Supplementary Materials: Differences in Learning Across the Lifespan Emerge via Resource-Rational Computations

Rasmus Bruckner[1,2,*], Matthew R. Nassar[3,4], Shu-Chen Li[5,6], Ben Eppinger[7,8]

[1]Department of Education and Psychology, Freie Universität Berlin; [2]Institute of Psychology, University of Hamburg; [3]Robert J. & Nancy D. Carney Institute for Brain Science, Brown University; [4]Department of Neuroscience, Brown University; [5]Faculty of Psychology, Technische Universität Dresden; [6]Centre for Tactile Internet with Human-in-the-Loop, Technische Universität Dresden; [7]Department of Psychology, University of Greifswald; [8]Department of Psychology, Concordia University; *Corresponding author: Rasmus Bruckner (rasmus.bruckner@fu-berlin.de)

Identical Picture and Spot-a-Word tests

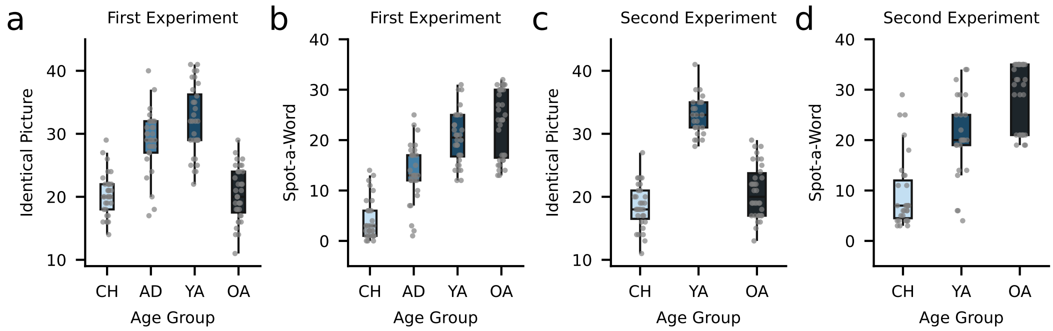

Figure S1. Processing speed and verbal knowledge across the two experiments. **a** Identical Picture test first experiment. **b** Spot-a-Word test first experiment. **c** Identical Picture test second experiment. **d** Spot-a-Word test second experiment. Abbreviations: children (CH), adolescents (AD), younger adults (YA), older adults (OA).

Across both laboratory experiments, participants completed the Identical Picture test measuring processing speed and the Spot-a-Word test measuring verbal knowledge (Lindenberger et al., 1993). As detailed below, our results are in line with previous studies showing age-related declines in fluid abilities and age-related increases in crystallized abilities (Li et al., 2004).

In the first experiment (Fig. S1a), an ANOVA showed significantly different Identical Picture scores between the age groups ($H = 74.034$, $p < 0.001$). Children (median = 20.0, IQR 18.0 to 22.0) had significantly lower scores compared to adolescents (median = 29.0, IQR 27.0 to 32.0; $u = 90.0$, $p < 0.001$, $R^2 = 0.483$) and younger adults (median = 32.0, IQR 29.0 to 36.25; $u = 31.0$, $p < 0.001$, $R^2 = 0.653$) but not compared to older adults (median = 20.0, IQR 17.5 to 24.0; $u = 564.0$, $p = 0.873$, $R^2 = 0.0$). Adolescents had significantly higher scores compared with older adults ($u = 914.0$, $p < 0.001$, $R^2 = 0.468$) but not compared to younger adults ($u = 336.0$, $p = 0.064$, $R^2 = 0.056$). Finally, younger adults had significantly higher scores compared to older adults ($u = 1068.0$, $p < 0.001$, $R^2 = 0.606$).

Regarding the Spot-a-Word test in the first experiment (Fig. S1b), an ANOVA yielded significant differences across the age groups ($H = 82.768$, $p < 0.001$). Scores of children (median = 3.0, IQR 1.0 to 6.0) were significantly lower compared to adolescents (median = 13.0, IQR 12.0 to 17.0; $u = 102.5$, $p < 0.001$, $R^2 = 0.453$), younger adults (median = 20.5, IQR 16.75 to 25.0; $u = 6.5$, $p < 0.001$, $R^2 = 0.719$), and older adults (median = 24.0, IQR 16.5 to 30.0; $u = 4.0$, $p < 0.001$, $R^2 = 0.727$). Adolescents had significantly lower scores compared to younger adults ($u = 176.0$, $p < 0.001$, $R^2 = 0.283$) and older adults ($u = 147.5$, $p < 0.001$, $R^2 = 0.367$). Finally, the test yielded a significant difference between younger and older adults ($u = 403.0$, $p = 0.049$, $R^2 = 0.058$).

In the second experiment (Fig. S1c), an ANOVA indicated significantly different Identical Picture scores ($H = 55.247$, $p < 0.001$). Children (median = 18.0, IQR 16.5 to 21.0) had significantly lower scores compared to younger adults (median = 33.0, IQR 31.0 to 35.0; $u = 0.0$, $p < 0.001$, $R^2 = 0.727$) but not compared to older adults (median = 20.0, IQR 17.0 to 23.75; $u = 382.0$, $p = 0.057$, $R^2 = 0.055$). Finally, younger adults had significantly higher scores than older adults ($u = 846.5$, $p < 0.001$, $R^2 = 0.707$).

Concerning the Spot-a-Word test in the second experiment (Fig. S1d), the age groups had significantly different scores as well ($H = 50.842$, $p < 0.001$). Children's scores (median = 7.0, IQR 4.5 to 12.0) were significantly lower compared to younger adults (median = 20.0, IQR 19.0

to 25.0; $u = 128.0$, $p < 0.001$, $R^2 = 0.325$) and compared to older adults (median = 32.0, IQR 21.0 to 35.0; $u = 41.5$, $p < 0.001$, $R^2 = 0.624$). Finally, younger adults had significantly lower scores compared to older adults ($u = 160.5$, $p < 0.001$, $R^2 = 0.278$).

Catch-trial control analysis

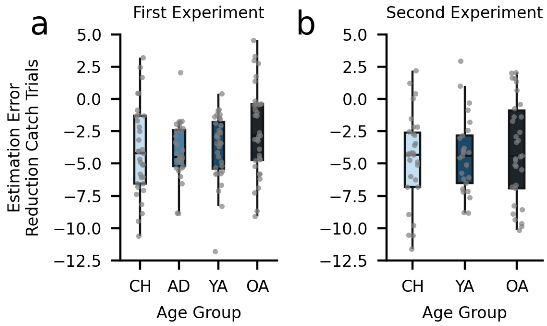

Figure S2. Catch-trial control analysis. **a** First experiment. **b** Second experiment. Abbreviations: children (CH), adolescents (AD), younger adults (YA), older adults (OA).

To test whether children and older adults can perform the task well without ambiguity about the helicopter location (catch trials), we compared estimation errors between catch trials and regular trials (with hidden helicopter). In the first experiment, all age groups showed significantly lower estimation errors on catch trials than on regular trials (Fig. S2a). In children, median = -4.189, IQR = -6.547 to -1.301, w = 34.0, $p < 0.001$, $R^2 = 0.588$. In adolescents, median = -4.494, IQR = -5.216 to -2.418, w = 5.0, $p < 0.001$, $R^2 = 0.728$. In younger adults, median = -3.393, IQR = -5.376 to -1.796, w = 1.0, $p < 0.001$, $R^2 = 0.756$. Finally, in older adults, median = -2.842, IQR = -4.727 to -0.405, w = 99.0, $p < 0.001$, $R^2 = 0.358$. The ANOVA did not yield significant differences between the groups ($H = 5.7$, $p = 0.127$), suggesting comparable performance levels.

Similarly, in the second experiment, all age groups showed lower estimation errors on catch trials (Fig. S2b). In children, median = -4.31, IQR = -6.811 to -2.618, w = 13.0, $p < 0.001$, $R^2 = 0.684$. In younger adults, median = -4.378, IQR = -6.522 to -2.823, w = 10.0, $p < 0.001$, $R^2 = 0.673$. In older adults, median = -4.424, IQR = -6.903 to -0.911, w = 74.0, $p < 0.001$, $R^2 = 0.429$. The ANOVA did not yield significant differences between the groups ($H = 0.532$, $p = 0.766$), suggesting similar performance levels as well.

Taken together, these results indicate that age groups had similar abilities to improve their performance based on explicit information about the mean of the outcome-generating distribution. This suggests a comparable understanding of the task instructions, similar motivation, and necessary visual and motor capabilities.

Supplementary results second experiment

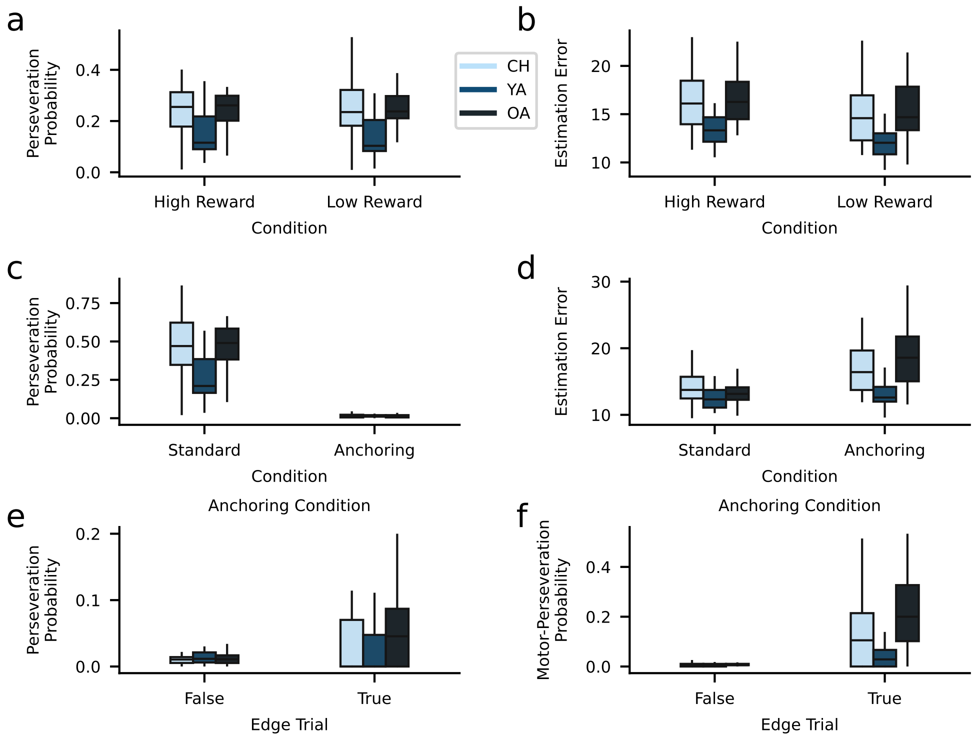

Figure S3. Supplementary results second experiment. a Estimated perseveration probability in the reward conditions. **b** Estimation errors in the reward conditions. **c** Estimated perseveration probability in the standard and anchoring conditions. **d** Estimation error in the standard and anchoring conditions. **e** Estimated perseveration probability plotted separately for regular and edge trials. **f** Estimated motor-perseveration probability plotted independently for regular and edge trials. Abbreviations: children (CH), younger adults (YA), older adults (OA).

According to Wilcoxon sign-rank tests, no age group showed a significant perseveration-probability difference between the high- and low-reward condition (Fig. S3a). In children, median = 0.007, IQR -0.022 to 0.03, w = 214.0, p = 0.505, R^2 = 0.014. In younger adults, median = 0.004, IQR -0.017 to 0.022, w = 144.0, p = 0.619, R^2 = 0.01. In older adults, median = -0.004, IQR -0.015 to 0.028, w = 291.0, p = 0.912, R^2 = 0.0. Because our primary interest in the second experiment was to explore which task manipulations were associated with decreased perseveration, we did not analyze the difference between high- and low-incentive trials in detail but focused on the difference between the standard and anchoring condition. For the sake of completeness, we here provide the age-related estimation-error differences between the high- and low-reward trials (Fig. S3b). In all age groups, Wilcoxon sign-rank tests yielded a significant estimation-error difference between the high- and low-reward condition (children: median = 1.177, IQR 0.304 to 2.406, w = 57.0, $p < 0.001$; younger adults: median = 1.289, IQR 0.967 to 2.145, w = 25.0, $p < 0.001$; older adults: median = 1.069, IQR 0.508 to 2.032, w = 37.0, $p < 0.001$).

We next provide additional details on the differences in the estimated perseveration probabilities in the standard compared to the anchoring condition. According to Wilcoxon sign-rank tests, children (median = -0.45, IQR -0.592 to -0.335, w = 0.0, $p < 0.001$, R^2 = 0.762), younger adults (median = -0.19, IQR -0.37 to -0.15, w = 0.0, $p < 0.001$, R^2 = 0.765), and older adults (median = -0.48, IQR -0.576 to -0.352, w = 0.0, $p < 0.001$, R^2 = 0.761) perseverated significantly less in the anchoring condition compared to the standard condition (Fig. S3c).

Moreover, children (median = 3.671, IQR 0.622 to 4.774, w = 72.0, p = 0.001, R^2 = 0.384) and older adults (median = 3.636, IQR 1.312 to 8.229, w = 42.0, $p < 0.001$, R^2 = 0.561) showed higher estimation errors in the anchoring condition compared with the standard condition. In younger adults, this test did not yield a significant estimation-error difference between the conditions (median = 0.107, IQR -0.773 to 1.088, w = 139.0, p = 0.527, R^2 = 0.016) (Fig. S3d).

Finally, we zoomed into the perseveration trials in the anchoring condition by dividing the trials into "edge" and regular trials. Edge trials reflect trials in which the starting position of the bucket was either $z_t = 0$ (left edge of the screen) or $z_t = 300$ (right edge of the screen). On these trials, the anchoring effect could be different from regular trials because the shift in the location could not affect the starting position beyond the edge. As shown in Fig. S3e, perseveration was practically absent in regular trials and thus primarily occurred on edge trials, suggesting that the few cases of perseveration (cf Fig. S3c, anchoring condition) were driven by edge trials in which the anchoring effect was truncated. A similar result was observed for motor-perseveration trials (Fig. S3f). Motor perseveration refers to the prediction in the initial bucket location (see SM Data preprocessing), which was absent in regular trials but occasionally took place on edge trials. Together, these results support the conclusion that our anchoring manipulation strongly reduced the perseverative behavior of our participants.

