

ARTICLE

Received 4 May 2015 | Accepted 12 Apr 2016 | Published 10 Jun 2016

DOI: 10.1038/ncomms11609

OPEN

Age differences in learning emerge from an insufficient representation of uncertainty in older adults

Matthew R. Nassar¹, Rasmus Bruckner^{2,3}, Joshua I. Gold⁴, Shu-Chen Li⁵, Hauke R. Heekeren³ & Ben Eppinger⁵

Healthy aging can lead to impairments in learning that affect many laboratory and real-life tasks. These tasks often involve the acquisition of dynamic contingencies, which requires adjusting the rate of learning to environmental statistics. For example, learning rate should increase when expectations are uncertain (uncertainty), outcomes are surprising (surprise) or contingencies are more likely to change (hazard rate). In this study, we combine computational modelling with an age-comparative behavioural study to test whether age-related learning deficits emerge from a failure to optimize learning according to the three factors mentioned above. Our results suggest that learning deficits observed in healthy older adults are driven by a diminished capacity to represent and use uncertainty to guide learning. These findings provide insight into age-related cognitive changes and demonstrate how learning deficits can emerge from a failure to accurately assess how much should be learned.

¹Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, Rhode Island 02912, USA. ²International Max Planck Research School LIFE, Max Planck Institute for Human Development, 14195 Berlin, Germany. ³Department of Education and Psychology, Freie Universität Berlin, 14195 Berlin, Germany. ⁴Department of Neuroscience, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA. ⁵Department of Psychology, TU Dresden, 01069 Dresden, Germany. Correspondence and requests for materials should be addressed to M.R.N. (email: matthew_nassar@brown.edu) or to B.E. (email: Benjamin.Eppinger@tu-dresden.de).

The ability to learn new outcome contingencies declines over the course of healthy aging in humans and animals^{1,2}. One proposed mechanism is age-related deficits in the ability to compute prediction errors³. However, evidence for this hypothesis is mixed^{1,2,4,5}. For example, learning ability and biomarkers of prediction error signalling are diminished for older versus younger adults in tasks that require learning probabilistic reward contingencies from experience^{3,5,6}. However, no such differences are found for gambling tasks with pre-specified reward contingencies^{7–12}. One possible explanation for these findings is that age-related deficits in error-driven learning do not reflect changes in how prediction errors are computed but rather how they are regulated according to environmental statistics^{13–16}. Such regulation should, for instance, enhance learning when contingencies are unknown (in periods of high uncertainty) but, if anything, suppress learning when contingencies are stable and known (low uncertainty)¹⁷.

Our goal was to test whether and how age differences in factors that regulate error-driven learning can account for age-related deficits in adaptive behaviour^{18–21}. We focused on three factors—uncertainty, surprise and hazard rate—that can have distinct effects on learning and are subserved, at least in part, by dissociable neural mechanisms^{14,15,22}.

Uncertainty about the state of a latent variable, such as the underlying market value of a company whose stock price fluctuates wildly, is termed estimation uncertainty^{19,23}. When estimation uncertainty is high, corresponding beliefs are unreliable, and should be revised according to new outcomes through learning^{13,18,23}. More precise prescriptions for the rate of learning are provided by relative uncertainty, which describes the contribution of estimation uncertainty to overall ignorance about upcoming events (for example, tomorrow's stock price). Relative uncertainty sets the optimal learning rate in a stable environment; thus a stock trader should be more sensitive to new stock prices when she is less certain about the underlying value of the company.

While uncertainty can provide a reasonable prescription for learning during periods of relative stability, efficient learning in dynamic environments also requires online detection of abrupt shifts in a latent state, such as might occur for a company with the announcement of a costly settlement²⁴. Such abrupt shifts are referred to as change points. They render the past irrelevant to the future, and thus require an immediate increase in learning to discard irrelevant information^{13,17}. While change points cannot always be perfectly identified, the probability of such an event can be efficiently estimated according to Bayes rule^{18,21,24}. In the Bayesian formulation, change-point probability depends critically on the likelihood of the newest observation (today's stock price) given the previous data (price history). The less likely the new observation, the more it indicates a fundamental change in the latent state. Thus a primary determinant of change-point probability, and consequently learning, is surprise^{17–19,21}. Surprise is greatest when outcomes deviate most substantially from predictions and thus can be measured by the absolute magnitude of prediction errors²⁵.

The probability that a surprising observation reflects a change point depends on the base rate of change points in the environment, termed the hazard rate²¹. The hazard rate acts as a prior that regularizes the number of detected change points; thus the lower the hazard rate, the greater the amount of surprise that is tolerated without an increase in learning. That is to say, a moderate decrement in the stock price of a company should be largely ignored if that company has been historically stable, but demands rapid recalibration of beliefs if that company has been prone to fundamental price shifts in the past.

Despite evidence that uncertainty, surprise and hazard rate are important mediators of learning in young adults, little is known about how subjective representations of these quantities affect learning across the adult lifespan. Older adults show deficits in learning tasks involving uncertainty and changes in task contingencies^{6,26–32}. In principle, such learning deficits could reflect a specific computational deficit in any of the normative learning factors listed above, but this idea has not been tested explicitly.

Here we provide such a test, using a three-step analytic approach. First, we used a normative model to show that deficits in processing uncertainty, surprise or hazard rate can lead to unique and diagnosable learning deficits. Second, we collected and analysed behavioural data to identify signatures of these different impairments and found that older adults have a deficit in representing and using uncertainty in the service of learning. Third, we confirmed this finding through quantitative model fitting and identified an additional age-related difference: older adults tended to be more variable in their assignments of learning rate than younger subjects. Altogether, these findings support the idea that age-related differences in learning reflect a selective impairment in the ability to represent uncertainty rather than a general deficit in the ability to learn from prediction errors.

Results

Task design and model predictions. Younger (20–30 years; $n=57$) and older (56–80 years; $n=57$) adults performed a modified version of a predictive inference task administered as part of a larger battery of behavioural tasks¹⁵. The task required participants to infer the position of a virtual helicopter based on the positions of bags that had previously fallen from it using a form of error-driven learning (Fig. 1a). On each trial, the participant placed a bucket under the inferred position of the helicopter. A bag then fell from the top of the screen and exploded to reveal contents that could be collected in the bucket. The participant could adjust the bucket placement for the subsequent trial in response to the indicated error between the previous bucket placement and the subsequent bag location. The key manipulation was the trial-by-trial bag position, which was normally distributed around a mean determined by the helicopter location. The mean location was stable for sequences of trials separated by abrupt, un-signalled change points at which the helicopter relocated to a random position. The s.d. of the bag distribution was consistent across change points within a block of trials but manipulated across blocks to give rise to varying levels of noise.

Effective performance in this task, as has been demonstrated by young adult participants in previous studies, is well described by a normative model that uses learning rates that are dynamically adjusted on each trial^{15,16,18,21,24} (Fig. 1b,c). Learning rate reflects the extent to which a given error is used to update the bucket position. A small learning rate (~ 0) implies that an error should be ignored, whereas a large learning rate (~ 1) implies that the error should be used in its entirety to place the bucket at the most recent bag location.

Dynamic learning rates prescribed by the normative model depend on two factors: change-point probability and relative uncertainty. The first factor, change-point probability, reflects the probability that the helicopter has relocated immediately before the most recent bag drop and can itself be dissociated into two separate components: hazard rate and surprise. The hazard rate reflects the frequency of change points, which is constant throughout the task but likely perceived differently across subjects (Fig. 1d, green). Surprise reflects the lack of correspondence

