



Published in final edited form as:

Nat Neurosci. ; 15(7): 1040–1046. doi:10.1038/nn.3130.

Rational regulation of learning dynamics by pupil–linked arousal systems

Matthew R. Nassar¹, Katherine M. Rumsey¹, Robert C. Wilson², Kinjan Parikh¹, Benjamin Heasley¹, and Joshua I. Gold¹

¹Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104

²Department of Psychology, Princeton University, Princeton, New Jersey 08540

Abstract

The ability to make inferences about the current state of a dynamic process requires ongoing assessments of the stability and reliability of data generated by that process. We found that these assessments, as defined by a normative model, were reflected in non–luminance–mediated changes in pupil diameter of human subjects performing a predictive–inference task. Brief changes in pupil diameter reflected assessed instabilities in a process that generated noisy data. Baseline pupil diameter reflected the reliability with which recent data indicated the current state of the data–generating process and individual differences in expectations about the rate of instabilities. Together these pupil metrics predicted the influence of new data on subsequent inferences. Moreover, a task– and luminance–independent manipulation of pupil diameter predictably altered the influence of new data. Thus, pupil–linked arousal systems can help regulate the influence of incoming data on existing beliefs in a dynamic environment.

Introduction

Many decisions, from foraging to financial, depend on the ability to infer a state of the world from both historical and newly arriving information. Such inferences are particularly challenging when they must account for multiple sources of uncertainty. When the uncertainty results from noise, reflecting random fluctuations in the information generated by an otherwise stable state, the average over all historical information is most predictive of future observations. In contrast, when the uncertainty results from a change in the state itself, only the most recent information pertains to the new state. Thus, historical information should be discounted and beliefs should be updated rapidly to maximize their predictive power. Under certain conditions, human subjects appear to encode and respond appropriately to these different forms of uncertainty when making inferences in a dynamic environment^{1–3}. Here we examined whether this ability is governed, at least in part, by

Users may view, print, copy, download and text and data- mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: http://www.nature.com/authors/editorial_policies/license.html#terms

Correspondence to: Joshua I. Gold.

Author contributions: M.N., J.G., and B.H. designed the experiment and tasks. M.N., K.R., and K.P. collected and analyzed data. M.N. and R.W. developed and applied the reduced Bayesian model. M.N. and J.G. wrote the manuscript.

arousal systems that affect pupil diameter, which are thought to include the noradrenergic brainstem nucleus locus coeruleus⁴⁻⁷.

Non-luminance-mediated changes in pupil diameter have long been used as indicators of clinical, cognitive, and arousal states⁸⁻¹¹. One interpretation of these pupil changes is that they reflect the amount of cognitive effort exerted at a given time, which can be related to task uncertainty¹¹. Accordingly, changes in pupil diameter can be elicited via manipulations of the uncertainty associated with possible actions in certain choice tasks^{6,12}. Changes in pupil diameter can also reflect perceived changes in the world, including perceptual switches during perceptual rivalry, detection of targets in oddball or near-threshold tasks, responses to low-probability go signals in a go/no-go task, and perceived changes in task utility that can affect task engagement^{7,12-15}.

These kinds of uncertainty- and change-related signals are thought to contribute to rational inference in a dynamic environment, including helping to regulate the relative influence of historical and newly arriving information on existing beliefs^{2,3}. Such regulation is a key feature of cognitive flexibility and can be equivalent to adjusting the learning rate in a reinforcement-learning framework^{1,16}. Our goal was to determine how such learning-rate adjustments relate to pupil-linked arousal systems. We show that the arousal system and possibly the locus coeruleus can play important and computationally complex roles in rationally regulating the influence of incoming information on beliefs about a dynamic world.

Results

We measured pupil diameter in thirty human subjects while they performed an isoluminant version of a predictive-inference task². Below we describe task performance, summarize a nearly optimal model that captures key features of performance, demonstrate that certain aspects of pupil diameter encode key variables in the model that can be used to predict performance, and finally show that a task-independent manipulation of arousal and pupil diameter can lead to predictable changes in task performance.

Behavior

The predictive-inference task required subjects to minimize errors in predicting the next number (outcome) in a series. The outcomes were picked from a Gaussian distribution with a mean that changed at random intervals (change points) and a standard deviation (set to either 5 or 10) that was stable over each block of 200 trials (Fig 1). After each prediction was recorded, the new outcome was shown using an iso-luminant display for 2 s, during which time the subject maintained fixation and pupil diameter was measured (Fig. 1). After this interval, the outcome disappeared and the previous prediction reappeared, to be updated for the subsequent trial. Payment scaled inversely with the subject's mean absolute error during the session².