Average learning rate second experiment

We tested which factors determined participants' belief updates in the second experiment. Similar to the first experiment, we analyzed average learning rates across all trials, indicating to which extent prediction errors influenced the belief update on average. The ANOVA showed that learning rates were significantly different between the age groups ($H = 21.951$, $p < 0.001$). The post-hoc tests indicated significantly lower learning rates in children (median = 0.73, IQR 0.65 to 0.838) compared to younger adults (median = 0.884, IQR 0.86 to 0.923; $u = 131.0$, $p < 0.001$, $R^2 = 0.318$) but not compared to older adults (median = 0.772, IQR 0.699 to 0.854; $u = 466.0$, $p = 0.427$, $R^2 = 0.01$). Similarly, younger adults had significantly higher learning rates compared to older adults ($u = 683.0$, $p < 0.001$, $R^2 = 0.264$). As in the first experiment, these results show that children and older adults use lower learning rates than younger adults, i.e., they use prediction errors to a lesser degree for updating their beliefs about the helicopter.

Computational modeling

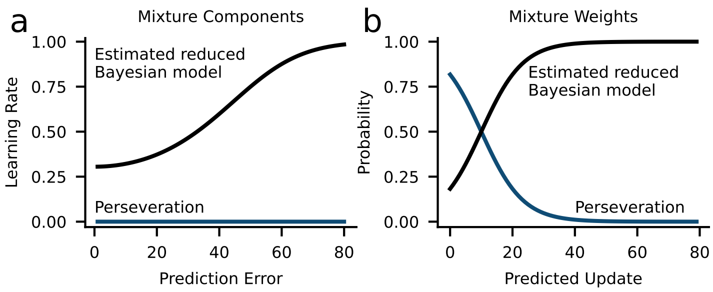

Figure S4. Mixture model. To account for age-related differences in perseveration within our model, the model contained a perseveration component that allowed for a better explanation of participant data than a control model without perseveration. **a** In our task, perseveration can computationally be described as a learning rate $\alpha_t = 0$, which leads the model to ignore the prediction error and, consequently, to an exact repetition of the previous prediction. **b** We modeled the contribution of this perseveration component (perseveration probability) as a function of the predicted update of the reduced Bayesian model. In effect, the model predicted a higher perseveration probability for smaller prediction errors.

Supplementary modeling results first experiment

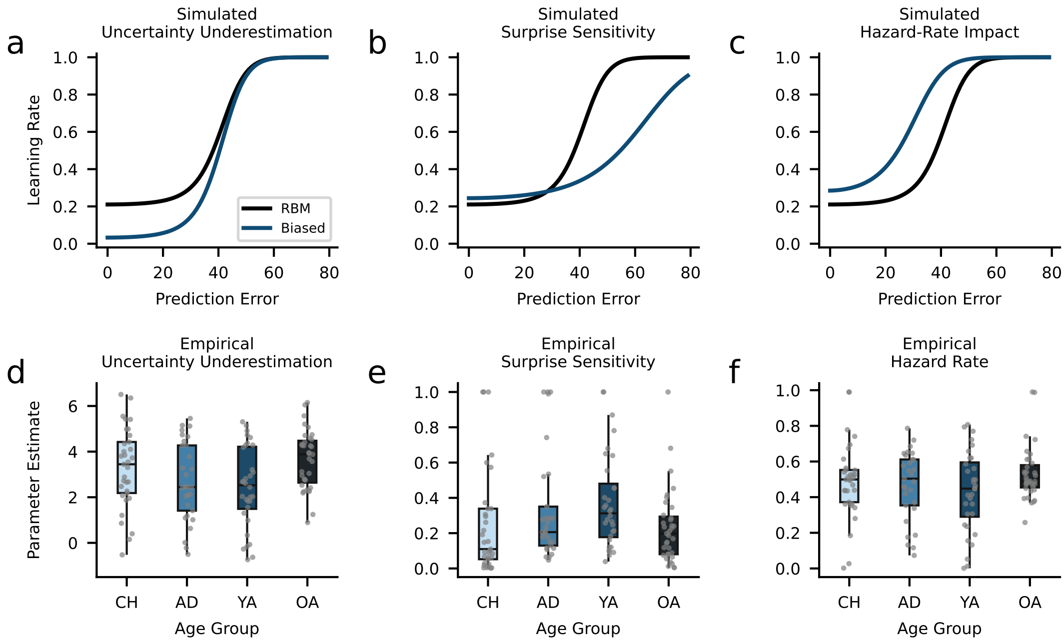

Figure S5. Hypothetical and estimated learning deficits resulting from specific deviations from the reduced Bayesian learning model. a-c The reduced Bayesian model (RBM) uses a high learning rate in response to large prediction errors that are more likely due to change points. In the range of small prediction errors more likely due to variability, it uses a lower learning rate to average out the variability in the outcomes. We hypothesized that learning deficits might stem from specific deviations from the learning-rate computations of the reduced Bayesian model, including an underestimation of uncertainty, lower surprise sensitivity, and a subjectively high hazard rate (Biased). **a** An underestimation of uncertainty leads to a specific learning deficit in the range of small prediction errors unlikely to be due to change points. That is, the model has a reduced ability to average out the variability in the outcomes. Insensitivity to surprise leads to a learning deficit in the range of large prediction errors that are likely to be due to change points. Consequently, the model has a reduced ability to respond to change points appropriately. **d** The hazard rate reflects the model's assumed frequency of change points. A subjectively higher hazard rate, i.e., overestimating how often change points occur, is associated with higher learning rates in the range of small and intermediate prediction errors. **d-f** Estimated parameters reflecting **d** uncertainty underestimation, **e** surprise sensitivity, and **f** subjective hazard rate for children (CH), adolescents (AD), younger adults (YA), and older adults (OA).

To better understand the computational underpinnings of differences in learning, we analyzed trial-by-trial dynamics of learning-rate adjustments in the different age groups using the mixture model. The uncertainty-underestimation parameter reflects the extent to which participants underestimate uncertainty during learning. Higher uncertainty-underestimation parameters lead to lower learning rates in the range of small prediction errors, which are most likely due to environmental variability (Fig. S5a). The mostly positive parameter estimates reflecting uncertainty underestimation shown in Fig. S5d suggest that all age groups tended to underestimate uncertainty because a full consideration of uncertainty is indicated with an uncertainty-underestimation parameter equal to 0. We next used an ANOVA to compare uncertainty underestimation across the four age groups. The test yielded significant parameter differences between the age groups ($H = 8.836$, $p = 0.032$). Older adults (median = 3.879, IQR 2.639 to 4.478) showed significantly higher uncertainty underestimation compared to adolescents (median = 2.443, IQR 1.411 to 4.271; $u = 358.0$, $p = 0.044$, $R^2 = 0.063$) and compared to younger adults (median = 2.524, IQR 1.49 to 4.214; $u = 348.0$, $p = 0.008$, $R^2 = 0.105$). That is, in line with our previous work, the results suggest reduced uncertainty-driven learning in older adults compared to their younger counterparts. However, the test did not yield a significant difference between children (median = 3.445, IQR 2.184 to 4.423) compared to adolescents (u = 568.0, $p = 0.209$, $R^2 = 0.025$), younger adults ($u = 668.0$, $p = 0.067$, $R^2 = 0.052$), and older adults ($u = 503.0$, $p = 0.364$, $R^2 = 0.012$). The test did not show a significant difference between adolescents and younger adults ($u = 506.0$, $p = 0.549$, $R^2 = 0.006$).

The surprise-sensitivity parameter indicates the influence of change-point probability on the learning rate. Lower surprise sensitivity reflects lower learning rates in the range of large prediction errors that are most likely due to change points (Fig. S5b). The surprise-sensitivity estimates shown in Fig. S5e suggest that all age groups considered change-point probability, albeit not as much as the reduced Bayesian model, which has a surprise sensitivity of 1. Moreover, the ANOVA indicated significantly different median surprise-sensitivity parameters between the age groups ($H = 10.195$, $p = 0.017$). Younger adults (median = 0.312, IQR 0.177 to 0.48) showed significantly higher surprise sensitivity compared to children (median = 0.11, IQR 0.052 to 0.339; $u = 324.0$, $p = 0.008$, $R^2 = 0.11$) and compared to older adults (median = 0.198, IQR 0.081 to 0.294; $u = 762.0$, $p = 0.011$, $R^2 = 0.095$). The test did not yield a significant difference between children and adolescents (median = 0.206, IQR 0.129 to 0.35; $u = 349.5$, $p = 0.07$, $R^2 = 0.053$), children and older adults ($u = 515.5$, $p = 0.45$, $R^2 = 0.008$), adolescents and younger adults ($u = 385.0$, $p = 0.257$, $R^2 = 0.021$), and between adolescents and older adults ($u = 596.5$, $p = 0.233$, $R^2 = 0.022$).

The hazard-rate parameter models assumptions about the frequency of change points. A higher hazard rate reflects the assumption that change points occur more frequently in the task and leads to higher levels of change-point probability and, thus, higher learning rates for small and intermediate prediction errors (Fig. S5c). The true but unknown hazard rate was approximately 0.1, i.e., the hazard-rate parameter estimates shown in Fig. S5f suggest that most participants overestimated the frequency of change points in the task. However, the ANOVA ($H = 2.573$, $p = 0.462$) did not yield significant hazard-rate parameter differences between children (median = 0.499, IQR 0.372 to 0.553), adolescents (median = 0.504, IQR 0.354 to 0.612), younger adults (median = 0.448, IQR 0.29 to 0.595), and older adults (median = 0.49, IQR 0.454 to 0.58).

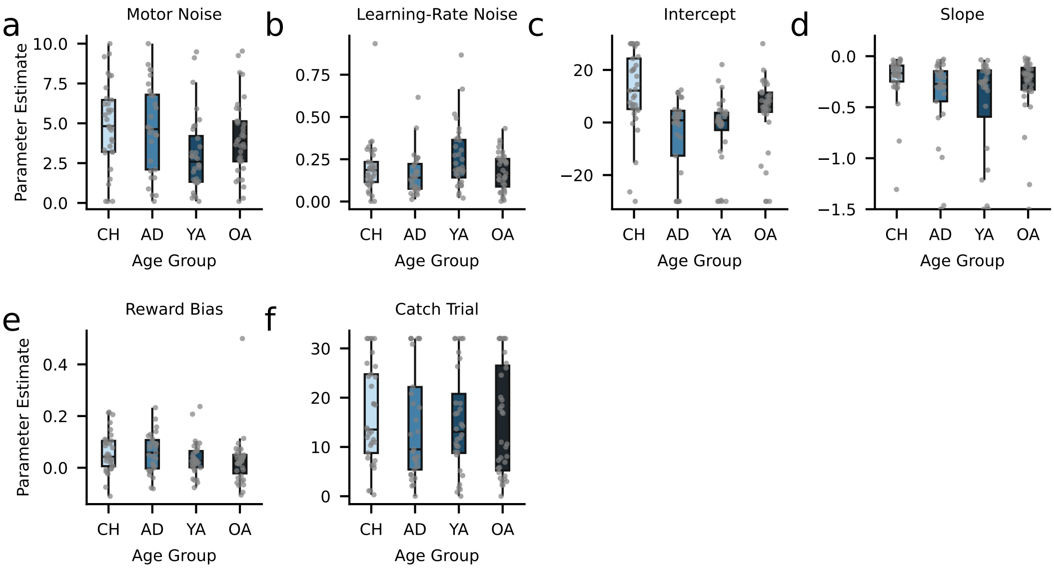

Figure S6. Additional model-parameter estimates (first experiment). **a** Motor noise. **b** Learning-rate noise. **c** Intercept of the logistic function. **d** Slope of the logistic function. **e** Reward-bias parameter. **f** Catch-trial parameter. Abbreviations: children (CH), adolescents (AD), younger adults (YA), older adults (OA).

Here, we also present the additional parameters of the mixture model. The motor-noise parameter o_0 modeled prediction noise that was independent of the update magnitude (Fig. S6a). An ANOVA yielded a significant difference between the groups ($H = 10.661$, $p = 0.014$). The test showed significantly lower motor noise in younger adults (median = 2.591, IQR 1.328 to 4.203) compared to children (median = 4.82, IQR 3.214 to 6.465; $u = 752.5$, $p = 0.003$, $R^2 = 0.133$), adolescents (median = 4.631, IQR 2.099 to 6.794; $u = 623.5$, $p = 0.022$, $R^2 = 0.086$), and older adults (median = 3.868, IQR 2.61 to 5.129; $u = 395.5$, $p = 0.04$, $R^2 = 0.063$). The test did not yield a significant difference between children and adolescents ($u = 494.0$, $p = 0.832$, $R^2 = 0.001$), children and older adults ($u = 690.5$, $p = 0.167$, $R^2 = 0.028$), and adolescents and older adults ($u = 585.5$, $p = 0.296$, $R^2 = 0.017$).

The learning-rate noise parameter o_1 modeled prediction noise depending on the update magnitude (Fig. S6b). The ANOVA did not indicate significant differences between children (median = 0.187, IQR 0.115 to 0.234), adolescents (median = 0.141, IQR 0.075 to 0.222), younger adults (median = 0.197, IQR 0.142 to 0.365), and older adults (median = 0.167, IQR 0.088 to 0.252) ($H = 6.785$, $p = 0.079$).

The intercept parameter b_0 of the logistic function was indicative of perseveration (Fig. S6c). An ANOVA yielded significant median intercept differences between the age groups ($H = 29.834$, $p < 0.001$). Post-hoc tests indicated a significant difference between children (median = 12.117, IQR 5.114 to 24.347) and adolescents (median = 0.828, IQR -12.628 to 4.456; $u = 773.0$, $p < 0.001$, $R^2 = 0.277$), children and younger adults (median = 0.856, IQR -2.886 to 3.695; $u = 849.0$, $p < 0.001$, $R^2 = 0.272$), children and older adults (median = 7.49, IQR 4.035 to 11.389; $u = 745.5$, $p = 0.04$, $R^2 = 0.062$), adolescents and older adults ($u = 276.0$, $p = 0.002$, $R^2 = 0.152$), younger adults and older adults ($u = 301.0$, $p = 0.001$, $R^2 = 0.157$), but not between adolescents and younger adults ($u = 454.0$, $p = 0.891$, $R^2 = 0.0$).