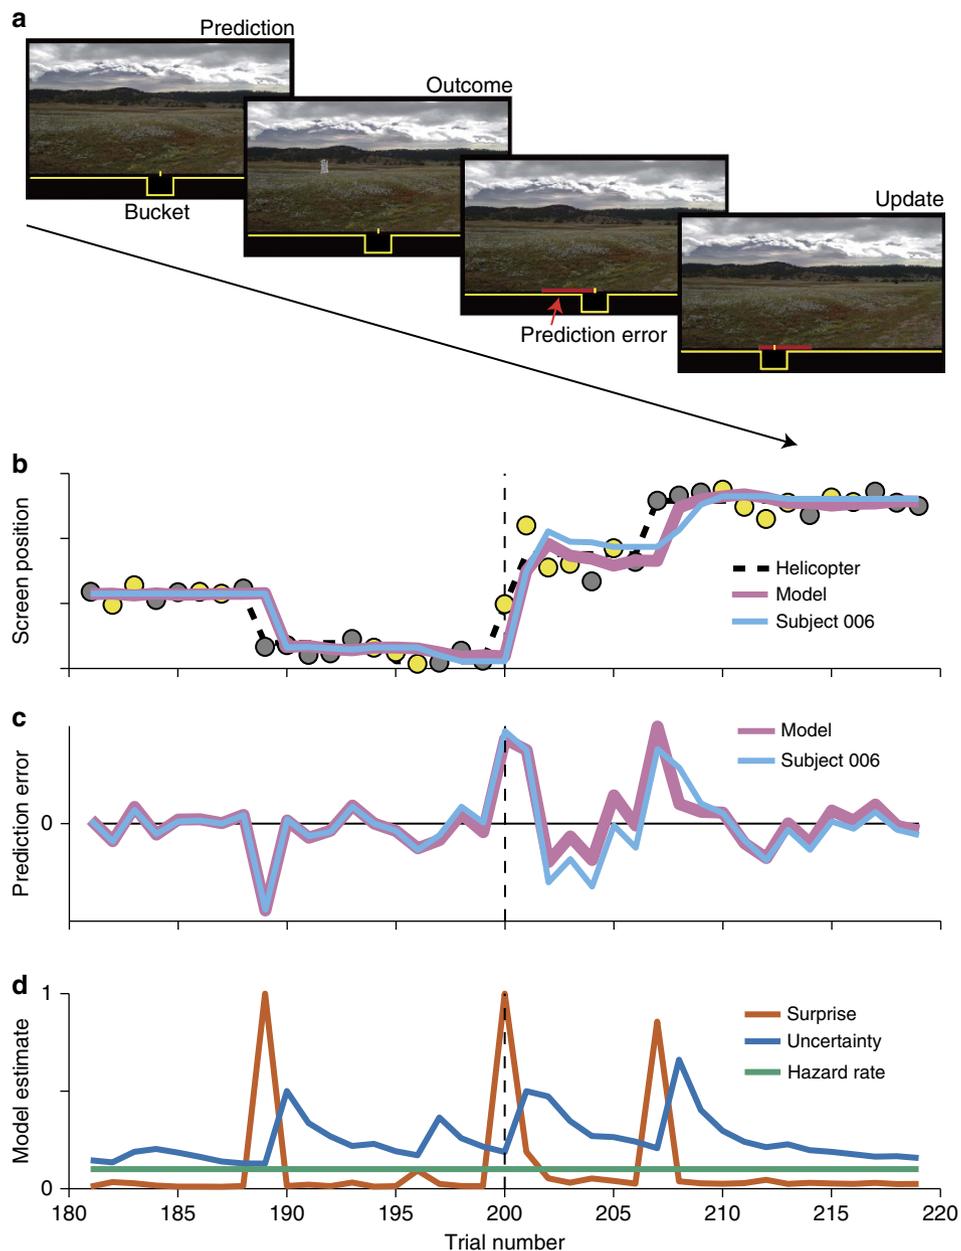


Figure 1 | Predictive inference task. (a) Participants were instructed to place a bucket under the most likely location of a helicopter to collect gold coins (prediction). On each trial the helicopter dropped a bag and some of the contents from this bag fell in the bucket placed by the participant (outcome). After bag contents were revealed and distributed (not shown) tick marks were displayed to mark the bucket and bag locations from the previous trial along with a red line indicating the difference between these locations (prediction error). Participants then updated the position of the bucket to maximize earnings on the subsequent trial. (b) Screen positions (ordinate) of various task events are plotted across trials (abscissa) for part of a task session. On most trials the helicopter (black dotted line) remained in the same location as the previous trial, though occasionally it would move to a completely new location of the screen (change point). On each trial the helicopter would drop a bag containing either gold or rocks (yellow and grey points, respectively) that was offset from the helicopter position according to a noise function that was manipulated blockwise (vertical dotted line indicates block transition). The participant placed the bucket on each trial according to previous bag locations (blue line). Participant bucket placements were qualitatively consistent with a normative learning model (pink line). (c) Spatial prediction errors (bag location–bucket location) for the same example participant (blue) and model (pink) were relatively small during periods of stability but extreme at change points. (d) Normative learning requires information about hazard rate, surprise and uncertainty on each trial. Hazard rate (green) was constant for the entire task leading the model to rely on the surprise associated with each prediction error to estimate trial-to-trial variability in change-point probability (orange). Likely change points lead to elevated estimates of relative uncertainty (blue) that decay with each well-predicted observation.

between predictions and outcomes and varies from trial-to-trial according to the absolute magnitude of prediction errors (Fig. 1d; orange). Thus, trial-to-trial fluctuations in change-point probability can be thought of as a normative prescription for surprise-driven learning.

The second factor, relative uncertainty, reflects the imprecision with which helicopter position can be estimated based on previous bag locations. Relative uncertainty is greatest after the model has observed only a single bag drop from a new helicopter location and decreases with each additional (unsurprising) bag

leading to learning rates that decay during periods of stability (Fig. 1d, blue).

One important feature of our model is that it can provide quantitative predictions of the distinct effects that specific computational impairments have on learning (Fig. 2). Insensitivity to surprise is characterized by a decrease in the steepness of the function relating relative prediction error to learning rate. Thus, surprise insensitivity would lead to a specific reduction in learning from large errors (Fig. 2a). In contrast, underestimation of the hazard rate increases the threshold for surprise necessary to infer a change point, without affecting the slope of the function. This leads to reduced learning rates across a wide range of moderately surprising outcomes that are less likely to be interpreted as change points if a low hazard rate is assumed. Underestimation of uncertainty leads to reduced learning specifically from small, unsurprising prediction errors that are unlikely to reflect change points in the helicopter location (Fig. 2c). All of these effects can also be distinguished from overall changes in learning rate, such as those expected from decreased prediction error magnitude (Fig. 2a–c, green).

Age-related differences in learning. To test whether learning differences across healthy aging reflect one or more of these specific computational deficits, we first analysed participant behaviour using a regression model. The model described learning behaviour according to principles of error-driven learning. Specifically, the distance that participants moved the bucket on each trial (the update) was explained using the difference between the position of the bag and the previous bucket position, called the prediction error (Fig. 1a). The slope of such a linear function is equivalent to a fixed learning rate in an error-driven learning model. The regression model included

separate interaction terms that could account for learning rates that were adjusted on each trial according to surprise, uncertainty and hazard rate (Fig. 3a). To diagnose the computational impairments highlighted above, we applied the regression model in sliding windows of trial outcomes ordered from the least to the most surprising as indexed by relative error magnitude (see Methods section).

Learning rates described by the regression model tended to increase with relative error magnitude, as predicted by the normative model, but with age-related differences for relatively small errors (Fig. 3b–d). On average, younger adults were more influenced than older adults by small errors (permutation test for H_0 : equal mean learning coefficients, $n = 57$ per group, $P < 0.05$). This reduced sensitivity to small errors in older adults did not reflect age-related differences in visual acuity. A deficit in visual acuity would have led to impairments in the processing of small errors independently of the noise level. However, we found that the age-related differences in behaviour were highly sensitive to expected variability, with pronounced differences between young and older participants for moderate spatial errors when the variability in the bag distribution was large (high noise; Fig. 3d).

These differences in overall learning rate were accompanied by age-related differences in how specific computational factors governed learning. The normative model suggests that learning should be greater during periods of uncertainty or after observing a surprisingly large prediction error (Fig. 1d). Our regression model captured such behaviour with interaction terms reflecting the effects of uncertainty and surprise on prediction error-driven learning. In Fig. 4 we plot the coefficients for these interaction terms for each sliding window of data sorted from least to most surprising. Positive coefficient values for these interaction terms indicate higher learning rates on trials where trial-by-trial estimates of relative uncertainty (uncertainty) or change-point

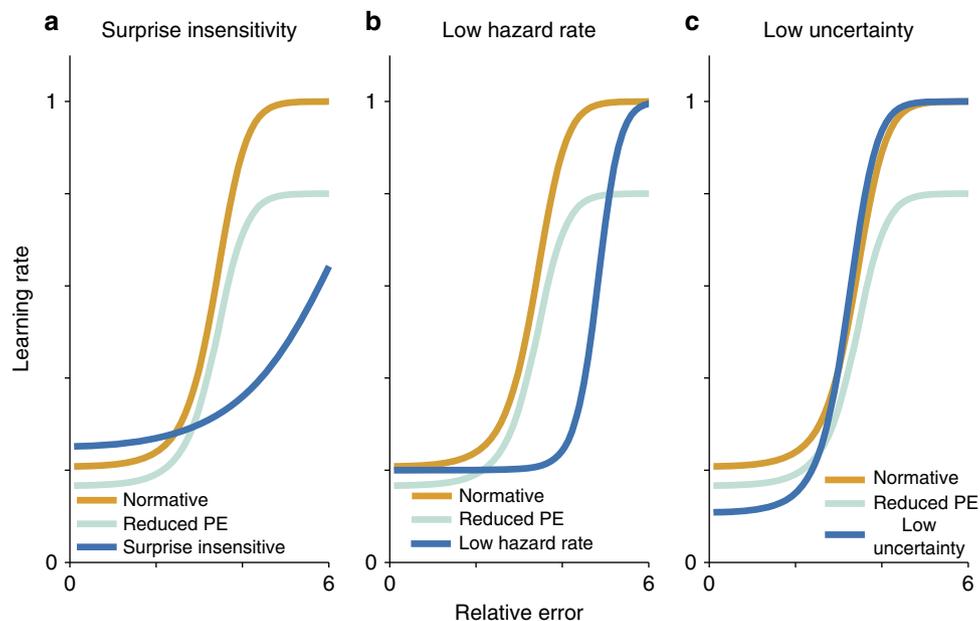


Figure 2 | Learning deficits stemming from hypothetical deviations from normative behaviour. Normative learning (orange line, all panels) requires identifying and responding appropriately to surprising outcomes while maintaining internal representations of uncertainty. While deficits in learning could emerge from a failure to fully represent prediction errors (green line, all panels), they could also emerge from a more selective deficit in appropriately choosing a learning rate such as: (a) Insensitivity to surprising outcomes, (b) Underestimation of the true base rate of change points (hazard rate) or (c) Insufficient representation of uncertainty. Such deviations from normative behaviour can all reduce the rate of learning, but do so only under a specific set of conditions. Insensitivity to surprise should lead to learning deficits most notable after surprisingly large errors (a). Underestimation of the hazard rate should result in learning deficits for moderately surprising errors (b). Reduced uncertainty should lead to selective learning impairments for the least surprising outcomes (c).

