. The regressors include those presented in the main text along with one indicating the probed trial's distance from the reversal. We allowed this new regressor to interact with the three regressors encoding distance experience (choice on the probed trial, reward on the probed trial, and rewards received across the probed context) and their interactions with age.

whether we were underpowered, we conducted a power analysis to estimate the sample size necessary to achieve 80% power to detect a three-way interaction between reward, transition type, and age ($\alpha = 0.05$). We used effect sizes from Potter et al. (2017) as a conservative benchmark, as that study yielded the smallest effect sizes of the three previous studies from our lab using this task in developmental populations [Decker et al., 2016, Potter et al., 2017, Nussenbaum et al., 2020]. For the three-way interaction, we determined the proportion of simulations in which $p < \alpha$ across 5,000 simulations for sample sizes ranging from 30 to 200. We found that a sample size of at 170 was required to achieve at least 80% power (Fig. S6), confirming that we are indeed underpowered to detect the three-way interaction effect.

S5 Relationships between memory-guided decision making, memory precision, and forward planning

For the regressions reported in this section, we extracted the random slopes from the logistic regression model described in the main text, but first, we re-estimated the model with age and its interactions excluded. We omitted age from this model so that we could examine how age interacted with our predictors of interest (e.g., the influence of the probed trial's reward).

Results from separate linear regressions predicting individuals' random slopes for (S5) the "average reward across the probed context", (S6) "reward on probed trial", and (S7) "choice on probed trial". Each regression included an intercept term and a regressor encoding lure discrimination index.

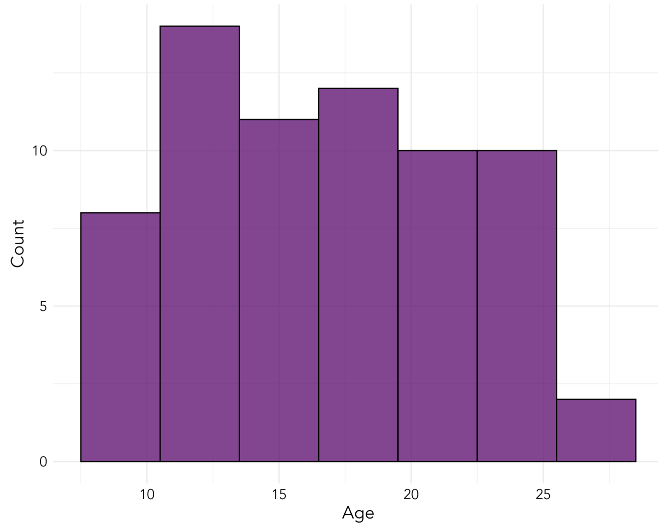

Figure S3: The distribution of ages for the 67 participants who completed all three tasks (main task, the Mnemonic Similarity Task, and the Two Step Task).

term	estimate	SE	p-value
Intercept	-0.09	0.06	0.16
LDI	0.30	0.14	0.04

Table S5: Predicting the influence of probed context

Results from separate linear regressions predicting individuals' extent of planning. Each regression included an intercept term and effects of age and one of the random slopes from Table S5-7 (S8: probed context, S9: probed trial reward, S10: probed trial choice), as well as an interaction term.

S6 Confidence judgments in the main regression

References

Johannes H Decker, A Ross Otto, Nathaniel D Daw, and Catherine A Hartley. From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychol. Sci.*, 27(6):848–858, June 2016.

Tracey C S Potter, Nessa V Bryce, and Catherine A Hartley. Cognitive components underpinning the development of model-based learning. *Dev. Cogn. Neurosci.*, 25:272–280, June 2017.

Kate Nussenbaum, Maximilian Scheuplein, Camille V Phaneuf, Michael D Evans, and Catherine A Hartley. Moving developmental research online: Comparing in-lab and web-based studies of model-based reinforcement learning. *Collabra Psychol.*, 6(1), November 2020.