We quantified the extent to which each new outcome influenced the subsequent prediction as the learning rate in a simple delta-rule model (Eq. 3)². The learning rate was equal to the magnitude of change in the prediction expressed as a fraction of the error made on the

previous prediction. Thus, a learning rate of one indicated abandonment of the previous prediction in favor of the most recent outcome. A learning rate of zero indicated maintenance of the previous prediction despite a non-zero prediction error.

Subjects tended to use variable learning rates that spanned the entire allowed range, from zero to one. Within this range, learning rates tended to be higher for larger errors, scaled by the noise of the generative distribution (Fig. 2A). Learning rates also tended to be highest on the trial after a change point and then decay for several trials thereafter (Fig. 2B). These basic trends were similar across subjects, although individual subjects used dramatically different distributions of learning rates (Fig. 2C).

Reduced Bayesian model

The learning rates used by subjects were consistent with both a full and a simplified version of the optimal (Bayesian) model^{2,17–19} One advantage of the reduced Bayesian model is that it updates beliefs according to a delta rule in which the learning rate is computed according to only two parameters computed per trial: change-point probability and relative uncertainty (Fig. 3A).

Change-point probability approximates the posterior probability that the mean of the generative distribution changed since the previous trial, given all previous data. If the mean did change, then previous outcomes should be unrelated to future ones and not contribute to an updated prediction. Accordingly, the model uses learning rates that scale linearly towards one (thus discarding historical information) as change-point probability approaches one (Fig. 3A). Change-point probability is computed by comparing the probability of each new outcome given either the current predictive distribution or the occurrence of a change point (Eq. 5). Its value increases monotonically as a function of the absolute difference between predicted and actual outcome, scaled according to the standard deviation of the generative distribution (Eq. 6, Fig. 3B).

Relative uncertainty is a function of total uncertainty, which in our task arises from two sources. The first source, noise, reflects the unreliability with which a single sample can be predicted from a distribution with a known mean. The second source reflects the unreliability of the current estimate of the mean, which decreases as more data are observed from a distribution. Relative uncertainty is the magnitude of this second form of uncertainty as a fraction of total uncertainty, analogous to the gain in a Kalman filter. Relative uncertainty determines the learning rate when change-point probability is zero and sets the y-intercept of the relationship between change-point probability and learning rate otherwise (Fig. 3A). The effects of relative uncertainty on model learning rates are greatest on the trials following a change point, when its value peaks at 0.5 and then decays over several trials (Eq. 7; Fig 3C).

Like the human subjects, the model tended to compute learning rates that were highest just following a change point in the mean of the generative distribution and then decayed for several trials independently of noise. When applied to the exact same outcome sequences as the subjects, the model also tended to produce similar learning rates (Fig. 3D).

We related change–point probability and relative uncertainty computed in the model to the mean pupil diameter (“pupil average”) and change in pupil diameter (“pupil change”) measured during the 2–s outcome–viewing period (Fig. 1 inset), using two linear regression models. The first, simpler model had four parameters: change–point probability and relative uncertainty computed from the reduced Bayesian model, the standard deviation of generative distribution, and a binary variable describing whether or not the prediction error was exactly zero. The second model included all of these parameters, as well as several potential confounding factors such as eye position and velocity (see Methods). The models are complementary: the first avoids potential interactions between large numbers of parameters and thus has coefficients that are more readily interpretable, whereas the second avoids missing out on the many factors that in principle could affect our pupil measurements. Both models captured a significant amount of variability in the pupil data (For pupil average/pupil change data, an F –test rejected the null model relative to the small model for 27/15 of the 30 subjects, and a nested F –test rejected the small model relative to the large model for 29/19 of the 30 subjects, $p < 0.05$).

Below we first report the most prominent effects from these regression analyses, which were similar for the two models and include roughly monotonic relationships between pupil change and change–point probability and between pupil average and relative uncertainty. We later show that these relationships were in fact slightly more complicated and included a dependence on baseline pupil diameter that helps us to interpret the results in terms of known properties of the arousal system.

Pupil change reflected change–point probability

The change in pupil diameter during the outcome–viewing period, like change–point probability in our model, tended to increase as a function of error magnitude, scaled as a function of noise (Fig. 4A; compare to Figs. 3B). Accordingly, when computed by the model using the same sequence of outcomes experienced by each subject, change–point probability tended to be positively predictive of z –scored pupil change (Fig. 4B ordinate). The complement was also true: change–point probability varied systematically as a function of pupil change for data pooled across the population (Fig. 4C). In contrast, there was no consistent relationship between change–point probability and pupil average (Fig. 4B abscissa).