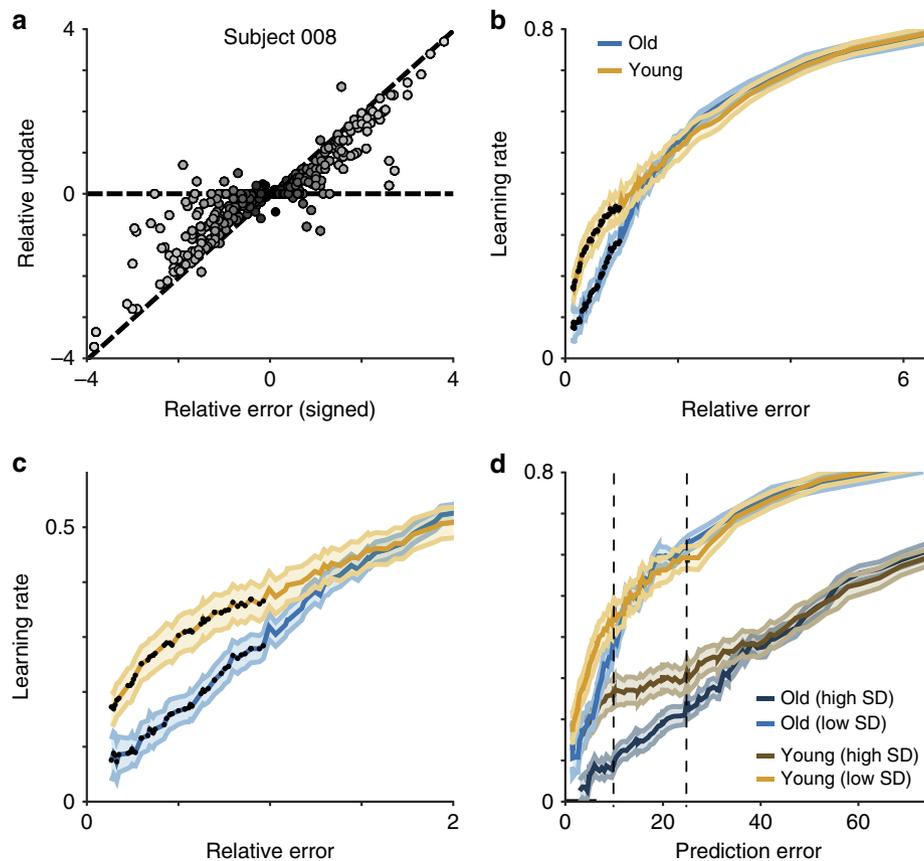


Figure 3 | Younger participants adjusted expectations more after making relatively small errors. (a) Relative bucket updates (bucket update divided by noise; ordinate) are plotted against signed relative error (error magnitude divided by noise; abscissa) for each trial completed by a single example subject. Trials are divided into bins according to unsigned relative error (absolute error magnitude divided by noise; lighter colours for larger relative errors). Consistent with normative learning, lighter points tend to be closer to the unity line (dotted diagonal) and darker points tend to deviate towards the zero update line (dotted horizontal). (b and c) Regression coefficients describing the effect of prediction errors on updates are plotted separately for older (blue) and younger (orange) participants across all sliding window bins of unsigned relative error (b) and for a subset of smaller relative error bins (c). Significant differences (permutation tests for cluster mass, $n = 57$ per group, $P < 0.05$) are marked with dark points. (d) Differences in learning depended on the noise distribution. Learning rates predicted by the regression model for old (blue) and young (orange) participants are plotted separately for high (dark) and low (light) standard deviation conditions across bins of spatial error magnitude. Learning differences between groups emerged for both conditions for errors smaller than one standard deviation of the noise distribution (marked by dotted lines). In all panels, lines and shading reflect mean and s.e.m. respectively.

probability (surprise) were greater. Consistent with normative updating, we found that both age groups learned more when uncertainty was high, especially for intermediate relative prediction errors (Fig. 4a; permutation test for H_0 : uncertainty coefficient = 0, $n = 57$, $P < 0.001$ for young participants and $P < 0.005$ for older participants) albeit with substantial heterogeneity across individuals (Fig. 4d). Similarly, both groups increased learning as a function of surprise (Fig. 4b,e; permutation test for H_0 : surprise coefficient = 0, $n = 57$, $P < 0.001$ for both groups). However, younger participants adjusted learning according to uncertainty more than their older counterparts (Fig. 4a,d; permutation test for H_0 : equal mean uncertainty coefficients, $n = 57$ per group, $P < 0.005$). In contrast, older participants adjusted learning more in response to surprise than their younger counterparts (Fig. 4b, permutation test for H_0 : equal mean surprise coefficients, $n = 57$ per group, $P < 0.005$). Thus, our data show an age-related double dissociation with respect to the factors that govern adaptive learning: older adults show a reduced sensitivity to uncertainty but an enhanced sensitivity to surprise than younger adults.

Behaviour on ‘catch trials’ (trials on which the helicopter was visible) also showed age-related learning differences. Participants

could use this information as an additional cue to identify the true mean of the helicopter location. Using additional terms in our regression model, we found that both younger and older participants tended to make appropriate, additional updates towards the visible helicopter (Fig. 4c,f; permutation test for H_0 : coefficient = 0, $n = 57$, $P < 0.001$ for both groups) and away from the most recent bag position (Supplementary Fig. 3; permutation test for H_0 : coefficient = 0, $n = 57$, $P < 0.001$ for both groups). However, the groups differed in the extent to which they updated towards the helicopter: older participants showed less pronounced updating, particularly after the smallest errors (Fig. 4c; permutation test for H_0 : equal mean helicopter update coefficients, $n = 57$ per group, $P < 0.05$).

These learning differences between groups were reflective of age rather than differences in fluid intelligence or working memory. An explanatory model that included each of the age-related learning differences identified above (learning rate for unsurprising outcomes, uncertainty, surprise and helicopter updating) could explain differences in age across subjects ($F = 6.37$, $n = 114$, $P < 0.001$) and generated predictions that correlated with fluid intelligence, as assessed with Raven’s progressive matrices ($r = -0.31$, $P < 0.001$) and working

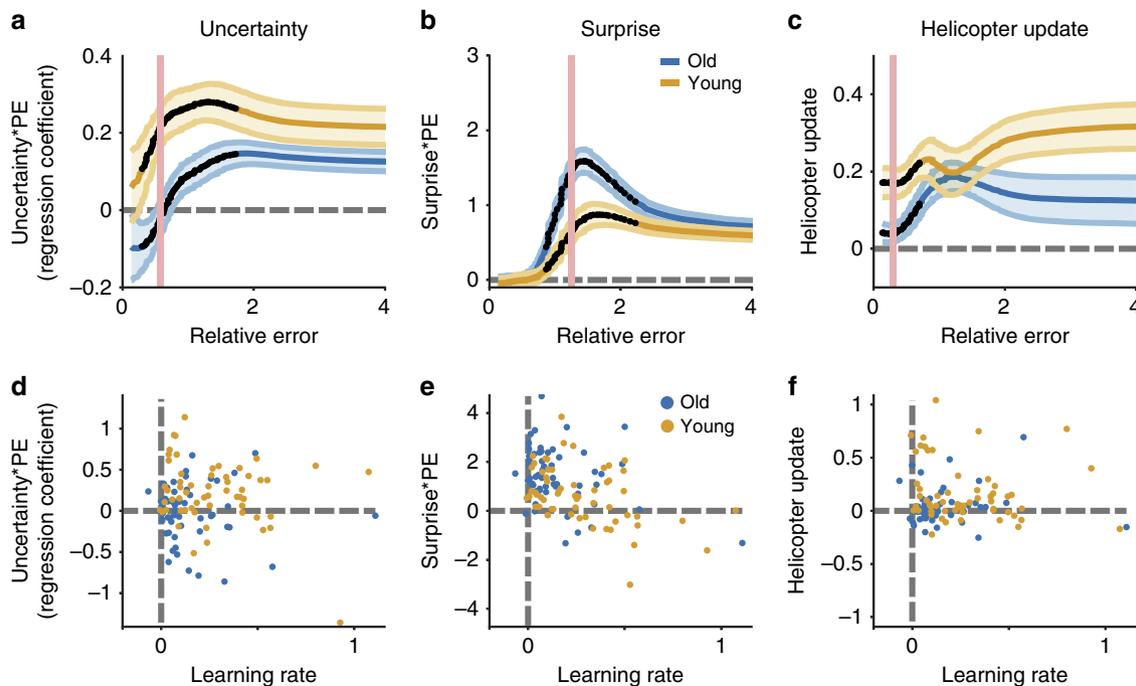


Figure 4 | Age differences in uncertainty and surprise-driven learning. Regression coefficients for interaction terms were estimated in sliding windows of outcomes ordered from the least to the most surprising. For **a-c**, regression coefficients are plotted on the ordinate against the median binned relative error for each sliding window on the abscissa. **(a)** Uncertainty regression coefficients (ordinate) for older (blue) and younger (orange) participants differed across a range of relative errors (abscissa). **(b)** Surprise regression coefficients (ordinate) were larger in older, relative to younger, participants for moderately large errors (abscissa). **(c)** Helicopter update coefficients (ordinate) show that both groups moved the bucket towards the helicopter on catch trials (positive values), but that young participants did so more than old participants after making small relative errors (abscissa). Mean/s.e.m. coefficient values are reflected by line/shading and dark points indicate bins including a significant group difference (permutation tests for cluster mass, $n = 57$ per group, $P < 0.05$). **(d-f)** Regression coefficients for individual young/old participants (orange/blue points) in relative error bins with maximal group differences (pink lines in **a-c**; ordinate) are plotted against learning rate coefficients. Note that group coefficient differences are not driven by the small group of subjects with extremely high learning rates.

memory as assessed by the operation span task ($r = -0.35$, $P < 0.001$). However, learning differences were not simply reflecting aspects of fluid intelligence or working memory that co-vary with age, as the task measurements were related to age even after accounting for these covariates (nested $F = 3.29$, $n = 114$, $P = 0.01$).

Learning differences are simulated by reduced uncertainty. To qualitatively assess the contribution of uncertainty underestimation to age-related differences in task performance, we modified the normative model by artificially reducing its estimate of uncertainty on each trial. We used both the normative model and the low-uncertainty model to generate data for the task sessions completed by our human participants and analysed the data from these models using the same regression framework. Furthermore, we followed the same procedure to simulate task behaviour using two other sub-optimal models, one insensitive to surprise and the other with under-estimated hazard rates, to compare the simulated behaviour to our empirical results.