Perseveration was also reflected in the slope parameter b_1 of the logistic function (Fig. S6d). An ANOVA did not yield significant median slope differences between the age groups ($H = 7.068$, $p = 0.07$). In children, median = -0.167, IQR -0.248 to -0.092. In adolescents, median

= -0.274, IQR -0.443 to -0.147. In younger adults, median = -0.256, IQR -0.594 to -0.14. In older adults, median = -0.195, IQR -0.328 to -0.114. The combination of the intercept and slope parameters is reflected in higher perseveration probabilities in children and older adults, which are especially evident in the range of small and intermediate updates.

The reward-bias parameter q modeled the influence of high versus low reward on the learning rate (Fig. S6e). In children, median = 0.042, IQR 0.006 to 0.104. In adolescents, median = 0.058, IQR -0.003 to 0.107. In younger adults, median = 0.026, IQR 0.002 to 0.065. In older adults, median = 0.025, IQR -0.022 to 0.05. An ANOVA did not yield a significant difference between the age groups ($H = 6.288$, $p = 0.098$).

Finally, the catch-trial parameter σ_H modeled the influence of catch trials on the update magnitude (Fig. S6f) In children, median = 13.53, IQR 8.774 to 24.748. In adolescents, median = 9.508, IQR 5.417 to 22.161. In younger adults, median = 13.082, IQR 8.779 to 20.771. In older adults, median = 10.963, IQR 5.228 to 26.513. An ANOVA did not yield significant differences between the age groups ($H = 0.873$, $p = 0.832$).

Control model

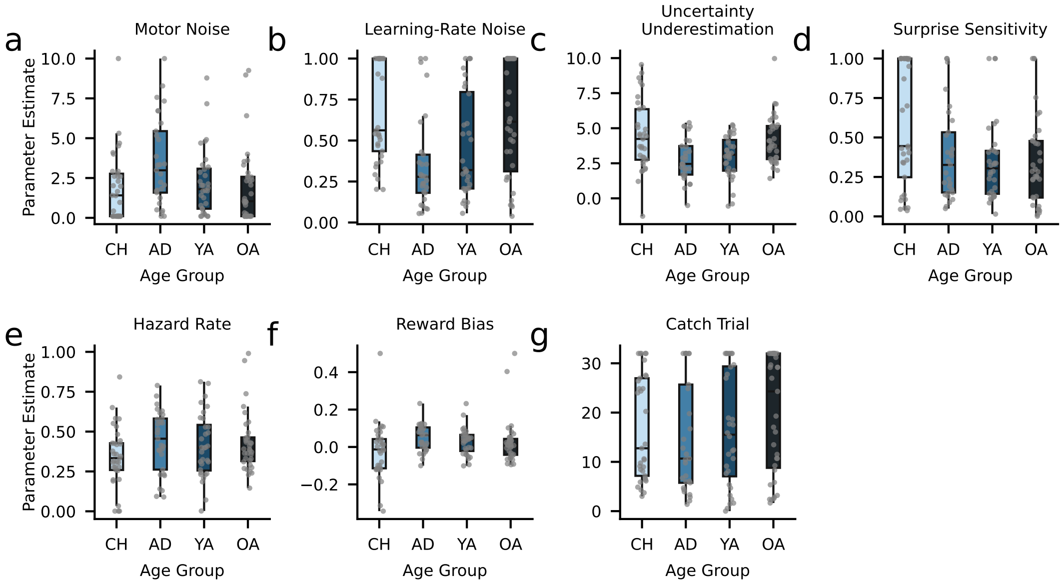

Figure S7. Model-parameter estimates of the control model (first experiment). **a** Motor noise. **b** Learning-rate noise. **c** Uncertainty underestimation. **d** Surprise sensitivity. **e** Hazard rate. **f** Reward-bias parameter. **g** Catch-trial parameter. Abbreviations: children (CH), adolescents (AD), younger adults (YA), older adults (OA).

In our previous study, we focused on uncertainty underestimation in older compared to younger adults using a model without perseveration component (Nassar et al., 2016). To examine if the age-related differences in uncertainty underestimation might be more clearly pronounced in a model without perseveration, we here additionally examined uncertainty underestimation in our control model.

In Fig. S7, we show the parameter estimates of the control model without perseveration, which fit the data worse compared to the mixture model that captured perseveration (see Model comparison). To test if the uncertainty-underestimation parameters in this model were different between the age groups, we also applied an ANOVA ($H = 14.409$, $p = 0.002$). Post-hoc tests showed a significant difference between children (median = 4.238, IQR 2.753 to 6.361) and adolescents (median = 2.456, IQR 1.692 to 3.728; $u = 675.0$, $p = 0.006$, $R^2 = 0.123$), children and younger adults (median = 3.108, IQR 1.972 to 4.167; $u = 701.0$, $p = 0.024$, $R^2 = 0.079$), adolescents and older adults (median = 3.732, IQR 2.797 to 5.16; $u = 277.0$, $p = 0.002$, $R^2 = 0.15$), and younger and older adults ($u = 369.0$, $p = 0.017$, $R^2 = 0.085$). The test did not yield a significant difference between children and older adults ($u = 604.0$, $p = 0.75$, $R^2 = 0.001$), and adolescents and younger adults ($u = 414.0$, $p = 0.475$, $R^2 = 0.008$).

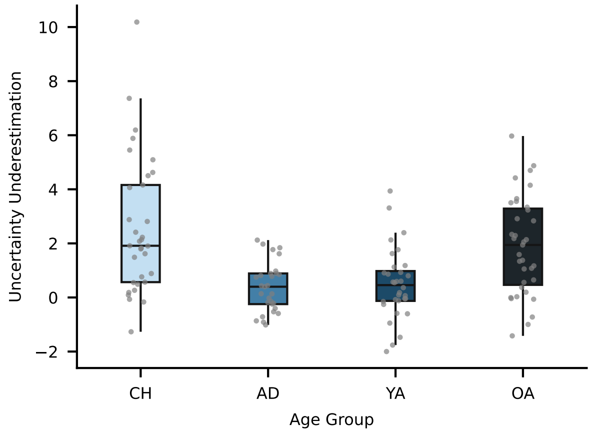

Figure S8. Uncertainty underestimation in simulation-based control analysis. We examined whether perseveration can be attributed to uncertainty underestimation in a control model that does not include the perseveration component. We first performed a simulation study using the reduced Bayesian model with perseveration, where the parameter values were based on the empirical model-based estimates. Therefore, perseveration probability was comparable between the simulated and empirical data sets (including the age-related differences with more perseveration in children and older adults). Crucially, in all simulations, we set the uncertainty-underestimation parameter to zero. We then fitted the control model to the simulated data to examine whether the uncertainty-underestimation parameter picks up variance that is due to perseveration. The uncertainty-underestimation parameters are higher in simulated children (CH) and older adults (OA) compared to younger adults (YA), although the ground-truth parameters were zero, indicating that perseveration (which was higher in simulated children and older adults) can give rise to the appearance of uncertainty underestimation when perseveration is not modeled explicitly.

These results suggest that the perseveration component of our best-fitting model partly captures variance that the uncertainty-underestimation parameter captures in the control model without perseveration. Both uncertainty underestimation and perseveration lead to an under-adjustment of predictions. Uncertainty underestimation leads to an update that is smaller than prescribed by the normative version of the reduced Bayesian model, while perseveration leads to a no-update altogether. Moreover, both effects occur primarily after small to intermediate prediction errors. To test this idea more directly, we next simulated data based on the model that included the perseveration component. The parameter values of the uncertainty-underestimation parameter were 0 (indicating no underestimation) for all simulated participants, and all other values corresponded to the model-based estimates obtained from model fitting (including age-related difference in preservation). We then fitted the control model without perseveration component to this simulated data set and, in line with our argument, found that uncertainty underestimation was larger in simulated children and older adults compared to the other two groups (Fig. S8).

Model comparison

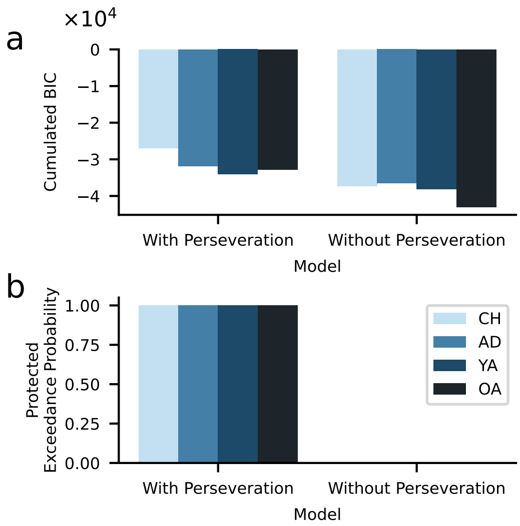

Figure S9. Model comparison. **a** Cumulated Bayesian information criterion (BIC) of the reduced Bayesian model with perseveration and the reduced Bayesian model without perseveration, plotted separately for each age group. Higher BIC values indicate a better model fit, that is, in all age groups, the reduced Bayesian model with perseveration fits the data better than the model without perseveration. **b** This is supported by the larger protected exceedance probabilities (pEP, all > 0.99) of the model with perseveration. Abbreviations: children (CH), adolescents (AD), younger adults (YA), older adults (OA).

We used the cumulated Bayesian information criterion (BIC; Stephan et al., 2009) and the corresponding protected exceedance probabilities (Rigoux et al., 2014) for model comparison (see Model comparison). Here, higher BIC values indicate a better model fit (Fig. S9a). We used protected exceedance probabilities for model comparison on the group level, which indicates the probability that a particular model is more likely than any other model of the model space (Fig. S9b). Together, this indicated that in each age group, the mixture model fits the data better than the reduced Bayesian model without perseveration. We finally counted the number of participants, in which the mixture model fits the data better than the reduced Bayesian model without perseveration. This indicated that in 121 out of 129 participants, the model comparison favored the mixture model.

Supplementary modeling results second experiment

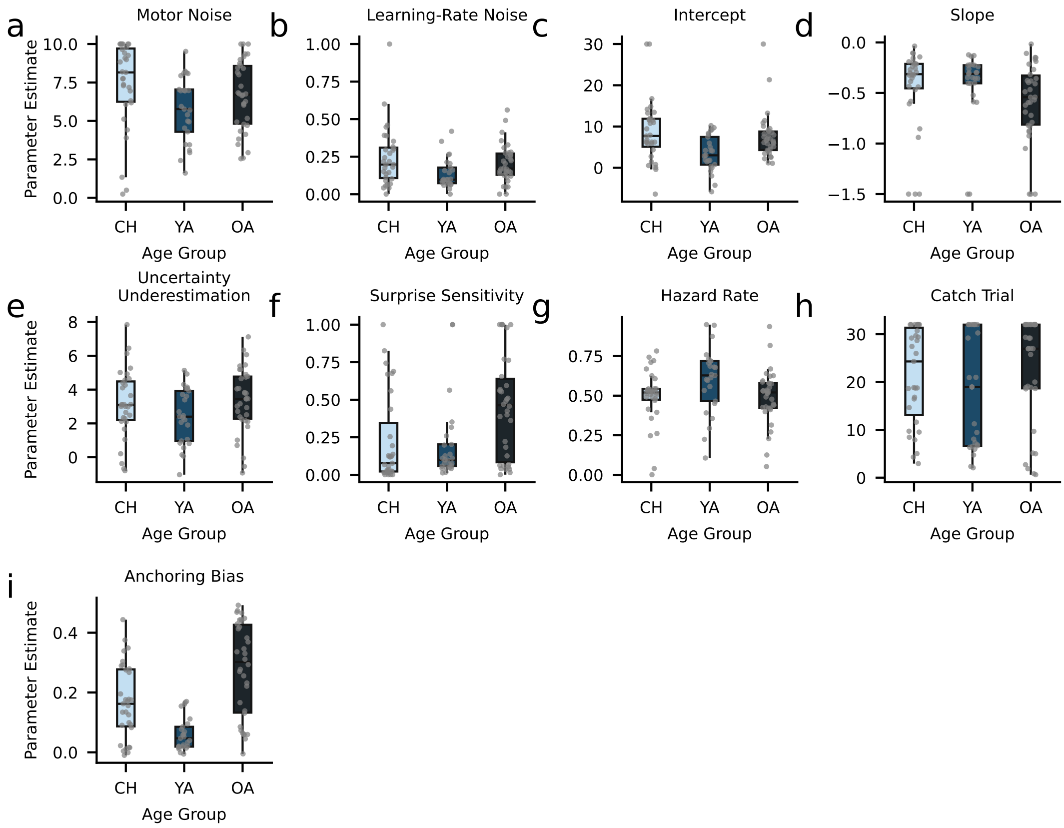

Figure S10. Model-parameter estimates (second experiment). **a** Motor noise. **b** Learning-rate noise. **c** Intercept of the logistic function. **d** Slope of the logistic function. **e** Uncertainty underestimation. **f** Surprise sensitivity. **g** Hazard rate. **h** Catch-trial parameter. **i** Anchoring bias. Abbreviations: children (CH), younger adults (YA), older adults (OA).

To confirm that age-related differences in anchoring capture performance differences beyond the computational factors of our model, we added an anchoring-bias parameter. In this extended version, the model's prediction of the next outcome was biased towards the anchor. This bias was governed by an additional free parameter, where higher parameter values indicate a stronger bias (Fig. S10i). According to an ANOVA, the parameters were significantly different between the age groups ($H = 27.835$, $p < 0.001$). Post-hoc tests indicated a significant difference between children (median = 0.162, IQR 0.087 to 0.277) and younger adults (median = 0.047, IQR 0.019 to 0.085; $u = 583.0$, $p = 0.001$, $R^2 = 0.184$), children and older adults (median = 0.302, IQR 0.133 to 0.426; $u = 323.0$, $p = 0.008$, $R^2 = 0.11$), and between younger and older adults ($u = 101.0$, $p < 0.001$, $R^2 = 0.417$).

Moreover, an ANOVA yielded a significantly different motor-noise parameter o_0 across the age groups ($H = 9.132$, $p = 0.01$; Fig. S10a). Children (median = 8.156, IQR 6.245 to 9.711) showed significantly higher motor noise compared to younger adults (median = 5.77, IQR 4.292 to 7.046; $u = 565.0$, $p = 0.003$, $R^2 = 0.152$), but not compared to older adults (median = 6.72, IQR 4.819 to 8.572; $u = 664.5$, $p = 0.071$, $R^2 = 0.05$). The test did not yield a significant difference between younger and older adults ($u = 330.0$, $p = 0.147$, $R^2 = 0.036$).