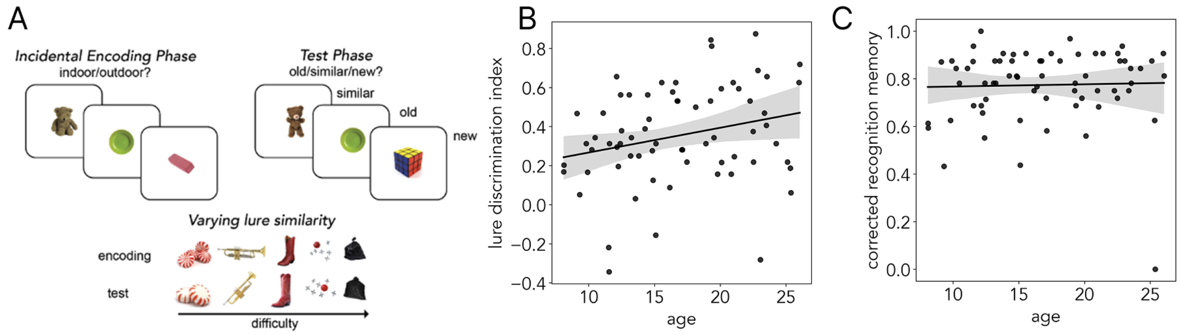

Figure S4: **Mnemonic Similarity Task. A. Task.** Participants completed two phases: an incidental encoding phase and a test (retrieval) phase. During the incidental encoding phase, participants perform a cover task in which they are sequentially shown 64 object images and are asked to indicate whether the object belongs indoors or outdoors. During the test phase, participants were presented sequentially with 96 object images and asked to respond whether they saw an object in the previous phase ("old"), saw a similar but not identical object ("similar"), or the object was completely novel ("new"). 32 test objects were target images (identical to an image shown during the encoding phase), 32 were lures (highly similar to objects shown during the encoding phase), and 32 were foils (images that were dissimilar to any encoding phase objects). Lure objects varied in similarity, falling into 5 bins of discrimination difficulty. **B-C.** Results. Lure discrimination significantly improved with age ($\rho(65) = 0.25$, p=.04) while corrected recognition memory did not ($\rho(65) = 0.13$, p=.29). Markers indicate individual estimates of lure discrimination index/corrected recognition memory plotted against age. Shaded regions indicate 95% confidence intervals.