The low-uncertainty model, unlike the other sub-optimal models, generated predictions that differed from those of the normative model in a manner that mimicked four key differences between younger and older participants. First, the low-uncertainty model showed lower learning rates specifically after small errors (Fig. 5a). Second, the low-uncertainty model modulated learning less according to uncertainty (Fig. 5b). Third, the low-uncertainty model was more sensitive to changes in surprise, particularly after the moderately sized errors where older

and younger participants differed most (Fig. 5c). Fourth, the low-uncertainty model was less prone to adjust expectations towards an additional cue representing the visible helicopter, especially in the absence of a surprisingly large prediction error (Fig. 5d). Neither of the other two sub-optimal models could reproduce the specific learning deficit for small errors or the double dissociation between uncertainty- and surprise-driven learning effects.

Reduced uncertainty best describes learning in older adults. To quantitatively test the idea that learning differences across the age groups arise as a result of uncertainty underestimation, we fit a flexible version of the normative learning model directly to participant behaviour. This flexible model contained free parameters to describe each of the computational deficits that could impact learning rate: (1) insensitivity to surprise, (2) mis-estimated hazard rate and (3) uncertainty underestimation. In addition to these learning terms, the model also contained two parameters to allow for behavioural variability: one term that allowed for variability in update magnitude (allowing for imprecise bucket placements) and another that allowed for variability in learning rate (allowing for bucket precision to decrease with error magnitude). For each subject, all parameters were fit simultaneously using maximum likelihood estimation and best fitting parameters were interpreted as quantitative estimates of the latent factors governing learning behaviour. Note that parameter fits from this model should be viewed as complimentary to, but not independent from, the results from

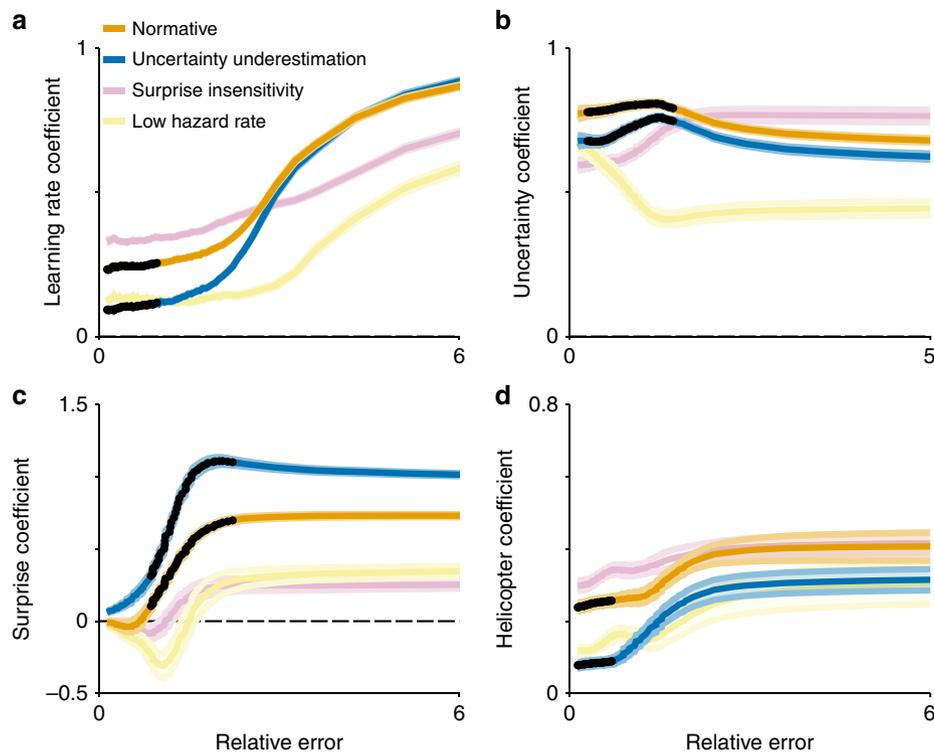


Figure 5 | Uncertainty underestimation mimics altered behaviour in older participants. Behaviour was simulated using a normative model (orange) and three deviant models: (1) surprise insensitivity (pink), (2) hazard rate underestimation (yellow) and (3) uncertainty underestimation (blue). Colours were chosen to emphasize that uncertainty underestimation deviates from the normative model in the same manner that older participants deviate from younger ones. **(a)** Mean/s.e.m. learning rate regression coefficients (line/shading) describing the effect of the prediction errors on the subsequent updates made by these models (see Fig. 3b). Only the uncertainty underestimation model shows a selective learning deficit in these bins. **(b)** Mean/s.e.m. uncertainty regression coefficients (line/shading) are plotted separately for each simulated model across all relative error bins (compare with Fig. 4a) **(c)** Mean/s.e.m. surprise regression coefficients (line/shading) are plotted separately for each simulated model across relative error bins (compare with Fig. 4b). **(d)** Mean/s.e.m. regression coefficients (line/shading) describing the effects of the helicopter over simulated model updates for trials when it was visible (for example, how much does the model adjust estimates towards the observed helicopter; compare with Fig. 4c). In all panels, dark points mark bins with significant differences between younger and older participants (permutation tests for cluster mass, $n = 57$ per group, all P values < 0.05) and lines/shading reflect mean/s.e.m. of simulated model behaviour.

our descriptive analysis as the parameter fits provide an aggregate measure of computational deficits based on the same underlying data used for the regression analysis above.

Parameter estimates for young and old participants confirmed age differences in uncertainty along with additional differences (Fig. 6). Maximum likelihood estimates of the uncertainty-underestimation parameter tended to be positive, indicating that most participants were best fit by models that failed to represent normative levels of uncertainty (mean \pm s.e.m. parameter estimates were 2.69 ± 0.23 and 1.52 ± 0.19 for old and young participants, respectively). Consistent with our descriptive results, this bias towards underestimating uncertainty was more pronounced in older participants ($t = 3.92$, $n = 57$ per group, $P < 0.001$). In addition, we found age-related differences in surprise sensitivity, which was higher in older adults (mean \pm s.e.m. parameter estimates were 0.51 ± 0.04 and 0.40 ± 0.03 for young and old participants, respectively; $t = 2.11$, $n = 57$ per group $P < 0.05$), and learning rate variability, which also was elevated in older adults as compared to younger adults ($t = 3.42$, $n = 57$ per group, $P < 0.001$). In contrast, we found no evidence for differences in hazard rate across the two groups ($t = 0.21$, $n = 57$ per group, $P > 0.9$).

One potential concern with using model-based parameter estimation to infer latent computational properties is that estimated parameters can be highly sensitive to the model in which they are embedded. In particular, the failure of a model to

account for key sources of variability can lead to biased parameter estimates³³. To examine the robustness of our modelling results, we constructed three additional models that progressively relax the assumptions about subjective perceptions of noise and in turn improved the ability of the model to account for behaviour. The first two models consider the possibility that subjective perceptions of noise are scaled (model 1) or scaled and offset (model 2) relative to the ground truth. The third model considers the possibility that noise is itself uncertain and predictive distributions are composed of a mixture of Gaussian distributions, each having a different width.

Consistent with previous work suggesting variability in perceptions of noise within and between subjects¹⁸, the most complex model containing within and between subject variability in noise estimates fit better than all simpler models, even after penalizing for additional parameters (Supplementary Fig. 4). The key advantage of this complex model appears to be that it effectively captures the shape of the relationship between learning rate and error magnitude (Supplementary Fig. 6). However, one disadvantage of this additional complexity is that it leads to parameter tradeoffs that decrease the recoverability of individual parameters, particularly the surprise sensitivity parameter (compare Supplementary Figs 5 and 7). Consistent with this lack of identifiability, surprise sensitivity did not differ across the age groups in any of these three models (all $P > 0.05$). Nonetheless, the uncertainty underestimation and learning rate

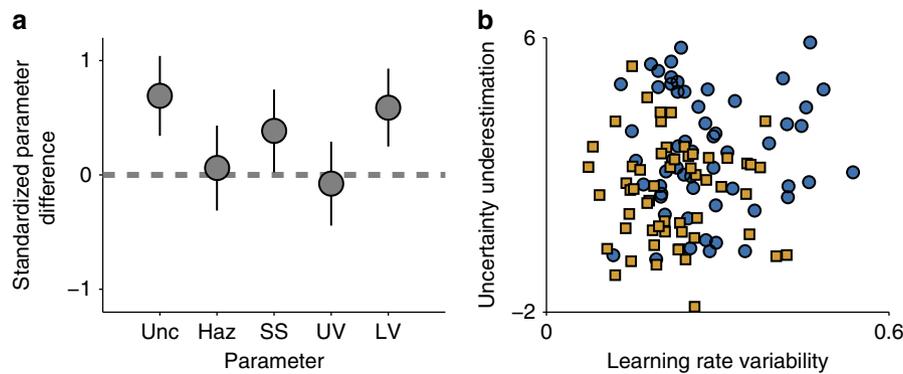


Figure 6 | Age-related differences in uncertainty underestimation. Participant updating behaviour was fit with a flexible version of the normative model. **(a)** Mean and 95% confidence intervals on standardized differences (young-old) for maximum likelihood parameter estimates of: (1) uncertainty underestimation (*Unc*), (2) expectations about the hazard of change points (*Haz*), (3) sensitivity to surprising outcomes (*SS*), (4) update variability (*UV*), (5) learning rate variability (*LRV*). **(b)** Maximum likelihood estimates of uncertainty underestimation (ordinate) and learning rate variability (abscissa) for young (orange squares) and old (blue circles) participants.

variability parameters were elevated for older relative to younger subjects in each of the three additional models (all $t > 3.4$, all $P < 0.001$ for *UV*, all $t > 3.1$ all $P < 0.005$ for *LRV*), suggesting that age differences in uncertainty underestimation and learning rate variability are robust to specific modelling choices.