An ANOVA yielded a significantly different learning-rate noise parameter o_1 between the age groups (H = 6.659, p = 0.036). Post-hoc tests showed a significant difference between children (median = 0.198, IQR 0.107 to 0.31) and younger adults (median = 0.097, IQR 0.073 to 0.178; u = 527.5, p = 0.021, R^2 = 0.094) and between younger adults and older adults (median = 0.171, IQR 0.129 to 0.271; u = 282.0, p = 0.029, R^2 = 0.081). The test did not indicate a significant difference between children and older adults (u = 564.0, p = 0.632, R^2 = 0.004) (Fig. S10b).

The ANOVA also yielded a significant age-related difference in the intercept parameter b_0 of the logistic function (H = 11.163, p = 0.004). Follow-up tests indicated a significant difference between children (median = 7.704, IQR 5.083 to 11.877) and younger adults (median = 3.015, IQR 0.749 to 7.464; u = 566.0, p = 0.003, R^2 = 0.154) and between younger and older adults (median = 6.547, IQR 4.283 to 8.789; u = 243.0, p = 0.005, R^2 = 0.131). The test did not indicate a significant difference between children and older adults (u = 603.0, p = 0.321, R^2 = 0.015) (Fig. S10c).

With respect to the slope of the logistic function b_1, an ANOVA showed significant differences between the age groups (H = 6.532, p = 0.038). Children (median = -0.314, IQR -0.454 to -0.213) had significantly higher slopes than older adults (median = -0.542, IQR -0.813 to -0.327; u = 684.0, p = 0.04, R^2 = 0.065). Younger adults (median = -0.311, IQR -0.404 to -0.226) had significantly higher slopes than older adults (u = 575.0, p = 0.022, R^2 = 0.089). Finally, the slopes of children and younger adults were not significantly different (u = 404.0, p = 0.792, R^2 = 0.001) (Fig. S10d).

For the uncertainty-underestimation parameter u, an ANOVA did not indicate significant differences between the age groups (H = 3.847, p = 0.146). In children, median = 3.102, IQR 2.204 to 4.483. In younger adults, median = 2.401, IQR 0.963 to 3.924. In older adults, median = 3.456, IQR 2.282 to 4.77 (Fig. S10e).

Moreover, the test yielded a significant difference in the surprise-sensitivity parameter s (H = 10.269, p = 0.006). Post-hoc tests indicated a significant difference between children (median = 0.077, IQR 0.022 to 0.346) and older adults (median = 0.463, IQR 0.084 to 0.64; u = 312.5, p = 0.005, R^2 = 0.122), and between younger adults (median = 0.099, IQR 0.058 to 0.203) and older adults (u = 261.0, p = 0.012, R^2 = 0.107). The test did not yield a significant difference between children and younger adults (u = 336.0, p = 0.4, R^2 = 0.013) (Fig. S10f).

Next, the test yielded a significant age-related difference in the hazard-rate parameter h (H = 6.851, p = 0.033) (Fig. S10g). Children (median = 0.52, IQR 0.475 to 0.543) had significantly lower hazard-rate parameters compared to younger adults (median = 0.628, IQR 0.465 to 0.718; u = 261.0, p = 0.038, R^2 = 0.077) but not compared to older adults (median = 0.52, IQR 0.423 to 0.579; u = 539.0, p = 0.88, R^2 = 0.0). Younger adults had significantly higher hazard-rate parameters compared to older adults (u = 588.0, p = 0.013, R^2 = 0.105).

Finally, the ANOVA did not yield a significant age-related difference in the catch-trial parameter σ_H (H = 2.008, p = 0.366) between children (median = 24.305, IQR 13.135 to 31.358), younger adults (median = 18.964, IQR 6.656 to 32.0), and older adults (median = 26.98, IQR 18.673 to 32.0) (Fig. S10h).

Parameter-recovery analysis

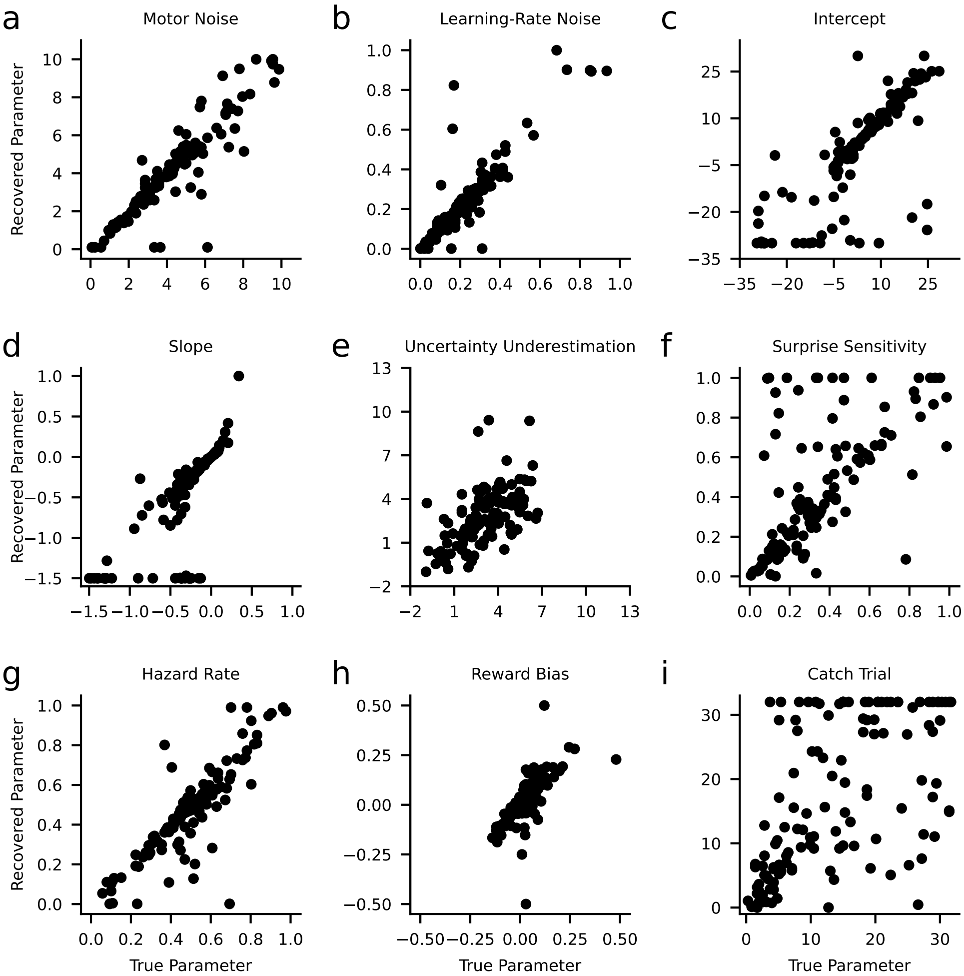

Figure S11. Parameter recovery first experiment. **a** Motor noise. **b** Learning-rate noise. **c** Intercept of the logistic function. **d** Slope of the logistic function. **e** Uncertainty underestimation. **f** Surprise sensitivity. **g** Hazard rate. **h** Reward-bias parameter. **i** Catch-trial parameter.

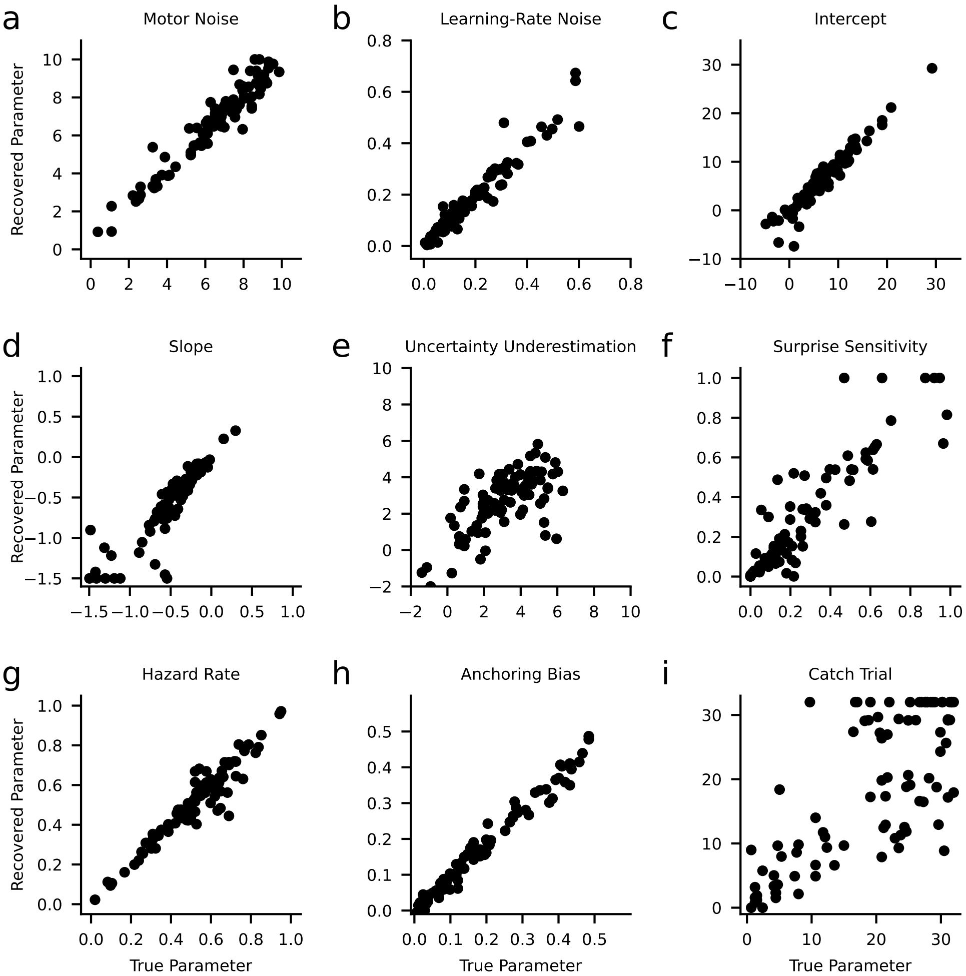

Figure S12. Parameter recovery second experiment. **a** Motor noise. **b** Learning-rate noise. **c** Intercept of the logistic function. **d** Slope of the logistic function. **e** Uncertainty underestimation. **f** Surprise sensitivity. **g** Hazard rate. **h** Anchoring bias. **i** Catch-trial parameter.

To test if the parameters of the mixture model between the reduced Bayesian model and the perseveration model can accurately be estimated, we conducted a parameter-recovery study. We used randomly selected parameters for the simulations to examine if a broad range of parameter combinations is recoverable. These parameters were based on the participants' empirical parameter distributions and additionally constrained between the approximate minimum and maximum empirical parameter estimates. The parameter-recovery results for the first (Fig. S11) and second experiment (Fig. S12) suggest that parameters can sufficiently be estimated except for the catch-trial parameter, which was often estimated at the upper boundary. This result indicates that the empirical results regarding catch trials should be treated cautiously.

Model validation

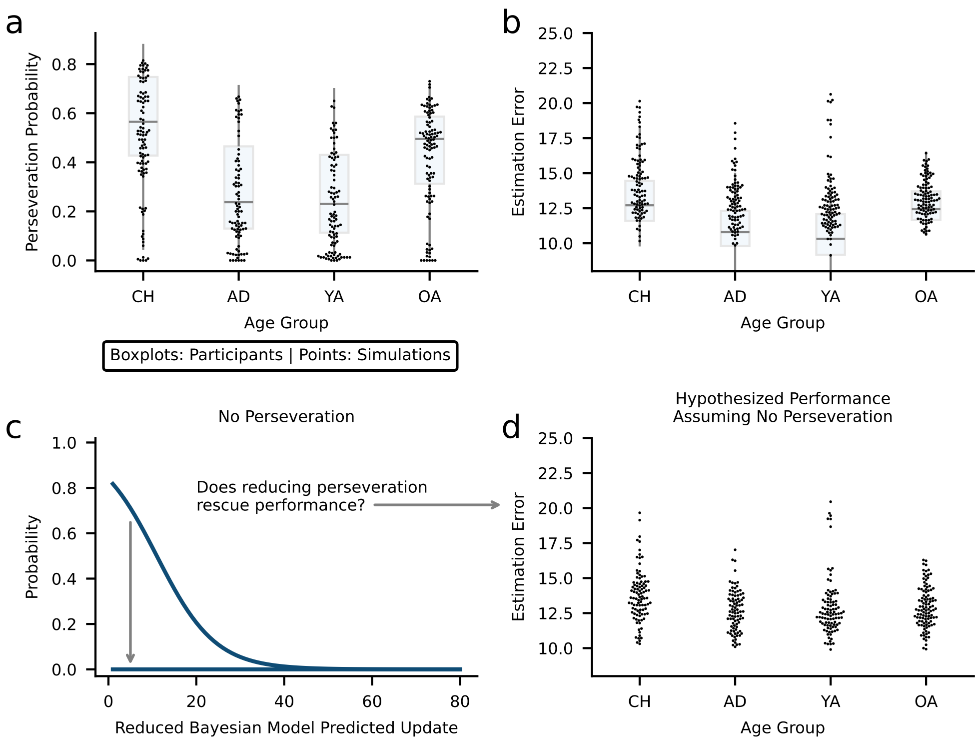

Figure S13. Simulation results. a-b To test whether the mixture model can capture age differences in perseveration and estimation errors, we simulated learning behavior for each age group. Each dot represents a simulated participant, and the boxplots show the empirical data for comparison. The results suggest that the mixture model qualitatively captures the pattern of **a** perseveration and **b** estimation errors across the age groups. **c** To test whether perseveration might contribute to age-related performance differences, we simulated learning behavior and set the perseveration probabilities to zero in all cases. **d** The simulation results suggest that, without perseveration, the four age groups would perform at similar levels. Abbreviations: children (CH), adolescents (AD), younger adults (YA), older adults (OA).