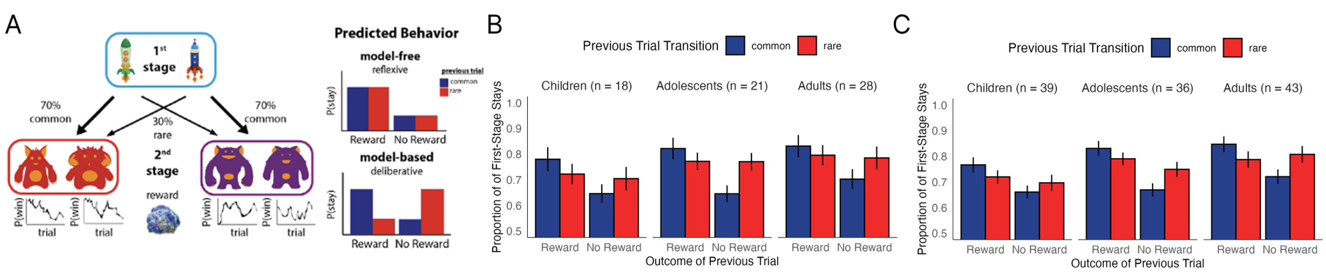

Figure S5: **Two step Task. A. Task.** Participants completed 200 trials. On each trial, they made two choices. At the first stage, they chose between two spaceships that would travel to one of two planets with differing frequency. On a planet at the second stage, they would then choose between two aliens, each with a unique drifting probability of reward. This task allows for the dissociation between model-free and model-based behavior. A model-free learner would be more likely to repeat their first stage action from the last trial if it led to reward, regardless of the transition structure. **B-C.** Results. Proportion of first-stage choice repetitions ("stays") as a function of the previous trial's outcome and transition type, shown separately for each age group. Children were younger than 13 years; adolescents were 13–17.99 years. In the subset of 67 participants, the reward × transition × age interaction was not significant ($b_{reward\,X\,transition\,X\,age}$=-0.30, SE=0.23, p =.19), but the interaction emerged as significant in the full sample of 118 participants ($b_{reward\,X\,transition\,X\,age}$-0.34, SE=0.17, p =.045).

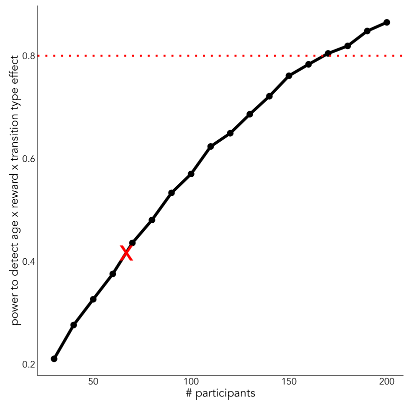

Figure S6: Simulated power (5,000 runs per n) to detect the reward × transition × age interaction on stay choices in the two-step task, using effect sizes from Potter et al. (2017). The red dotted line indicates 80% power at α=0.05. Achieving 80% power requires approximately 170 participants. The red “” marks our sample size (N = 67), indicating that the study is underpowered to detect this interaction.

term	estimate	SE	p-value
Intercept	-0.01	0.03	0.70
LDI	0.05	0.06	0.45

Table S6: Predicting the influence of reward on probed trial

term	estimate	SE	p-value
Intercept	0.08	0.06	0.18
LDI	-0.23	0.14	0.10

Table S7: Predicting the influence of choice on probed trial

term	estimate	SE	p-value
Intercept	0.31	0.04	0.00
probed context	0.19	0.13	0.14
age	0.06	0.04	0.12
probed context*age	0.14	0.15	0.36

Table S8: Predicting the extent of planning with age and sensitivity to probed context

term	estimate	SE	p-value
Intercept	0.31	0.04	0.00
probed trial reward	0.59	0.31	0.07
age	0.05	0.04	0.20
probed trial reward*age	0.81	0.38	0.04

Table S9: Predicting the extent of planning with age and sensitivity to the reward received on the probed trial

term	estimate	SE	p-value
Intercept	0.32	0.04	0.00
probed trial choice	0.03	0.14	0.85
age	0.05	0.04	0.20
probed trial choice*age	0.26	0.14	0.06

Table S10: Predicting the extent of planning with age and sensitivity to the reward received on the probed trial

term	estimate	SE	p-value
Intercept	-0.65	0.01	$p < .001$
last choice	0.99	0.015	$p < .001$
reward_{-1}	0.71	0.047	$p < .001$
reward_{-2}	0.26	0.026	$p < .001$
reward_{-3}	0.18	0.024	$p < .001$
choice probed trial	0.082	0.021	$p = .0001$
reward probed trial	0.04	0.021	$p = .05$
probed context	0.14	0.046	$p = .002$
age	0.015	0.01	$p = .15$
last choice x age	0.28	0.015	$p < .001$
reward_{-1} x age	0.10	0.047	$p = .034$
reward_{-2} x age	0.0036	0.024	$p = .89$
reward_{-3} x age	0.018	0.024	$p = .46$
choice probed trial x age	-0.045	0.021	$p = .034$
reward probed trial x age	-0.0096	.021	$p = .64$
probed context x age	0.096	0.045	$p = .034$

Table S11: Estimated coefficients from a model predicting Test Phase choices. For predictors based on distant experience (choice on the probed trial, reward on the probed trial, and probed context), analyses were restricted to trials in which participants reported high confidence in having seen the probe object.