Model estimates of uncertainty underestimation also offer a parsimonious description of age-related changes in learning. Subject-specific estimates of the uncertainty underestimation parameter were correlated with each of the four learning differences identified by our descriptive analysis (Pearson's $r = -0.49$, -0.20 , 0.46 and -0.29 for learning rate, uncertainty, surprise and helicopter coefficients, respectively; all P values < 0.05), suggesting that this metric can account for each of the features of the data that were identified using the descriptive model. Uncertainty underestimation fits were negatively correlated with fluid intelligence ($r = -0.268$, $P < 0.01$), but not with working memory ($r = -0.1287$, $P = 0.17$). Uncertainty underestimation also explained variance in age even after accounting for these potential covariates (nested $F = 5.92$, $P < 0.05$) and coefficients relating uncertainty underestimation to age were positive in a model that included these terms (mean \pm 95% confidence interval: 1.59 ± 1.28 years per unit uncertainty underestimation). Thus uncertainty underestimation offers a parsimonious explanation for age-related learning differences beyond the well-documented age differences in fluid abilities.

Discussion

Age-related deficits in learning have previously been attributed to differences in the computation of prediction errors^{3,5,6}. However, the data to support this hypothesis are somewhat contradictory and point to a more complicated scenario³⁴. Here we examined one such scenario: deficits in learning result from differences in how older individuals assign influence to new information according to environmental statistics. We identified three plausible computational changes that could give rise to such a learning deficit: (1) insensitivity to surprise, (2) underestimation of uncertainty and (3) underestimation of the hazard rate of change points. To formalize these predictions, we simulated learning behaviour in a predictive inference task using a normative model and independently manipulated each of these factors. We then asked younger and older adults to perform the task and tested the model predictions empirically.

We found that older adults displayed a selective deficit in learning that is best described by reduced uncertainty. Older

adults learned less from unsurprising outcomes than their younger counterparts, qualitatively matching the predictions from a low-uncertainty model of learning (compare Figs 2c and 4b). Moreover, older adults, like low-uncertainty learners, relied more heavily on surprise and less heavily on uncertainty in adjusting learning rates (compare Figs 4a,b and 5b,c). Consistent with the qualitative results of the regression analysis our quantitative model fitting suggests that age differences in learning can be explained by systematic underestimation of uncertainty by older adults (Fig. 6). Reduced uncertainty, as measured by parameter estimates from the flexible learning model, could explain differences in age even after accounting for age-related differences in working memory and reasoning abilities. These results support the overarching hypothesis that changes in learning across healthy aging result from changes in upstream computations necessary for determining how much to learn in a given situation. Furthermore, they suggest that age-related deficits in uncertainty-driven learning can be dissociated from a general decline in fluid abilities with age. Taken together, our data suggest that cognitive aging reduces subjective representations of uncertainty, which in turn diminishes learning under specific circumstances.

In addition, older adults in our study used more variable learning rates than their younger counterparts. This result is consistent with previous findings of age-related increases in behavioural variability³⁵. Older adults also tended to have relatively enhanced sensitivity to surprise. However, we were unable to clearly dissociate this effect from their reduced uncertainty estimates. If future work could confirm such an advantage, surprise sensitivity might represent a compensatory mechanism that could be exploited in designing learning environments that are tailored to the needs of the elderly population.

While our results are the first to demonstrate a failure of older adults to recruit requisite levels of uncertainty in the service of learning, previous work has hinted that this might be the case. For example, recent work suggests that older adults show deficits in probabilistic compared with deterministic learning tasks^{6,26–29}, in situations where feedback is ambiguous^{30,31} and during reversal learning^{27,32}. Despite these deficits, older adults show similar competence to younger adults in decision tasks that explicitly describe risks, suggesting that the deficit is in learning from probabilistic cues, rather than failing to act appropriately according to learned values^{9,11,12}. Our results suggest that these age-related impairments in learning can be attributed to a specific

deficit in the ability to recruit requisite levels of uncertainty to appropriately guide learning.

Why do older adults fail to represent sufficient levels of uncertainty? One possibility is that representing appropriate levels of uncertainty requires a cognitive and/or biological resource that decays across healthy aging. One obvious candidate for such a resource is working memory capacity³⁶. Working memory capacity declines as a function of healthy aging and has been identified as a potential source of age-related deficits in probabilistic learning^{37–39}. While previous work has focused on the role of working memory for the storage and selective recall of action–outcome contingencies, another use for such a system might be to store and recall plausible hypotheses about latent task states^{38,40}. In such a regime, decrements in working memory capacity would lead to fewer stored hypotheses, and in the extreme where only a single hypothesis is stored, a complete lack of uncertainty⁴¹. However, within our study we did not find strong evidence for a relationship between working memory capacity and uncertainty underestimation. Moreover, age-related deficits in uncertainty-driven learning persisted after controlling for age differences in working memory, arguing against this potential interpretation.

Another possibility is that older adults fail to represent sufficient levels of uncertainty because they have an aversion to uncertainty or the mental effort required to represent it. In descriptive lottery tasks younger and older adults are similarly averse to uncertainty about probabilities associated with possible gains suggesting that this explanation is unlikely to play a major role⁴². More generally, it is possible that older subjects are more averse to the expenditure of mental effort required to maintain an accurate uncertainty estimate and instead rely on a simpler learning strategy⁴³. From our computational model it is not clear why representing low levels of uncertainty would be any easier than representing any other fixed level of uncertainty: the computational costs within our model are associated with updating, rather than maintaining, uncertainty estimates. One potential source of cognitive costs could be in the representation of the task model itself. There is some evidence that older adults tend to rely less on model-based strategies for learning⁴⁴. In contrast, they tend to rely more on external cues to guide behaviour, which, in our task, could correspond to adopting a strategy like the ‘win-stay lose-shift’ heuristic commonly used by older adults in choice tasks^{45,46}. This general idea corresponds well to the double dissociation noted in our descriptive analysis of learning behaviour: Older subjects are more responsive to environmental learning cues (surprise) but much less responsive to internally generated ones (uncertainty).

The crucial factor limiting uncertainty representations in older adults could, and at some level must be, biological in nature. One candidate for such a limiting factor is norepinephrine, a neuromodulator thought to track uncertainty in changing environments¹⁷. There is some evidence that low-level arousal systems including brainstem noradrenergic nucleus locus coeruleus (LC) are modulated by uncertainty and may mediate its effects on learning. Pupil diameter, which is thought to reflect LC activity, is modulated by uncertainty in both learning and exploration tasks^{16,47,48}. Activation of LC, and the corresponding increase in cortical norepinephrine, increases the signal-to-noise ratio in sensory neurons and may enhance learning by increasing the gain of cortical units representing newly arriving information^{17,49–51}. A similar theory has already been proposed to explain age-related differences in learning^{2,52}. Whereas previous accounts of the change in gain of information processing across age have focused on dopamine depletion as a possible explanation, there is also evidence that noradrenergic

signalling may be dampened across healthy aging^{53–55}, providing a potential link between our findings, the relationship between arousal and uncertainty, and the decreased cortical gain theory of aging.

Although it is tempting to link age-related changes in representing uncertainty to reduction of a single neurotransmitter, several alternative biological accounts exist. Functional imaging studies have highlighted uncertainty representations in prefrontal areas including the anterior prefrontal cortex (aPFC) and orbitofrontal cortex (OFC)^{15,56–58}. There is substantial evidence for changes in the function and structure of the prefrontal cortex across healthy aging⁵⁹. For example, prefrontal regions are substantially under-recruited in older versus younger adults when learning higher-order task structures⁶⁰. Such under-recruitment could limit top-down activation of a cortical learning network during uncertain regimes¹⁵. Alternatively, diminished prefrontal recruitment during uncertainty could contribute directly to muted dopaminergic prediction error signalling, as prefrontal inputs play critical roles in shaping these signals⁶¹.

The existence of supporting evidence for both prefrontal and noradrenergic correlates of uncertainty representation highlights the need for future work combining computational methods with biological measurements and interventions that could unravel the underlying causal relationships between these factors and learning. Future studies should also investigate the extent to which biomarkers for uncertainty interact with the magnitude of reward-prediction error signals in the striatum. Taken in the context of previous work, our findings suggest that reward-prediction error signals in the striatum may be enhanced by neuromodulatory (LC) or cortical (aPFC) uncertainty signals. This would explain why prediction error signalling is diminished in older adults only under conditions of uncertainty, as these are the only conditions where reduced uncertainty representations would come into play^{3,5–8,58}. While this mechanism provides a parsimonious explanation for previous findings, alternative accounts involving uncertainty representations computed locally in the striatum or the ventral tegmental areas are also plausible and should be tested through age comparisons of uncertainty modulations in the BOLD signal^{62,63}.

Our results may also have implications beyond healthy aging for the understanding of a number of mental disorders for which learning deficits are a hallmark. In some cases, such deficits may not reflect an inability to learn but rather specific deficits in computational processes that govern how much to learn from new information. In addition to our results regarding aging, there has been recent support for this idea with respect to the effects of trait anxiety on learning⁶⁴, but to date it is unknown whether learning abnormalities in conditions such as attention deficit hyperactivity disorder (ADHD), autism or schizophrenia are also driven by higher-order computational factors. Our task and modelling framework provide a means to address these issues⁶⁵.

To summarize, in the current study we apply a normative model of predictive inference to simulate possible effects of aging on three factors that are thought to govern learning: uncertainty, surprise and hazard rate. Using behavioural data and quantitative model fitting we show that learning deficits in older adults are best characterized by an underestimation of uncertainty rather than a generic reduction of learning. This finding provides a parsimonious mechanistic explanation for age-related impairments in learning across a variety of tasks. Furthermore, it highlights specific cortical and subcortical regions involved in representing uncertainty as candidates for mediating age-related learning deficits. We hope that this work facilitates future studies aiming to understand the neural underpinnings of limited uncertainty representation and age-related changes thereof.