First experiment

One potential concern with using computational models as data-analytic tools is that although they make quantitative predictions of the data, they might fail to reproduce the observed data in simulations qualitatively (Palminteri et al., 2017). We used the estimated mixture model to simulate learning behavior to test if our model-based analyses accurately captured participants' behavior. First, to test if the estimated reduced Bayesian model with an additional perseveration component qualitatively reproduces participant behavior, we applied it to the same sequence of data as experienced by our participants. To obtain smoother simulation results, we repeated this procedure across three cycles. The comparison between empirical and simulated data is shown in Fig. S13a (perseveration probability) and Fig. S13b (estimation errors).

In addition, as shown in Fig. S13c,d, we also used simulations to explore if the reduced Bayesian model without perseveration would predict similar performance levels between the age groups. Here, we conducted the simulations as described above but adjusted the parameters such that perseveration probability was approximately equal to zero. These results suggested similar performance levels and motivated our second experiment that tested if age-related performance differences can be eliminated with manipulations to reduce perseveration.

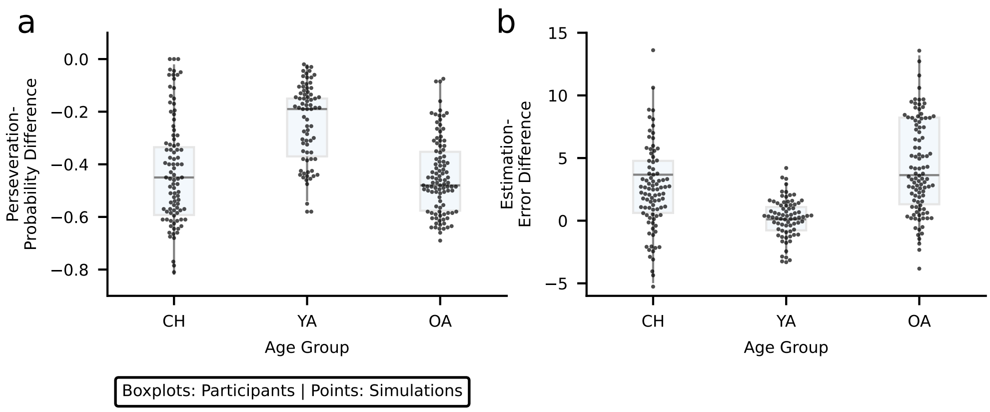

Figure S14. Model validation second experiment. To test if the reduced Bayesian model with the additional anchoring bias qualitatively captures the difference in perseveration and estimation errors between the anchoring and standard task condition, we used the model to simulate behavior across three cycles, where each point in the plot represents one simulation. **a** Comparison of the simulated and empirical perseveration probability. **b** Comparison of the simulated and empirical estimation errors. Abbreviations: children (CH), younger adults (YA), older adults (OA).

Second experiment

As in our model-validation analyses of the first experiment, we conducted the simulations with the estimated parameters and the same outcomes observed by our participants. These simulations revealed that the simulated perseveration (Fig. S14a) and estimation-error (Fig. S14b) differences qualitatively matched our empirically obtained results, which suggests that our computational model accurately captured participants' behavior.

Robust regression with age group

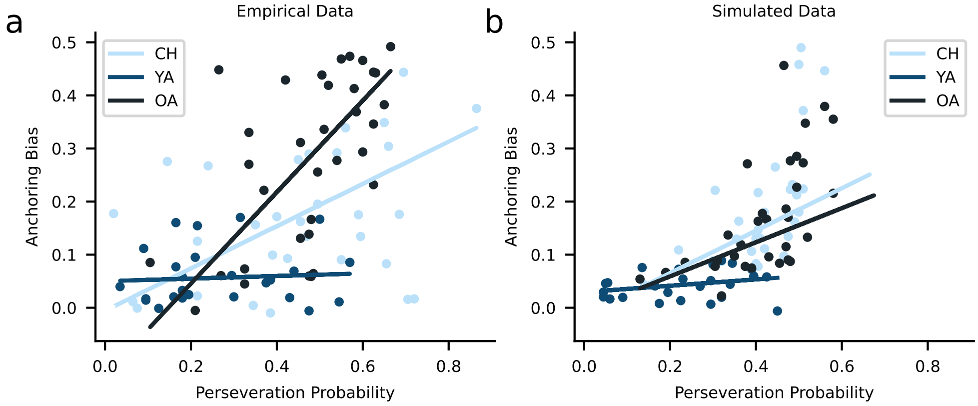

Figure S15. Robust linear regression with age group. a Empirical relationship between perseveration and anchoring in children (CH), younger adults (YA), and older adults (OA). **b** Simulations with the resource-rational sampling model recapture the relationship between perseveration and anchoring within the age groups.

We also applied a robust linear regression model with age group

$$d_i = \beta_0 + \beta_1 \hat{p}(\mathrm{n})_i + \beta_2^{ya} g_i + \beta_2^{ch} g_i + \beta_3^{ya} \hat{p}(\mathrm{n}) g_i + \beta_3^{ch} \hat{p}(\mathrm{n}) g_i \tag{30}$$

where i denotes the subject and g_i denotes age group that we coded as dummy variable with older adults as reference category. The model thus contained an additional main effect of age group with regression weight β_2 and an interaction between age group and perseveration probability with regression weight β_3.

This model revealed a significant main effect of perseveration probability $\beta_1 = 0.8617$ ($z = 4.874, p < 0.001$), and a significant interaction between perseveration probability and the level of younger adults of the age group variable β_3^{ya} = -0.8372 (z = -3.286, p = 0.001) and the level of children β_3^{ch} = -0.4635 (z = -2.173, p = 0.03). All other parameters were not significant: β_0 = -0.1269 (z = -1.453, p = 0.146), β_2^{ya} = 0.1769 (z = 1.699, p = 0.089), β_2^{ch} = 0.1213 (z = 1.146, p = 0.252). The coefficient of determination of the model was R^2 = 0.464. Hence, this robust regression with age group suggests that the anchoring bias is systematically related to perseveration and that this relationship is weaker in younger adults and children compared to older adults. For the predicted values according to this model, see Fig. S15a.

Sampling model

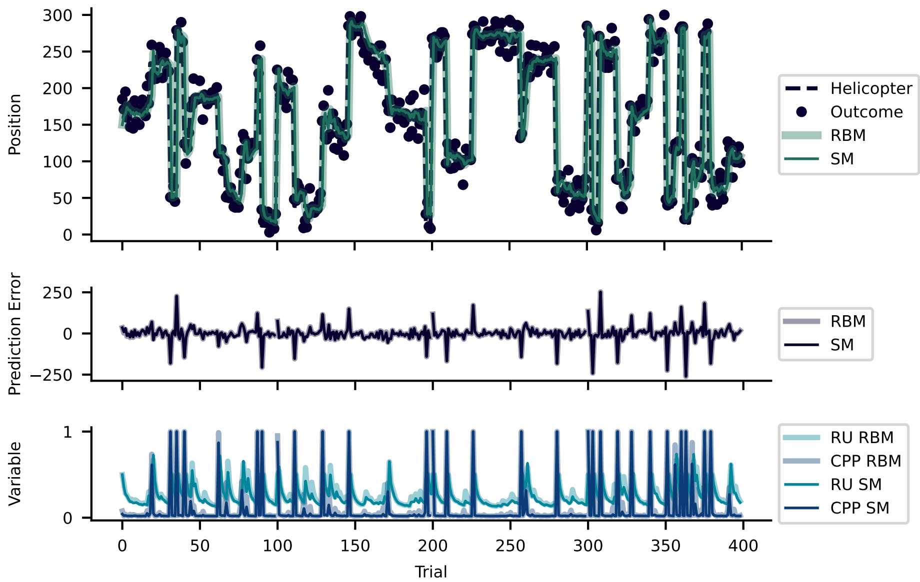

Figure S16. Comparison of the sampling model (SM) and the reduced Bayesian model (RBM). To validate the sampling model, we compared the computational variables between the sampling and the reduced Bayesian model. To approximate optimal inference, we used a large number of samples (2000 samples) without applying the stopping criterion. The comparison shows that the trial-by-trial beliefs, prediction errors, relative uncertainty, and change-point probability variables are similar between the two models. This illustrates that the sampling model can closely approximate near-optimal Bayesian inference if a sufficient number of samples is drawn.

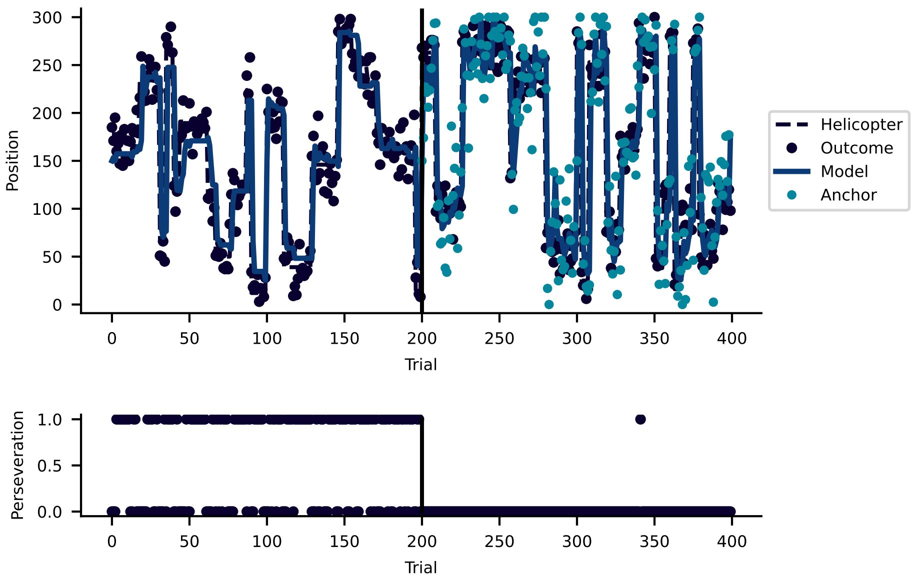

Figure S17. Illustration of the sampling model for the standard and anchoring conditions. The first 200 trials show the standard condition in which the sampling model frequently perseverates. The last 200 trials show the anchoring condition, where the anchor biases the model, but perseveration occurs rarely.

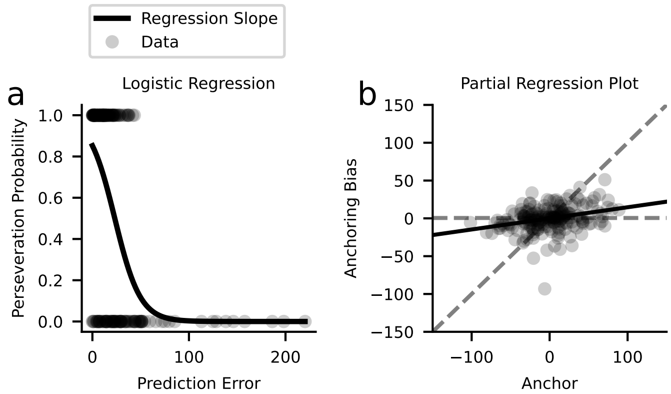

Figure S18. Illustration of perseveration and anchoring bias in the sampling model. **a** Simulated perseveration probability as a function of the prediction error. In line with our empirical results, perseveration is more likely for smaller prediction errors. **b** Simulated anchoring bias. In line with the empirical data, the model shows a stronger anchoring bias for larger anchors.

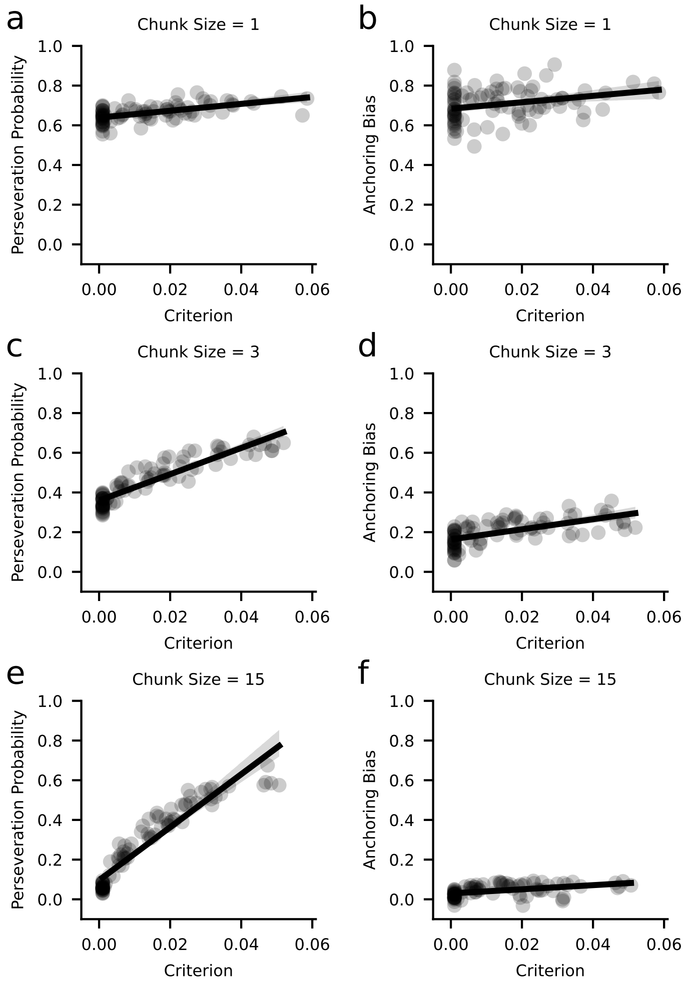

Figure S19. Joint effect of chunk size and criterion parameter on perseveration probability and anchoring bias. We illustrate the effects of varying parameter values to show how different parameter settings impact perseveration and anchoring. In particular, the plot demonstrates how the criterion parameter influences the biases under three different chunk-size settings. These simulation results reveal a general trend where both perseveration and anchoring tend to increase as the criterion values become larger. Moreover, the examples demonstrate that the effects of the parameters interact to some extent. Illustration of **a** perseveration and **b** anchoring for lower chunk-size parameters (chunk size = 1). Here, perseveration and the anchoring bias increase as a function of a larger criterion. When chunk size is larger (chunk size = 3), **c** perseveration and **d** anchoring increase as a function of the criterion, but the overall magnitude of the biases is lower. **e** Finally, when chunk size is considerably larger (chunk size = 15), perseveration still increases, **f** but the anchoring bias almost remains constant. This corresponds to a floor effect, where the larger number of samples overwrites the impact of the anchor. Consequently, these examples suggest a diminishing correlation between perseveration and anchoring as the chunk-size parameter increases from lower values (like in children and older adults) to higher values (like in younger adults).