Methods

Participants. 59 younger and 63 older adults took part in the study. Target sample size was pre-determined according to age effect sizes in previous studies of learning and decision-making. Six older and two younger adults were excluded because of insufficient data (<300 trials of predictive inference task completed). Thus, the effective sample consisted of 57 younger adults (mean age: 24.5 years, age range: 20–30 years, 29 female) and 57 older adults (mean age: 69.2 years, age range: 56–80 years, 26 female). Participants gave written informed consent. The Institutional Review Board of the Max Planck Institute for Human Development approved the study. In addition to the experimental task, participants completed a biographical and a personality questionnaire and several psychometric tests: (1) Identical pictures test; (2) Raven’s Progressive matrices⁶⁶; (3), Spot-a-Word test; and (4) the Operations span task (OSPAN)⁶⁷. As shown in Table 1 older adults had lower scores on the Identical pictures test, Raven’s matrices and the OSPAN task than younger adults (P values <0.001, η^2_G > 0.21). In contrast, older adults obtained higher scores than younger adults on the Spot-a-Word test (P < 0.001, η^2_G = 0.20). Consistent with previous findings from larger population-based samples, these results suggest age-related reductions in fluid intelligence and age-related improvements in crystallized intelligence⁶⁸.

Procedure. Participants performed two sessions, which were separated by a minimum of 1 week and a maximum of 3 weeks. In the first session, participants completed a biographical questionnaire, the BIS/BAS personality questionnaire, Raven’s progressive matrices⁶⁶ as well as a two-state Markov decision task⁶⁹, data of which are presented in ref. 44. In the second session, participants performed the predictive inference (Helicopter) task¹⁵, the OSPAN task, the Spot-a-Word and the Identical pictures test as well as a version of the two-state Markov task.

Predictive inference task. Participants completed two blocks (200 trials each) of a computerized predictive inference task programmed in Matlab (The MathWorks, Natick, MA) using MGL (<http://justingardner.net/mgl>) and snowDots (<http://code.google.com/p/snow-dots>) extensions. The predictive inference task required inferring the mean of a noisy variable that underwent occasional change points¹⁵. The problem was embedded in a cover story involving a virtual helicopter (mean) that moved occasionally (change points) and dropped a bag from the sky on each trial (noisy variable).

On each trial the participant moved a bucket to the most likely position of the helicopter using a keyboard (Fig. 1a). After the bucket position was confirmed through a key press, the participant observed a bag fall from the top of the screen followed by an explosion that revealed the contents of the bag (200 gold coins or silver rocks; randomized across trials) and the extent to which those contents were collected in the bucket (ranging from 0–200 depending on the distance between the bucket and the bag). Gold tokens (but not rocks) collected in the bucket were translated into incentive payments at the end of the task. The horizontal position of the bag was denoted with a grey tick mark on the screen and the distance between the bag and bucket (prediction error) denoted by a red line. These markings served to eliminate working memory requirements and allowed subjects access to all relevant information in choosing how much to adjust the bucket position for the subsequent trial.

The horizontal position of each bag (represented on a numerical scale from 0 to 300 for convenience) was drawn from a normal distribution with a mean corresponding to the position of a virtual helicopter hovering in the sky and a s.d. that was manipulated blockwise (10 or 25; counterbalanced for order). On most trials the helicopter would remain stationary, but on a small fraction of trials (ground truth hazard rate; 1/10) it would relocate to a new screen position. On the vast majority of trials the helicopter was ‘hidden’ by clouds. Occasionally, the helicopter was revealed visually (catch trials; 1/10). In principle, the visible helicopter could provide perfect information about the mean of the distribution, but in practice the centre of the helicopter was not obvious due to asymmetry in the cartoon helicopter image and the vertical distance between this image and that of

the bucket (Supplementary Fig. 1). Participants were instructed to infer the location of the helicopter based on previous observations (bag and helicopter positions) and to place the bucket directly underneath it.

Training. Before completing two blocks of the predictive inference task participants went through a series of training tasks that slowly built the complex task from simpler elements. As in the experimental session, every training task consisted of a low and high standard deviation (noise) block (counterbalanced for order). In the first training task the helicopter was completely visible and thus bag locations were not necessary to guide behaviour. To ensure that participants understood that the helicopter is the best outcome predictor we used a response criterion that required participants to put their bucket ten times exactly underneath the visible helicopter. Each noise block stopped after either the criterion was reached or after a maximum of 80 trials. In the second training task with two 50-trial runs clouds covered the helicopter and occasionally disappeared during catch trials. This version of the task was the same as the experimental task except that participants would not earn money for their collected coins. Overall performance, in terms of coins collected, did not differ across age groups (Supplementary Fig. 2).

Computational modelling. To dissociate surprise-driven updating from uncertainty-driven updating we extended an existing normative model for learning in a dynamic environment that has been described in detail previously^{15,16,18}. In brief, this model approximates optimal inference by tracking two factors that should drive learning: change-point probability (the probability with which a change in the helicopter location occurs) and uncertainty (the reliability with which an outcome reflects the true location of the helicopter). Here we extend this model in four ways. First, we develop a new method for estimating change-point probability and uncertainty that captures subjective differences in experienced surprise. Second, we extend the generative framework and corresponding inference equations of the model to incorporate catch trials. Third, we extend the normative model to allow for specific deviations from normativity including surprise insensitivity, incorrect hazard rate assumptions, and uncertainty underestimation. Finally, we extend the model to consider more complex models of behaviour that allow for subjective differences in the representation of noise.

The first extension of the previously described computational methods allowed for subjective estimates change-point probability and uncertainty. Previous studies have run the normative model over trial outcomes to get trial-by-trial estimates of these quantities¹⁶; however, one issue with this approach is that since participant and model predictions do not always perfectly match, an outcome that constitutes a small and unsurprising error for the model might actually be a large and rather surprising one for the participant. To avoid this potential problem we obtained subjective measures of change-point probability and uncertainty by running the normative model across the prediction errors experienced by participants, rather than the outcomes that generated them. Model variables were computed recursively by first determining the uncertainty about the current helicopter location according to the relative uncertainty, change-point probability and prediction error from the previous trial:

$$\sigma_{\mu}^2 = \Omega_t \sigma_N^2 + (1 - \Omega_t) \sigma_N^2 \tau_t + \Omega_t (1 - \Omega_t) \delta_t (1 - \tau_t) \tag{1}$$

Where σ_{μ}^2 is the variance on the predictive distribution over possible helicopter locations, σ_N^2 is the variance on the distribution over bag locations (noise), Ω_t is the probability of a change point on the previous trial (that is, the probability with which the helicopter has relocated between trials), τ_t is the relative uncertainty from the previous trial and δ_t is the prediction error from the previous trial. Relative uncertainty was computed by expressing uncertainty about the helicopter location as a fraction of total uncertainty about the location of the next bag:

$$\tau_{t+1} = \frac{\sigma_{\mu}^2}{\sigma_{\mu}^2 + \sigma_N^2} \tag{2}$$

Where τ_{t+1} is the relative uncertainty for trial $t + 1$. This relative uncertainty estimate, along with the variance on the bag distribution (noise; σ_N^2) was used to calibrate the change-point probability associated with each new prediction error:

$$\Omega_{t+1} = \frac{H/300}{H/300 + \mathcal{N}(\delta_{t+1} | 0, \frac{\sigma_N^2}{1 - \tau_{t+1}})(1 - H)} \tag{3}$$

Where H is the hazard of a change point (0.1) and δ_{t+1} is the new prediction error. Subjective estimates of change point probability and relative uncertainty were computed by evaluating equations 1 and 2 according to the trial-by-trial prediction errors made by each individual subject.

The second extension of the model was necessary to account for additional information provided on catch trials in which the helicopter is visible to participants. To maintain the deterministic nature of the model but also account for perceptual ambiguity associated with the helicopter image we treat the visible helicopter as a cue indicating a Gaussian likelihood function centred on the ground truth (mean of the bag distribution). We allow the variance of the Gaussian to be adjusted to account for behaviours ranging from completely trusting the helicopter information to completely ignoring it. Combining this additional cue with the information provided by the bag itself led to the following additional equations that

Psychometric measure	Younger adults (N = 57, mean age: 24.5, range: 20–30)	Older adults (N = 57, mean age: 69.1, range: 56–80)
	(mean, s.e.m.)	(mean, s.e.m.)
Raven (raw score)	12.3 (0.4)	6.3 (0.4)
Ospan total score	54.3 (1.7)	39.9 (2.2)
Processing speed	31.1 (0.6)	21.2 (0.5)
Spot-a-Word	19.9 (0.8)	25.7 (0.7)

were implemented at the end of each helicopter visible trial to update position estimates:

$$B_t = w_t B_t + (1 - w_t) \mu_t \quad (4)$$

$$w_t = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_H^2} \quad (5)$$

Where B_t is the belief of the model about the true mean of the distribution and w_t reflects the weight of the current belief in a weighted mixture of the current belief and the true mean (μ) as indicated by the helicopter. w_t is determined according to the relative variances on the current predictive (σ_μ) and helicopter centred likelihood distributions (σ_H).

In addition, the following equations were implemented to reduce the relative uncertainty estimates on trials where the helicopter was observable:

$$\tau_t = \frac{\sigma_*^2}{\sigma_*^2 + \sigma_N^2} \quad (6)$$

Where σ_*^2 is the variance on the predictive distribution over possible helicopter locations after correcting for additional information provided by the visible helicopter:

$$\sigma_*^2 = \left(\frac{\sigma_\mu^2}{1} + \frac{\sigma_H^2}{1} \right)^{-1} \quad (7)$$

Where σ_μ^2 and σ_H^2 are the variances associated with the internal prediction and the perceptual information provided by the visible helicopter, respectively.

The third extension of the normative model served to allow for specific deviations from optimal behaviour. We simulated behaviour from four versions of the normative model: (1) a version using the update equations described previously^{15,16} with the modifications described above, (2) a model with diminished sensitivity to surprise created by raising the change point likelihood to a power between 0 and 1 (0.2 for figures) as described previously¹⁸, (3) a low hazard rate model expecting change points to be rare (H was set to 0.001) and (4) an uncertainty underestimation model in which uncertainty was reduced after each observed bag drop by dividing the estimated variance on the predictive distribution over possible helicopter locations (σ_μ^2) by a constant on each trial (10 for simulations).

Flexible versions of the normative model were fit directly to behaviour and used to infer maximum likelihood estimates of (1) hazard rate, (2) surprise sensitivity and (3) uncertainty underestimation, which were then used to identify age-related differences in these computational factors. For the purposes of model fitting, participant updates were defined to be sampled from a normal distribution with a mean equal to the model predicted update and a s.d. that was a linear function of the absolute prediction error magnitude. The intercept and slope of this linear function were fit as free parameters and can be thought of as variability in the motor update and learning rate selection processes respectively. Thus, the minimally complex model contained five free parameters, three of which were related to learning and two of which were related to response variability. This model fit better than several more constrained ones in which parameters were fixed to their normative values (Supplementary Fig. 4).

In addition, more complex models were constructed that considered potential sources of variability related to the perception of noise. These complex models included all of the basic variables as well as one or more of the following free parameters: (1) a multiplicative scaling term to allow for scaled perceptions of noise, (2) an additive offset term allowing for subjective biases in overall levels of noise perception and (3) a noise variability term allowing for individual subjects to represent a distribution across possible noise values. Since there were only two noise conditions, including additive and multiplicative scale factors amounted to allowing the noise for each block type to be fit as a free parameter. Within the model that accounted for noise variability, the likelihood of observations was not drawn from a single normal distribution (as described in equation 3), but instead from a weighted mixture of normal distributions, where each component of the mixture had a mean of zero and a s.d. equal to a scaled version of the total uncertainty. Scale values were represented as uniformly spaced points on a grid (ranged 0.1–100) with associated probabilities drawn from an inverse gamma distribution. The shape term of the gamma distribution was fit as a free parameter and can be thought of as conveying the amount of evidence for the expected noise distribution, with lower values indicating more uncertainty over possible noise values.

All models were fit using a constrained search algorithm (fmincon in Matlab) that maximized the total log posterior probability of participant updates given participant prediction errors and parameter estimates. Weak priors favouring normative learning parameters were used to regularize parameter estimates. Uncertainty underestimation estimates were positively skewed and thus reported and analysed in log units. All model-fitting code will be made available on request.

Data analysis. Participant bucket placements and trial outcomes were used to compute trial-by-trial prediction errors (δ_t):

$$\delta_t = \chi_t - B_t \quad (8)$$

where χ_t and B_t are the locations of the dropped bag and placed bucket on trial t ,

respectively. The corresponding updates made by the participant in response to each prediction error were computed as:

$$\text{Update}_t = B_{t+1} - B_t \quad (9)$$

The first and last trials of each block were omitted from further analysis, as updates on these trials were likely to be influenced by block changes. Trials where the prediction error equalled zero were also omitted, as they provide no information about error-driven learning. In addition trials where bucket placement fell more than 15 screen units away from any possible delta rule update towards the previous bag or helicopter position were omitted, as they were considered to be governed by a process other than error-driven learning. 1.1% of trials were removed in this way.

Trial-by-trial updates were analysed with a regression model that included trial-by-trial prediction errors to account for overall learning rate, as well as the interaction of prediction error with five mean-centred factors: (1) surprise (change-point probability as computed above), (2) uncertainty (relative uncertainty as computed above), (3) noise (s.d. of bag distribution), (4) trial value (gold versus rocks) and (5) helicopter visibility (binary variable describing whether helicopter cue was provided). To allow for updates towards the visible helicopter on catch trials, the model also included the interaction between the true mean of the distribution and the helicopter visibility variable. An additional nuisance term was also included to account for a slight bias in bucket placements towards the centre of the screen. One potential shortcoming of this regression model is that the residuals are heteroscedastic; specifically, absolute residuals are larger on trials where participants made larger absolute prediction errors. To account for this, we used an initialization regression for each participant and pooled the residuals to compute the variance over residuals across sliding windows of absolute prediction error magnitude. These variance estimates were used to weight the errors in a weighted regression equation that also included a ridge penalty to regularize coefficient estimates:

$$b = (A^T P A + R^T R) / (A^T P \times \text{Update}) \quad (10)$$

where A is the explanatory matrix, P is the inverse variance matrix, and R is a regularization matrix constructed with the ridge parameter equal to 0.1.

To identify specific learning differences predicted from the normative model (Fig. 2), we applied the penalized weighted regression model to data that were binned in sliding windows according to the size of the absolute prediction error made by the participant divided by standard deviation of the bag distribution, which served as a proxy for surprise (Fig. 3a). Each bin contained 10% of the total data and successive bins had lower and upper bounds that were incremented by a single percentile resulting in 90 total bins.

Regression coefficients were smoothed across bins and t -tests were used to identify 'clusters' of contiguous bins for which the p -value was smaller than 0.05. This procedure was repeated for three separate tests to reject the null hypotheses that (1) coefficients from older participants are different than zero, (2) coefficients from younger participants are different from zero and (3) coefficients from younger participants are different from those from the older participants. For each cluster, we computed cluster mass as the size of the cluster (number of bins) times the average absolute t -statistic within that cluster. For each test statistic a null distribution over cluster mass was generated by creating 10,000 permutations of the data (using sign-flipping for single group tests and label-flipping for the group comparison). Cluster corrected permutation tests were conducted by comparing the observed cluster mass for each test statistic against the null distribution created through these permutations. See Supplementary Fig. 3 for estimates of parameters that did not differ across the age groups.

Single participant coefficients were extracted for each coefficient that was significantly different across age groups according to a leave-one-subject-out (LOSO) procedure: Coefficients for each participant were extracted from the error bin that corresponded to the maximum absolute t -statistic from a between groups t -test across all bins for all other participants. These LOSO coefficient estimates were included as explanatory variables in a regression on participant age. Specifically, we created four distinct explanatory models containing: (1) only an intercept term, (2) LOSO coefficients and an intercept, (3) LOSO coefficients, Raven's scores, OSPAN scores and an intercept. Nested F -tests were used to compare the fits of these different models while accounting for differences in complexity.

References

- Burke, S. N. & Barnes, C. A. Neural plasticity in the ageing brain. *Nat. Rev. Neurosci.* **7**, 30–40 (2006).
- Eppinger, B., Hämmerer, D. & Li, S.-C. Neuromodulation of reward-based learning and decision making in human aging. *Ann. NY Acad. Sci.* **1235**, 1–17 (2011).
- Chowdhury, R. et al. Dopamine restores reward prediction errors in old age. *Nat. Neurosci.* **16**, 648–653 (2013).
- Samanez-Larkin, G. R., Levens, S. M., Perry, L. M., Dougherty, R. F. & Knutson, B. Frontostriatal white matter integrity mediates adult age differences in probabilistic reward learning. *J. Neurosci.* **32**, 5333–5337 (2012).
- Eppinger, B., Schuck, N. W., Nystrom, L. E. & Cohen, J. D. Reduced striatal responses to reward prediction errors in older compared with younger adults. *J. Neurosci.* **33**, 9905–9912 (2013).