Parameter comparison

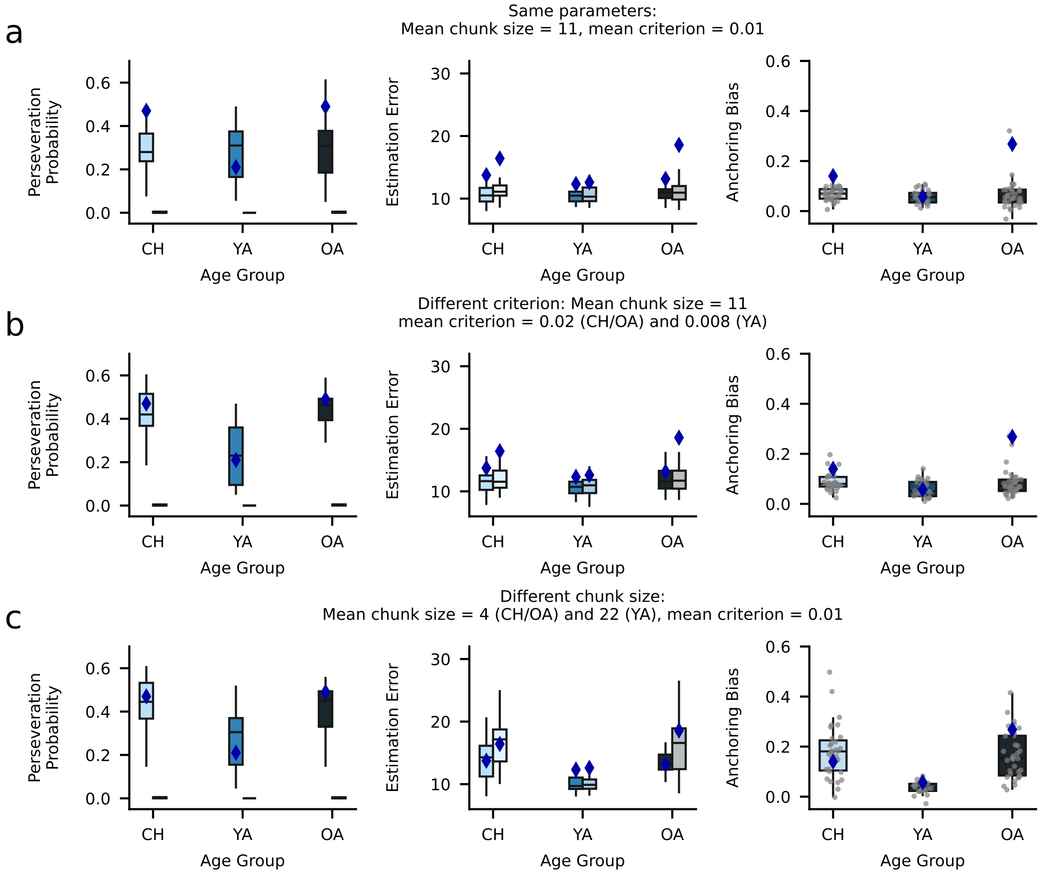

Figure S20. Parameter comparison. Simulations with different parameter settings. Boxplots show simulations. Empirical medians are plotted with blue diamonds to compare simulations to the empirical data. **a** Same criterion and chunk-size parameters. **b** Different criterion parameters but same chunk sizes. **c** Different chunk-size but same criterion parameters. Abbreviations: children (CH), younger adults (YA), older adults (OA).

To examine the resource-rational sampling model on the parameter level in more detail, we performed additional simulation studies (Fig. S20). In the first simulation, we assumed the same chunk size and criterion values across children, younger adults, and older adults (Fig. S20a). This model clearly fails to capture age-related differences in the three summary measures of perseveration probability, estimation error, and anchoring bias. In the second simulation, we assumed different criterion values but similar chunk-size parameters (Fig. S20b). One crucial insight of this simulation is that differences in the criterion parameter capture differences in perseveration probability; however, the model fails to explain differences in estimation errors and the anchoring bias. Third, we tested a model with different chunk sizes but similar criterion values. This model is equivalent to the simulations in the main paper and captures the results well (Fig. S20c). While this suggests that the age groups rely on similar criterion values, it is important to keep in mind that a model without criterion parameter fails to capture our results (see SM Resource-only model). Taken together, the key insight of this parameter comparison is that age-related differences in sampling behavior seem to be primarily driven by differences in the chunk-size parameter that determines the size of each set of samples.

Resource-only model

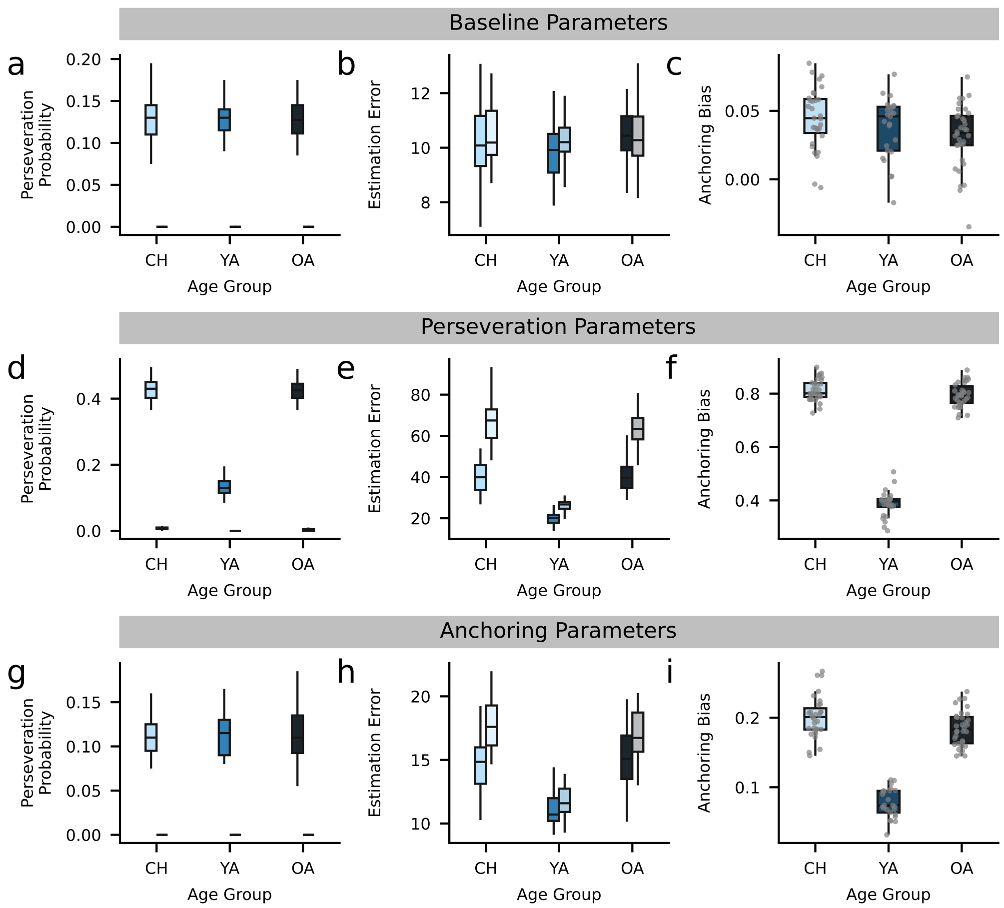

Figure S21. Resource-only model. The resource-only model assumed that age groups only differ in their cognitive resources. It allowed us to test if a model without a dynamic stopping criterion captures the behavioral results similar to the resource-rational sampling model. We first simulated data with baseline parameter settings leading to **a** moderate perseveration, **b** low estimation errors, and **c** a small anchoring bias. We then chose "perseveration" parameter values that **d** yielded age-related differences in perseverative behavior. However, these parameter values led to **e** strongly overestimated estimation errors and **f** anchoring biases. In opposition to these parameter settings, the "anchoring" parameters **g** did not capture perseveration but instead **h** age differences in estimation errors and **i** the anchoring bias. Together, these results suggest that a dynamic stopping criterion, like in the resource-rational sampling model, is necessary to explain our results. Abbreviations: children (CH), younger adults (YA), older adults (OA).

We performed a simulation study with the resource-only model to examine whether an alternative sampling model in which age groups only differ in their cognitive resources could explain our results. This model lacked a dynamic stopping criterion and instead invested a fixed amount of resources on each trial. As such, a sampling model without a dynamic stopping criterion does not naturally show perseverative behavior. Still, to test the model under fair conditions (in favor of the resource-only model), we added the logistic function explained above with fixed parameters that yielded perseverative behavior for smaller prediction errors (intercept $\beta_0 = 0.85$, slope $\beta_1 = -0.25$). We first simulated data with a baseline model, which generated a sufficient amount of samples to achieve high task performance (100 samples). In particular, the model showed moderate perseveration due to the logistic function (Fig. S21a), low estimation errors in both the standard and anchoring condition (Fig. S21b), as well as a minor anchoring bias (Fig. S21c).

We next tried to model age-related differences in perseveration and anchoring based on common parameter settings. When we assumed a very low number of samples in children and older adults (2 samples) and more samples in younger adults (10 samples), the simulations suggested perseveration differences between the groups (Fig. S21d). However, these parameter settings clearly overestimated estimation errors and the anchoring bias. Conversely, assuming more samples (children and older adults 20 samples, younger adults 50 samples) led to more realistic estimation errors and anchoring biases but underestimated age-related perseveration differences. Therefore, the resource-only control model did not offer a unifying explanation for the observed biases across the lifespan. As such, it suggests that a resource-rational stopping criterion, as proposed by our main sampling model, is necessary to capture age-related differences in perseveration and anchoring.

Motor model

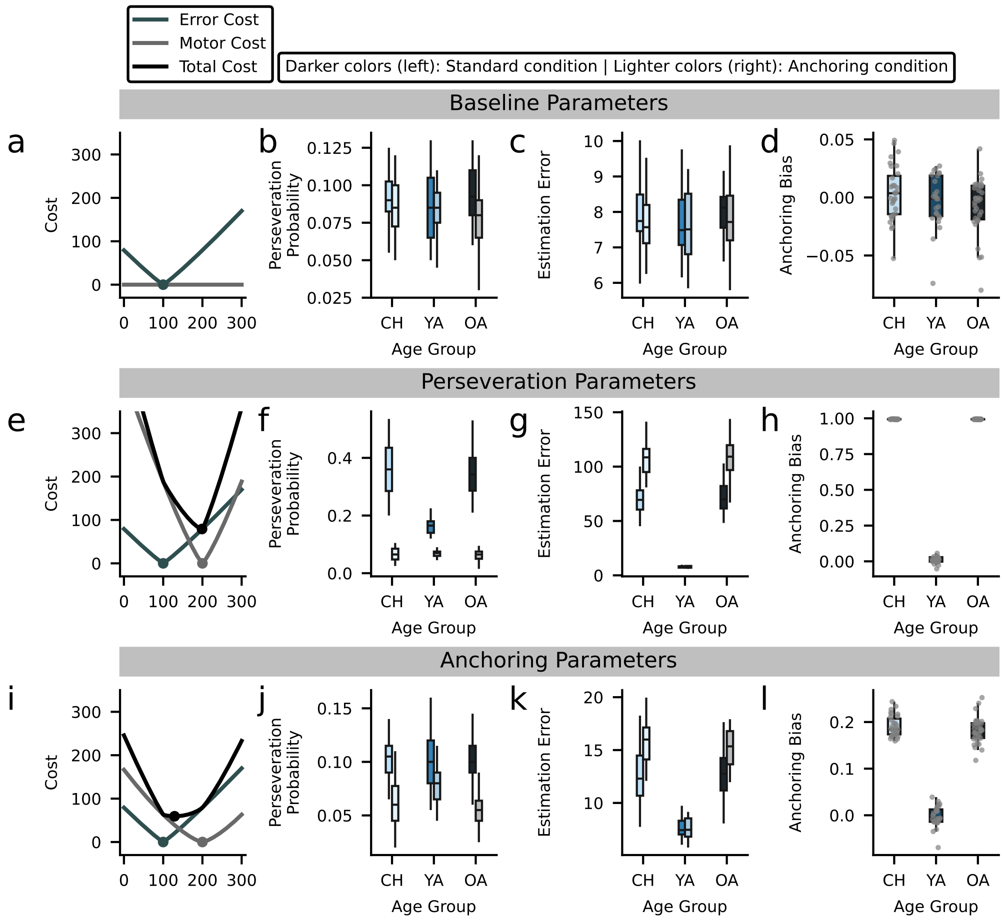

Figure S22. Motor model. The motor model allowed us to test whether adding motor costs to the reduced Bayesian model could explain perseveration and anchoring from a perspective different from our sampling model. **a** A baseline model only considered error costs. Error costs were higher when the bucket was placed further away from the optimal location. This version resembled the reduced Bayesian model. **b** The baseline model perseverated only rarely, **c** estimation errors were low, and **d** the model did not show systematic anchoring biases. **e** The "perseveration" version of the motor model had extremely high motor costs (and the same error costs as the baseline model). The three lines only reflect the parameters of simulated children and older adults. **f** Due to the high motor costs, particularly simulated children (CH) and older adults (OA) perseverated more frequently than younger adults (YA). **g** However, estimation errors were strongly overestimated, and **h** anchoring biases showed a ceiling effect. **i** In contrast, the "anchoring" version had lower motor costs (and the same error costs as before). The three lines only reflect the parameters of simulated children and older adults. **j** Due to the lower motor costs, perseveration was underestimated. **k,l** However, in this case, estimation errors and anchoring biases were closer to the empirical data.

To explore whether motor costs provide an alternative account of perseveration and anchoring, we developed a motor-control model. The motor model was based on the reduced Bayesian model and had additional motor costs associated with bucket adjustments. Larger updates of the bucket position were associated with higher motor costs. At the same time, the model considered "error" costs due to an under-adjustment of the bucket. Therefore, the model was confronted with a trade-off between motor and error costs. Not adjusting the bucket at all would minimize motor costs but would also be associated with high error costs. In contrast, updating the bucket optimally according to the reduced Bayesian model would minimize error costs but would be associated with higher motor cots. The model solved this trade-off by minimizing the combined costs (motor and error costs).