6. Samanez-Larkin, G. R., Worthy, D. A., Mata, R., McClure, S. M. & Knutson, B. Adult age differences in frontostriatal representation of prediction error but not reward outcome. *Cogn. Affect. Behav. Neurosci.* **14**, 672–682 (2014).
7. Samanez-Larkin, G. R. *et al.* Anticipation of monetary gain but not loss in healthy older adults. *Nat. Neurosci.* **10**, 787–791 (2007).
8. Schott, B. H. *et al.* Ageing and early-stage Parkinson's disease affect separable neural mechanisms of mesolimbic reward processing. *Brain* **130**, 2412–2424 (2007).
9. Mata, R., Josef, A. K., Samanez-Larkin, G. R. & Hertwig, R. Age differences in risky choice: a meta-analysis. *Ann. NY Acad. Sci.* **1235**, 18–29 (2011).
10. Samanez-Larkin, G. R. & Knutson, B. in *The Neuroscience of Risky Decision Making* (eds Reyna, V. F. & Zayas, V.) (American Psychological Association, 2014).
11. Denburg, N. L., Tranel, D. & Bechara, A. The ability to decide advantageously declines prematurely in some normal older persons. *Neuropsychologia* **43**, 1099–1106 (2005).
12. Zamarian, L., Sinz, H., Bonatti, E., Gamboz, N. & Delazer, M. Normal aging affects decisions under ambiguity, but not decisions under risk. *Neuropsychology* **22**, 645–657 (2008).
13. Behrens, T. E. J., Woolrich, M. W., Walton, M. E. & Rushworth, M. F. S. Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–1221 (2007).
14. Payzan-LeNestour, E., Dunne, S., Bossaerts, P. & O'Doherty, J. P. The neural representation of unexpected uncertainty during value-based decision making. *Neuron* **79**, 191–201 (2013).
15. McGuire, J. T., Nassar, M. R., Gold, J. I. & Kable, J. W. Functionally dissociable influences on learning rate in a dynamic environment. *Neuron* **84**, 870–881 (2014).
16. Nassar, M. R. *et al.* Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* **15**, 1040–1046 (2012).
17. Yu, A. J. & Dayan, P. Uncertainty, neuromodulation, and attention. *Neuron* **46**, 681–692 (2005).
18. Nassar, M. R., Wilson, R. C., Heasley, B. & Gold, J. I. An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* **30**, 12366–12378 (2010).
19. O'Reilly, J. X. Making predictions in a changing world—inference, uncertainty, and learning. *Front. Neurosci.* **1**–10 (2013).
20. Yu, A. & Dayan, P. in *Advances in Neural Information Processing* (eds Becker, S. T. S. & Obermayer, K.) (MIT Press, 2003).
21. Wilson, R. C., Nassar, M. R. & Gold, J. I. Bayesian online learning of the hazard rate in change-point problems. *Neural Comput.* **22**, 2452–2476 (2010).
22. O'Reilly, J. X. *et al.* Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proc. Natl Acad. Sci. USA* **110**, E3660–E3669 (2013).
23. Payzan-LeNestour, E. & Bossaerts, P. Risk, unexpected uncertainty, and estimation uncertainty: bayesian learning in unstable settings. *PLoS Comput. Biol.* **7**, e1001048 (2011).
24. Adams, R. P. & MacKay, D. *Bayesian Online Change-point Detection* (University of Cambridge, 2007).
25. Hayden, B. Y., Heilbronner, S. R., Pearson, J. M. & Platt, M. L. Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *J. Neurosci.* **31**, 4178–4187 (2011).
26. Eppinger, B., Kray, J., Mock, B. & Mecklinger, A. Better or worse than expected? Aging, learning, and the ERN. *Neuropsychologia* **46**, 521–539 (2008).
27. Weiler, J. A., Bellebaum, C. & Daum, I. Aging affects acquisition and reversal of reward-based associative learning. *Learn. Mem.* **15**, 190–197 (2008).
28. Hämmerer, D., Li, S.-C., Müller, V. & Lindenberger, U. Life span differences in electrophysiological correlates of monitoring gains and losses during probabilistic reinforcement learning. *J. Cogn. Neurosci.* **23**, 579–592 (2011).
29. Pietschmann, M., Endrass, T., Czerwon, B. & Kathmann, N. Aging, probabilistic learning and performance monitoring. *Biol. Psychol.* **86**, 74–82 (2011).
30. Eppinger, B. & Kray, J. To choose or to avoid: age differences in learning from positive and negative feedback. *J. Cogn. Neurosci.* **23**, 41–52 (2011).
31. Herbert, M., Eppinger, B. & Kray, J. Younger but not older adults benefit from salient feedback during learning. *Front. Psychol.* **2**, 171 (2011).
32. Mell, T. *et al.* Effect of aging on stimulus-reward association learning. *Neuropsychologia* **43**, 554–563 (2005).
33. Nassar, M. R. & Gold, J. I. A healthy fear of the unknown: perspectives on the interpretation of parameter fits from computational models in neuroscience. *PLoS Comput. Biol.* **9**, e1003015 (2013).
34. Samanez-Larkin, G. R. & Knutson, B. Decision making in the ageing brain: changes in affective and motivational circuits. *Nat. Rev. Neurosci.* **16**, 278–289 (2015).
35. MacDonald, S. W. S., Hultsch, D. F. & Dixon, R. A. Performance variability is related to change in cognition: evidence from the Victoria Longitudinal Study. *Psychol. Aging* **18**, 510–523 (2003).
36. Cowan, N. The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav. Brain Sci.* **24**, 87–114 discussion 114–85 (2001).
37. Salthouse, T. A. Working memory as a processing resource in cognitive aging. *Dev. Rev.* **10**, 101–124 (1990).
38. Collins, A. G. E. & Frank, M. J. How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *Eur. J. Neurosci.* **35**, 1024–1035 (2012).
39. Beitz, K. M., Salthouse, T. A. & Davis, H. P. Performance on the Iowa Gambling Task: From 5 to 89 years of age. *J. Exp. Psychol. Gen.* **143**, 1677–1689 (2014).
40. Collins, A. G. E., Brown, J. K., Gold, J. M., Waltz, J. A. & Frank, M. J. Working memory contributions to reinforcement learning impairments in schizophrenia. *J. Neurosci.* **34**, 13747–13756 (2014).
41. Daw, N. & Courville, A. The pigeon as particle filter. *Adv. Neural Inf. Process. Syst.* **20**, 369–376 (2008).
42. Tymula, A., Rosenberg Belmaker, L. A., Ruderman, L., Glimcher, P. W. & Levy, I. Like cognitive function, decision making across the life span shows profound age-related changes. *Proc. Natl Acad. Sci. USA* **110**, 17143–17148 (2013).
43. Westbrook, A., Kester, D. & Braver, T. S. What is the subjective cost of cognitive effort? load, trait, and aging effects revealed by economic preference. *PLoS ONE* **8**, e68210 (2013).
44. Eppinger, B., Walter, M., Heekeren, H. R. & Li, S.-C. Of goals and habits: age-related and individual differences in goal-directed decision-making. *Front. Neurosci.* **7**, 253 (2013).
45. Lindenberger, U. & Mayr, U. Cognitive aging: is there a dark side to environmental support? *Trends Cogn. Sci.* **18**, 7–15 (2014).
46. Worthy, D. A. Age-based differences in strategy use in choice tasks. *Front. Neurosci.* **5**, 1–10 (2012).
47. Joshi, S., Li, Y., Kalwani, R. M. & Gold, J. I. Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron* **89**, 1–15 (2015).
48. Jepma, M. & Nieuwenhuis, S. Pupil diameter predicts changes in the exploration-exploitation trade-off: evidence for the adaptive gain theory. *J. Cogn. Neurosci.* **23**, 1587–1596 (2011).
49. Waterhouse, B. D. & Woodward, D. J. Interaction of norepinephrine with cerebrocortical activity evoked by stimulation of somatosensory afferent pathways in the rat. *Exp. Neurol.* **67**, 11–34 (1980).
50. Servan-Schreiber, D., Printz, H. & Cohen, J. D. A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science* **249**, 892–895 (1990).
51. Aston-Jones, G. & Cohen, J. D. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* **28**, 403–450 (2005).
52. Li, S.-C. & Rieckmann, A. Neuromodulation and aging: implications of aging neuronal gain control on cognition. *Curr. Opin. Neurobiol.* **29**, 148–158 (2014).
53. Chan-Palay, V. & Asan, E. Quantitation of catecholamine neurons in the locus coeruleus in human brains of normal young and older adults and in depression. *J. Comp. Neurol.* **287**, 357–372 (1989).
54. Manaye, K. F., McIntire, D. D., Mann, D. M. & German, D. C. Locus coeruleus cell loss in the aging human brain: a non-random process. *J. Comp. Neurol.* **358**, 79–87 (1995).
55. Vijayashankar, N. & Brody, H. A quantitative study of the pigmented neurons in the nuclei locus coeruleus and subcoeruleus in man as related to aging. *J. Neuropathol. Exp. Neurol.* **38**, 490–497 (1979).
56. Badre, D., Doll, B. B., Long, N. M. & Frank, M. J. Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron* **73**, 595–607 (2012).
57. Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879 (2006).
58. Schultz, W. *et al.* Explicit neural signals reflecting reward uncertainty. *Phil. Trans. R Soc. B* **363**, 3801–3811 (2008).
59. Nyberg, L. *et al.* Longitudinal evidence for diminished frontal cortex function in aging. *Proc. Natl Acad. Sci. USA* **107**, 22682–22686 (2010).
60. Eppinger, B., Heekeren, H. R. & Li, S.-C. Age-related prefrontal impairments implicate deficient prediction of future reward. *Neurobiol. Aging* **36**, 2380–2390 (2015).
61. Takahashi, Y. K. *et al.* Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nat. Neurosci.* **14**, 1590–1597 (2011).
62. Tobler, P. N., Fiorillo, C. D. & Schultz, W. Adaptive coding of reward value by dopamine neurons. *Science* **307**, 1642–1645 (2005).
63. Franklin, N. T. & Frank, M. J. A cholinergic feedback circuit to regulate striatal population uncertainty and optimize reinforcement learning. *eLife* **4**, e12029 (2015).
64. Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X. & Bishop, S. J. Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nat. Neurosci.* **18**, 590–596 (2015).
65. Maia, T. V. & Frank, M. J. From reinforcement learning models to psychiatric and neurological disorders. *Nat. Neurosci.* **14**, 154–162 (2011).
66. Raven, J., Raven, J. C. & Court, J. H. *Manual for Raven's Progressive Matrices and Vocabulary Scales* (Oxford Psychologists Press, Oxford, UK, 1998).
67. Unsworth, N., Heitz, R. P., Schrock, J. C. & Engle, R. W. An automated version of the operation span task. *Behav. Res. Methods* **37**, 498–505 (2005).

68. Li, S.-C. *et al.* Transformations in the couplings among intellectual abilities and constituent cognitive processes across the life span. *Psychol. Sci.* **15**, 155–163 (2004).
69. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-based influences on humans' choices and striatal prediction errors. *Neuron* **69**, 1204–1215 (2011).

Acknowledgements

We thank Ben Heasley for programming the task and Michael Frank, Joseph Kable, Joseph McGuire, and Yin Li for helpful comments. This work was supported by NIH grant R01-MH098899 to J.I.G. and F32-MH102009-01A1 to M.R.N. and a travel grant provided by the Burroughs Wellcome Fund. S.-C.L. H.R.H. and B.E. were in part supported by grants from the German Federal Ministry of Education and Research (BMBF; Grant numbers: FKZ 01GQ0913, FKZ 01GQ1313).

Author contributions

M.R.N., J.I.G., S.-C.L., H.H. and B.E. designed the study, R.B. and B.E. collected the data, M.R.N. and R.B. analysed the data, M.R.N., R.B., J.I.G., S.-C.L., H.H. and B.E. wrote the manuscript.

Additional information

Supplementary Information accompanies this paper at <http://www.nature.com/naturecommunications>

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

How to cite this article: Nassar, M. R. *et al.* Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nat. Commun.* **7**:11609 doi: 10.1038/ncomms11609 (2016).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>