We first simulated data using a baseline model with parameter settings resembling the reduced Bayesian model (Fig. S22a-d). This version of the model only considered error costs. Error costs were minimal in the location of the optimal belief ($v_l = 0.5$, $e_v = 1.1$, $v_m = 0.0$, $e_m = 1.0$, Fig. S22a). Therefore, the model showed approximately optimal belief updating. We ensured that perseveration did only rarely occur (logistic function parameters modeling perseveration were $b_0 = 0.5$ and $b_1 = -5.0$) and found that estimation errors were low and anchoring biases absent.

The next simulations included additional motor costs and indicated that the motor model did not provide a unifying explanation for perseveration and anchoring. Assuming extremely high motor costs, the model showed perseveration probabilities that differed between younger adults and children and older adults (children and older adults: $v_m = 0.75$, $e_m = 1.2$; younger adults $v_m = 0.5$, $e_m = 1.0$; Fig. S22e). However, the resulting estimation errors and anchoring biases were unrealistically high. Assuming lower motor costs allowed us to successfully capture differences in anchoring and estimation errors (children and older adults: $v_m = 0.1$, $e_m = 1.4$; younger adults $v_m = 0.1$, $e_m = 1.0$; Fig. S22i). However, these parameter settings strongly underestimated perseveration in all age groups. Taken together, while the motor model can separately account for perseveration or anchoring, it cannot recover the two biases based on common parameter settings. Therefore, our control analyses do not support the idea that adding motor costs to the reduced Bayesian model provides an alternative explanation for perseveration and anchoring.

Third experiment

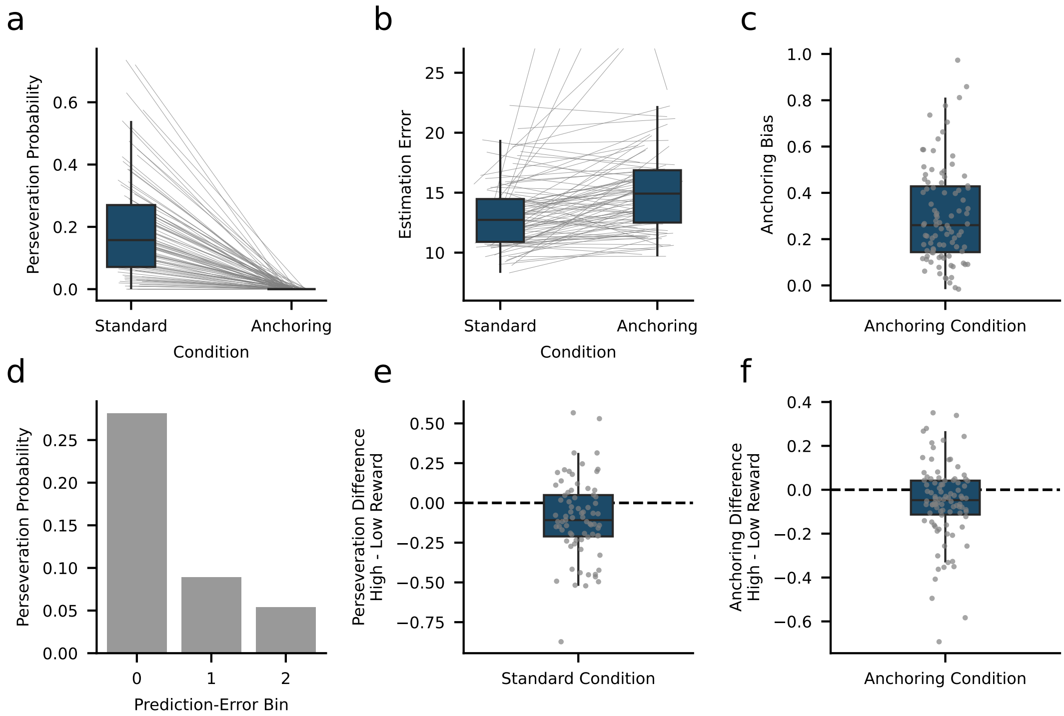

Figure S23. Third experiment. We conducted a third experiment with block-wise incentive manipulation. **a**. The anchoring manipulation strongly reduced perseveration. In the standard condition, perseveration levels were comparable to our second experiment, and perseveration was virtually at zero in the anchoring condition. **b** Similarly, estimation errors were higher in the anchoring compared to the standard condition. That is, the accuracy of the predictions was lower when the starting location of the bucket was randomly perturbed. **c** Consistent with this effect, participants showed an anchoring bias like in our previous experiment, which partly drives performance differences between the conditions. **d** We then examined the effects of the reward manipulation (high vs. low reward blocks) on perseveration. To do so, we divided trials according to absolute prediction errors. Each bin contained one-third of the trials. In line with the previous experiments, this analysis showed that perseveration was particularly frequent in bin 1 (smaller prediction errors) compared to the other bins (larger prediction errors). **e**. Subsequently, we examined the differences in perseveration probability between high- and low-reward blocks within bin 1 for subjects that perseverated on at least 5% of the low-reward trials in bin 1 (to avoid including subjects that did not perseverate). In line with our hypothesis, perseveration was reduced in high-reward compared to low-reward blocks (in the text, we also show that the result remains when all participants are considered). **f** Similarly, we found that the anchoring bias was reduced in high- compared to low-reward blocks, providing further support that people draw more samples when more reward is available.

While the results from our second experiment were largely in line with our overarching resource-rational theory, one result was, on its face, inconsistent with our theory, leading us to further investigate it through an additional experiment. In particular, our second experiment included an incentive manipulation aimed at reducing perseveration. However, this manipulation had no effect on behavior, neither in younger adults nor in children or older adults. We speculated that the manipulation was ineffective because participants had to switch between high and low rewards on a trial-by-trial basis. Tracking and responding to the reward magnitude on each trial might itself be cognitively demanding, and thus, participants may have taken a resource-rational strategy that we did not initially anticipate: ignoring the reward cue altogether.

In order to test this idea, we performed an online experiment in younger adults that matched the basic experimental design but used a block-wise reward manipulation. This decision was based on research on incentive effects on cognitive control, where block-wise reward manipulations are much more common than trial-by-trial manipulations (Shenhav et al., 2021). In such a design, participants could potentially learn to allocate resources rationally within each block type without dedicating resources toward interpreting cues, thereby allowing us to better test whether perseveration and anchoring are reduced in a high-reward compared to a low-reward condition.

The anchoring manipulation reduced perseveration. Perseveration probability was significantly higher in the standard compared to the anchoring condition (Fig. S23a; median = 0.158, IQR 0.071 to 0.27, $w = 0.0$, $p < 0.001$, $R^2 = 0.722$). Moreover, estimation errors were significantly lower in the standard compared to the anchoring condition (Fig. S23b; median = -1.569, IQR -3.402 to 0.299, $w = 817.0$, $p < 0.001$, $R^2 = 0.303$). These performance differences were partly driven by an anchoring bias similar to the second experiment (Fig. S23c; median = 0.26, IQR 0.144 to 0.428, $w = 4.0$, $p < 0.001$, $R^2 = 0.751$). We next examined the distribution of perseveration probability across different prediction-error magnitudes. To do so, we divided trials into three bins of prediction-error magnitude ordered from low to high. Like in the previous experiments, perseveration was particularly frequent for smaller prediction errors (bin 1; Fig. S23d).

To test if the reward manipulation affected perseveration probability, we examined the difference between the high- and low-reward blocks in participants who perseverated in at least 5% of the low-reward-block trials in bin 1. This analysis indicated that perseveration was significantly lower during high-reward blocks (Fig. S23e; median = -0.108, IQR -0.211 to 0.049, $w = 797.0$, $p = 0.001$, $R^2 = 0.147$). We also compared perseveration between the two reward conditions in bin 1 across all participants independent of our somewhat arbitrary perseveration criterion of the previous analysis. This comparison similarly showed significantly reduced perseveration in high-reward blocks (median = -0.054, IQR -0.192 to 0.054, $w = 1323.0$, $p = 0.008$, $R^2 = 0.074$). Finally, we compared the anchoring bias between the high- and low-reward conditions. Similar to the previous results, the anchoring bias was significantly lower in the high-reward condition (Fig. S23f; median = -0.047, IQR -0.113 to 0.042, $w = 1390.0$, $p = 0.001$, $R^2 = 0.107$). Together, the results of the third experiment show that reward prospects reduce perseveration and anchoring. These findings suggest that participants exert more cognitive effort to obtain higher rewards, supporting our model's prediction that reward motivates people to draw more samples for accurate belief updating.

Optimal and approximate Bayesian inference

Optimal inference

Optimal Bayesian inference over the mean μ_t of the outcome-generating Gaussian distribution conditional on the observed outcomes $x_{1:t}$ can be performed by considering the conditional distribution

$$p(\mu_t|x_{1:t}) = \frac{p(\mu_t, x_{1:t})}{p(x_{1:t})} \tag{31}$$

where the joint distribution of μ_t and $x_{1:t}$ is the result of the following derivation:

$$\begin{aligned}
p(\mu_t, x_{1:t}) &= \sum_{c_t} \int p(x_{1:t}, \mu_t, \mu_{t-1}, c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|x_{1:t-1}, \mu_t, \mu_{t-1}, c_t) p(x_{1:t-1}, \mu_t, \mu_{t-1}, c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|\mu_t) p(x_{1:t-1}, \mu_t, \mu_{t-1}, c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|\mu_t) p(\mu_t|x_{1:t-1}, \mu_{t-1}, c_t) p(x_{1:t-1}, \mu_{t-1}, c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t) p(x_{1:t-1}, \mu_{t-1}, c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t) p(\mu_{t-1}|x_{1:t-1}, c_t) p(x_{1:t-1}, c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t) p(\mu_{t-1}|x_{1:t-1}) p(x_{1:t-1}, c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t) p(\mu_{t-1}|x_{1:t-1}) p(x_{1:t-1}|c_t) p(c_t) d\mu_{t-1} \\
&= \sum_{c_t} \int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t) p(\mu_{t-1}|x_{1:t-1}) p(c_t) d\mu_{t-1}
\end{aligned} \tag{32}$$

and the marginal distribution over outcomes $x_{1:t}$ can be obtained by the integral

$$p(x_{1:t}) = \int p(\mu_t, x_{1:t}) d\mu_t \tag{33}$$

which together thus yields

$$p(\mu_t|x_{1:t}) = \frac{\sum_{c_t} \int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t) p(\mu_{t-1}|x_{1:t-1}) p(c_t) d\mu_{t-1}}{\int p(\mu_t, x_{1:t}) d\mu_t} \tag{34}$$

Expanding the sum in the numerator, i.e.,

$$p(\mu_t|x_{1:t}) = \\
\frac{\int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t = 0) p(\mu_{t-1}|x_{1:t-1}) p(c_t = 0) d\mu_{t-1} + \int p(x_t|\mu_t) p(\mu_t|\mu_{t-1}, c_t = 1) p(\mu_{t-1}|x_{1:t-1}) p(c_t = 1) d\mu_{t-1}}{\int p(\mu_t, x_{1:t}) d\mu_t} \tag{35}$$

shows that this inference scheme considers the case that a change point did not occur (first term) and that a change point occurred (second term). To keep track of this mixture distribution over many trials, message-passing algorithms that consider all possible mixtures between change-point and no-change-point trials have been applied in previous work (Adams & MacKay, 2007; Wilson et al., 2010). However, these methods require many costly computations, which arguably go beyond human-level capacities (Nassar et al., 2010). Therefore, Nassar and colleagues proposed a reduced Bayesian model that requires fewer computations and shares the core assumptions of the optimal Bayesian solution.

Approximate inference

The reduced Bayesian model approximates the mixture over all combinations of no-change-point and change-point trials with a Gaussian distribution with a similar mean and variance as the mixture distribution computed in optimal Bayesian inference. Notably, the reduced model also considers the scenario of "no change point" and "change point"; however, this only concerns the current trial t instead of all trials $1 : t$.

We first show the derivation of the mixture component that assumes the absence of a change point (Murphy, 2012). In the absence of change points, one can sequentially update a conjugate Gaussian distribution over the mean of the outcome-generating Gaussian distribution according to

$$\begin{aligned} \mu_{t+1} &= \frac{\sigma^2}{\sigma^2 + \sigma_t^2}\mu_t + \frac{\sigma_t^2}{\sigma^2 + \sigma_t^2}x_t \\ &= (1 - \tau_t)\mu_t + \tau_t x_t \\ &= \mu_t + \tau_t x_t - \tau_t \mu_t \\ &= \mu_t + \tau_t(x_t - \mu_t) \end{aligned} \tag{36}$$

where

$$\tau_t := \frac{\sigma_t^2}{\sigma^2 + \sigma_t^2} \tag{37}$$

refers to the relative uncertainty over the estimated Gaussian mean parameter. σ_t^2 called estimation uncertainty can be updated according to

$$\sigma_t^2 = \tau_t \sigma^2 \tag{38}$$

Next, we consider the case that a change point occurred. Recall that a change point renders the past irrelevant, and therefore, under the assumption of a flat prior expectation over the parameter, the best prediction is the most recent outcome x_t.

To appropriately weight these two mixture components, the reduced Bayesian model computes change-point probability ω_t, defined as

$$\omega_t := p(c_t = 1|x_t) = \frac{p(c_t = 1, x_t)}{p(x_t)} \tag{39}$$

where

$$p(x_t) = \sum_{c_t} p(x_t|c_t)p^h(c_t) \tag{40}$$

and thus

$$
\begin{aligned}
p(c_t = 1|x_t) &= \frac{p(x_t|c_t = 1)p^h(c_t = 1)}{\sum_{c_t} p(x_t|c_t)p^h(c_t)} \\
&= \frac{p(x_t|c_t = 1)p^h(c_t = 1)}{p(x_t|c_t = 0)p^h(c_t = 0) + p(x_t|c_t = 1)p^h(c_t = 1)} \\
&= \frac{(1/300)h}{N(x_t; \mu_t, \sigma_t^2 + \sigma^2)(1 - h) + (1/300)h} \\
&= \frac{(1/300)h}{N(\delta_t; 0, \sigma_t^2 + \sigma^2)(1 - h) + (1/300)h}
\end{aligned}
\tag{41}
$$

The reduced Bayesian model then computes the mean of the outcome-generating distribution as a function of the two mixture components weighted by change-point probability

$$
\mu_{t+1} = (\mu_{t+1}|c_t = 0)(1 - \omega_t) + (\mu_{t+1}|c_t = 1)\omega_t \tag{42}
$$

which can be reformulated to an error-driven learning rule as follows

$$
\begin{aligned}
\mu_{t+1} &= (\mu_{t+1}|c_t = 0)(1 - \omega_t) + (\mu_{t+1}|c_t = 1)\omega_t \\
&= ((1 - \tau_t)\mu_t + \tau_t x_t)(1 - \omega_t) + x_t\omega_t \\
&= (\mu_t + \tau_t x_t - \tau_t\mu_t)(1 - \omega_t) + x_t\omega_t \\
&= (\mu_t + \tau_t(x_t - \mu_t))(1 - \omega_t) + x_t\omega_t \\
&= (\mu_t + \tau_t(x_t - \mu_t))(1 - \omega_t) + (\mu_t + (x_t - \mu_t))\omega_t \\
&= (\mu_t + \tau_t\delta_t)(1 - \omega_t) + (\mu_t + \delta_t)\omega_t \\
&= \mu_t - \mu_t\omega_t + \tau_t\delta_t - \tau_t\delta_t\omega_t + \mu_t\omega_t + \delta_t\omega_t \\
&= \mu_t - \mu_t\omega_t + \mu_t\omega_t + \tau_t\delta_t - \tau_t\delta_t\omega_t + \delta_t\omega_t \\
&= \mu_t + \tau_t\delta_t - \tau_t\delta_t\omega_t + \delta_t\omega_t \\
&= \mu_t + (\tau_t - \tau_t\omega_t + \omega_t)\delta_t \\
&= \mu_t + (\omega_t + \tau_t - \tau_t\omega_t)\delta_t \\
&= \mu_t + \alpha_t\delta_t
\end{aligned}
\tag{43}
$$

where

$$
\alpha_t := \omega_t + \tau_t - \tau_t\omega_t \tag{44}
$$

Finally, estimation uncertainty, that is, the variance of the approximated mixture distribution σ_{t+1}^2, is the ω_t-weighted mixture of the no-change-point and change-point component

$$
\sigma_{t+1}^2 := \omega_t\sigma^2 + (1 - \omega_t)\tau_t\sigma^2 + \omega_t(1 - \omega_t)(\delta_t(1 - \tau_t))^2 \tag{45}
$$

which follows from substituting the conditional means and variances into the equation of the variance of a mixture of two Gaussian distributions

$$
\begin{aligned}
\sigma_{t+1}^2 &= (\sigma_{t+1}^2|c_t=0)(1-\omega_t) + (\sigma_{t+1}^2|c_t=1)\omega_t + (1-\omega_t)\omega_t((\mu_{t+1}|c_t=0) - (\mu_{t+1}|c_t=1))^2 \\
&= \tau_t\sigma^2(1-\omega_t) + \sigma^2\omega_t + \omega_t(1-\omega_t)((\mu_t+\tau_t\delta_t) - (\mu_t+\delta_t))^2 \\
&= \tau_t\sigma^2(1-\omega_t) + \sigma^2\omega_t + \omega_t(1-\omega_t)(\tau_t\delta_t - \delta_t)^2 \\
&= \tau_t\sigma^2(1-\omega_t) + \sigma^2\omega_t + \omega_t(1-\omega_t)(\delta_t(\tau_t-1))^2 \\
&= \omega_t\sigma^2 + (1-\omega_t)\tau_t\sigma^2 + \omega_t(1-\omega_t)(\delta_t(1-\tau_t))^2
\end{aligned}
\tag{46}
$$

Catch-trial derivations

On catch trials, we assumed that the provided helicopter cue that was centered on the true mean was perceived as an ambiguous cue that corresponded to a Gaussian probability density function (PDF) where the mean corresponded to the true helicopter location and the variance that expressed the reliability of the cue was assumed to be different between participants. In effect, the updated mean and variance followed from the multiplication of the Gaussian PDF over the learned helicopter location based on the history of outcomes and the Gaussian PDF over the presented helicopter position, where the covariance between the two distributions corresponds to

$$C := \frac{1}{\frac{1}{\sigma_t^2} + \frac{1}{\sigma_H^2}} \tag{47}$$

and the mean corresponds to

$$
\begin{aligned}
\mu_{t+1} &= C\frac{1}{\sigma_t^2}\mu_{t+1} + C\frac{1}{\sigma_H^2}\mu_H \\
&= \frac{1}{\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}}\frac{1}{\sigma_t^2}\mu_{t+1} + \frac{1}{\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}}\frac{1}{\sigma_H^2}\mu_H \\
&= \frac{1}{\sigma_t^2\Big(\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}\Big)}\mu_{t+1} + \frac{1}{\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}}\frac{1}{\sigma_H^2}\mu_H \\
&= \frac{1}{\sigma_t^2\Big(\frac{\sigma_H^2}{\sigma_t^2\sigma_H^2}+\frac{\sigma_t^2}{\sigma_t^2\sigma_H^2}\Big)}\mu_{t+1} + \frac{1}{\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}}\frac{1}{\sigma_H^2}\mu_H \\
&= \frac{1}{\sigma_t^2\frac{\sigma_H^2+\sigma_t^2}{\sigma_t^2\sigma_H^2}}\mu_{t+1} + \frac{1}{\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}}\frac{1}{\sigma_H^2}\mu_H \\
&= \frac{\sigma_t^2\sigma_H^2}{\sigma_t^2(\sigma_H^2+\sigma_t^2)}\mu_{t+1} + \frac{1}{\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}}\frac{1}{\sigma_H^2}\mu_H \\
&= \frac{\sigma_H^2}{\sigma_H^2+\sigma_t^2}\mu_{t+1} + \frac{1}{\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}}\frac{1}{\sigma_H^2}\mu_H \\
&= \frac{\sigma_H^2}{\sigma_H^2+\sigma_t^2}\mu_{t+1} + \frac{1}{\sigma_H^2\Big(\frac{1}{\sigma_t^2}+\frac{1}{\sigma_H^2}\Big)}\mu_H \\
&= \frac{\sigma_H^2}{\sigma_H^2+\sigma_t^2}\mu_{t+1} + \frac{1}{\sigma_H^2\Big(\frac{\sigma_H^2}{\sigma_t^2\sigma_H^2}+\frac{\sigma_t^2}{\sigma_t^2\sigma_H^2}\Big)}\mu_H \\
&= \frac{\sigma_H^2}{\sigma_H^2+\sigma_t^2}\mu_{t+1} + \frac{1}{\sigma_H^2\frac{\sigma_H^2+\sigma_t^2}{\sigma_t^2\sigma_H^2}}\mu_H \\
&= \frac{\sigma_H^2}{\sigma_H^2+\sigma_t^2}\mu_{t+1} + \frac{\sigma_t^2\sigma_H^2}{\sigma_H^2(\sigma_t^2+\sigma_H^2)}\mu_H \\
&= \frac{\sigma_H^2}{\sigma_H^2+\sigma_t^2}\mu_{t+1} + \frac{\sigma_t^2}{\sigma_t^2+\sigma_H^2}\mu_H \\
&= \Big(1-\frac{\sigma_t^2}{\sigma_t^2+\sigma_H^2}\Big)\mu_{t+1} + \frac{\sigma_t^2}{\sigma_t^2+\sigma_H^2}\mu_H \\
&= (1-w)\mu_{t+1} + w\mu_H
\end{aligned}
\tag{48}
$$

where

$$
w := \frac{\sigma_t^2}{\sigma_t^2+\sigma_H^2} \tag{49}
$$

Task instructions

Before completing the experimental task, participants were instructed verbally and using written material on the computer screen. We built up the helicopter cover story step-by-step, with practice blocks in between. Every practice phase consisted of a low- and high-variability block. In the first practice phase, the helicopter was always visible. To ensure that participants understood that the helicopter was the best predictor of the outcomes, we used a response criterion that required participants to put their bucket ten times exactly underneath the visible helicopter. These blocks either stopped after the criterion was reached or after a maximum of 80 trials. In the second practice phase, clouds covered the helicopter, and participants had to infer its location based on the outcomes. Occasionally, the helicopter was visible during catch trials, and participants could use this information for learning. These blocks consisted of 50 trials and were the same as in the real experiment. After the practice session, we emphasized again that participants would earn money for their collected points.

German instructions

- Slide 1: Ein Helikopter wirft Goldsäcke ab. Versuche, möglichst viele aufzufangen!
- Slide 2: Mit der Tastatur kannst du den Eimer bewegen. Mit den Tasten F und G bewegst du den Eimer nach links und mit H und J bewegst du den Eimer nach rechts.
- Slide 3: Die Goldsäcke werden immer in der Nähe des Helikopters landen, ihre exakte Position ist jedoch von der Windstärke abhängig. Die beste Strategie ist, den Eimer immer genau unter den Helikopter zu stellen! Der Helikopter bleibt meist längere Zeit am selben Ort, manchmal fliegt er aber auch ein Stück weiter und verändert damit seine Position.
- Slide 4: Jeder Sack enthält entweder Gold oder Steine. Für das von dir aufgefangene Gold, erhälst du Geld!! Steine hingegen sind nichts wert und du erhälst dafür auch kein Geld. Ob ein Sack Gold oder Steine enthält, erfährst du jeweils erst, wenn er auf dem Boden landet.
- Slide 5: Erster Übungsdurchgang: Zur Erinnerung: Um möglichst viele Säcke zu fangen, solltest du den Eimer immer unter den Helikopter stellen. Um den kommenden Übungsdurchgang erfolgreich abzuschließen, musst du den Eimer in mindestens zehn Fällen genau unter den Helikopter gestellt haben.
- Slide 6: Probiere es jetzt selber aus!
- [Participants completed the first practice session]
- Slide 7: Nun wird der Helikopter von Wolken verdeckt. Ansonsten bleibt alles gleich.
- Du solltest weiterhin versuchen, den Eimer direkt unter den Helikopter zu stellen, auch wenn du ihn nun nicht mehr sehen kannst.
- Slide 8: Die rote Leiste am unteren Bildschirmrand zeigt dir jeweils die Distanz zwischen der letzten Position des Eimers und der Stelle, wo der letzte Sack gelandet ist. Es ist sinnvoll, den Eimer an einer Stelle entlang der roten Leiste zu positionieren.
- Slide 9: Manchmal verschwinden die Wolken und der Helikopter ist für einen Moment sichtbar. Dies hilft dir, den Eimer richtig hin zu stellen.

- Slide 10: Jetzt folgen zwei Übungsblöcke und dann geht‘s los!
- [Participants completed the second training session and then the experiment started]
- In the second experiment, we introduced the reward manipulation in the following way: Jeder Sack enthält entweder Gold oder Steine. Gold ist 4 mal so viel wert wie Steine. Ob ein Sack Gold oder Steine enthält, kannst du an der Farbe des „+“ erkennen. Insgesamt erhältst du für jeden gefüllten Eimer 50 Cent, die du am Ende des Versuchs wirklich ausgezahlt bekommst! Versuche also möglichst viele Säcke zu fangen!

English translation

- Slide 1: A helicopter drops bags of gold. Try to catch as many as possible!
- Slide 2: Use the keyboard to move the bucket. Use the F and G keys to move the bucket to the left and with H and J you move the bucket to the right.
- Slide 3: The gold bags will fall near the helicopter, but their exact position depends on the strength of the wind. The best strategy is to always put the bucket exactly underneath the helicopter! The helicopter usually stays in the same place, but occasionally it flies to a new location.
- Slide 4: Each bag contains either gold or stones. For the gold that you collect, you get money!! Stones, in contrast, are worth nothing, and you get no money for them. You will only find out if a bag contains gold or stones when it lands on the ground.
- Slide 5: First practice run: Reminder: To catch as many bags as possible, you should always place the bucket underneath the helicopter. To successfully complete the next practice run, you must place the bucket exactly under the helicopter at least ten times.
- Slide 6: Give it a try!
- [Participants completed the first practice session]
- Slide 7: Now the helicopter will be hidden by clouds. Everything else is exactly the same.
- You should still place your bucket directly under the helicopter even though you can no longer see it.
- Slide 8: The red bar at the bottom of the screen shows the distance between the last position of the bucket and the last bag location.
- Your best strategy is to position the bucket within the red bar.
- Slide 9: Sometimes the clouds disappear, and the helicopter is visible for a moment. This helps you to position the bucket correctly.
- Slide 10: Now you will complete two more practice sessions, and then the task starts!
- [Participants completed the second training session and then the experiment started]
- In the second experiment, we introduced the reward manipulation in the following way: Each bag contains either gold or stones. Gold is worth 4 times as much as stones. You can tell whether a bag contains gold or stones by the color of the "+". You receive a total of 50 cents for each filled bucket, which you really get paid out at the end of the experiment! So try to catch as many bags as possible!

References

Adams, R. P., & MacKay, D. J. C. (2007, October). Bayesian online changepoint detection. http://arxiv.org/abs/0710.3742

Li, S.-C., Lindenberger, U., Hommel, B., Aschersleben, G., Prinz, W., & Baltes, P. B. (2004). Transformations in the couplings among intellectual abilities and constituent cognitive processes across the life span. *Psychological Science*, *15*(3), 155–163. https://doi.org/10.1111/j.0956-7976.2004.01503003.x

Lindenberger, U., Mayr, U., & Kliegl, R. (1993). Speed and intelligence in old age. *Psychology and Aging*, *8*(2), 207–220. https://doi.org/10.1037/0882-7974.8.2.207

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press.

Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, *7*(1), 11609. https://doi.org/10.1038/ncomms11609

Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, *30*(37), 12366–12378. https://doi.org/10.1523/JNEUROSCI.0822-10.2010

Palminteri, S., Wyart, V., & Koechlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences*, *21*(6), 425–433. https://doi.org/10.1016/j.tics.2017.03.011

Rigoux, L., Stephan, K. E., Friston, K., & Daunizeau, J. (2014). Bayesian model selection for group studies — Revisited. *NeuroImage*, *84*, 971–985. https://doi.org/10.1016/j.neuroimage.2013.08.065

Shenhav, A., Prater Fahey, M., & Grahek, I. (2021). Decomposing the motivation to exert mental effort. *Current Directions in Psychological Science*, *30*(4), 307–314. https://doi.org/10.1177/09637214211009510

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*(4), 1004–1017. https://doi.org/10.1016/j.neuroimage.2009.03.025

Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Computation*, *22*(9), 2452–2476. https://doi.org/10.1162/NECO_a_00007