One notable exception to the positive relationship between pupil change and error magnitude occurred for trials in which the error was exactly zero, which corresponded to relatively large pupil changes (left–most data in Fig. 4A). Accordingly, a binary variable added to the linear model that described whether or not the subject correctly predicted the outcome was related to pupil change (the mean value of the regression coefficient was 0.180 z_{PC} for the four–parameter regression model and 0.156 z_{PC} for the larger model; $p < 0.05$ for H_0 : mean=0 for each model) but not pupil average (mean regression coefficient=–0.076 and –0.092 z_{PA} for the smaller and larger regression models, respectively, $p > 0.05$). Thus, pupil change reflected not only change–point probability, but also whether or not the subject correctly predicted the observed outcome.

Average pupil diameter reflected belief uncertainty

The average pupil diameter during the outcome-viewing period, like relative uncertainty in our model, tended to peak on the trial after a change point and then diminish in magnitude as more relevant information reinforced the existing belief (Fig 5A; compare to Figs. 2B and 3C). Accordingly, when computed by the model using the same sequence of outcomes experienced by each subject, relative uncertainty tended to be positively predictive of pupil average (Fig. 5B abscissa). This result did not simply reflect differences in motor output following change points (e.g., longer button presses to choose a learning rate near one), because similar results were obtained in a control experiment in which subject predictions were reset using a learning rate of 0.5 on each trial, thus requiring the same motor act to choose a learning rate of either zero or one (mean regression coefficient=0.30 and 0.35 $z_{PA/RU}$ for the smaller and larger regression models, respectively, $p<0.05$). The complement was also true: relative uncertainty varied systematically as a function of pupil average for data pooled across the population (Fig. 5C). In contrast, there was no consistent relationship between relative uncertainty and pupil change (Fig. 5B ordinate).

Overall uncertainty in our task depends on not only relative uncertainty but also noise, which we manipulated by varying the standard deviation of the generative distribution in blocks (STD=5 or 10). Consistent with our model, in which noise is only used to compute change-point probability (Eqs. 5 and 6), these manipulations of noise were reflected in pupil change but only insofar as pupil change represented change-point probability (Fig. 4A). These manipulations of noise did not have any other systematic effects on either pupil change or pupil average ($p>0.1$ for H_0 : a mean value of zero for the regression coefficient describing the influence of noise on the given pupil measurement for both regression models). Thus, for this task pupil average did not appear to reflect overall uncertainty about a future outcome but rather a specific form of uncertainty that arises after change points and signals the need for rapid learning.

Pupil metrics reflected individual learning differences

As noted above (Fig. 2C), there was a great deal of variability in the average learning rates used by individual subjects. These individual differences are thought to reflect biases that govern the extent to which subjects tend to interpret the cause of prediction errors in terms of either noise or change points². One advantage of our reduced model is that it can simulate these individual differences in terms of the subjective hazard rate, which is the expected rate at which change points will occur. Accordingly, fitting the model to behavioral data from individual subjects with subjective hazard rate as a single free parameter yielded fit values that varied systematically with average learning rates ($r=0.93$, H_0 : $r=0$, $p<0.001$; Fig 6A).

These individual differences in the inferred (fit) subjective hazard rates corresponded to individual differences in both the temporal dynamics and magnitude of outcome-locked pupil responses. We quantified the temporal dynamics using an index that related the pupil response on a given trial to a mean-subtracted version of the template shown in Fig. 6B. This template describes the strength of the across-subject, linear relationship between pupil diameter and hazard rate in a sliding time window. This relationship was strongest soon after outcome onset, thus likely reflecting prior expectations about the newly arriving outcome.

There was a positive relationship between the mean value of this index and fit hazard rate for individual subjects ($r=0.51, p<0.01$). In addition, there was a positive relationship between pupil average and fit hazard rate for individual subjects ($r=0.40, p<0.05$).

Based on these relationships, we constructed a linear regression model using the temporal–dynamics index and pupil average to explain individual differences in task performance. The model yielded strong, pupil–based predictions of per–subject values of both fit hazard rate ($r=0.59, p<0.001$) and average learning rate ($r=0.59, p<0.001$; Fig 6C). Thus, individual differences in average learning rate, which can be described computationally as differing expectations about the rate of change–points, could be predicted from the temporal dynamics and average magnitude of pupil diameter measured during outcome viewing.

Pupil metrics predicted trial–by–trial learning rates

The relationships between pupil metrics and parameters of the reduced Bayesian model suggest that measurements of pupil diameter during the outcome–viewing period can be used to predict the subsequent learning rate. For example, we found positive relationships between pupil change and change–point probability (Fig. 4) and between pupil average and relative uncertainty (Fig. 5). Thus, observing relatively high values of either pupil metric on a given trial should indicate that the subject will use a larger–than–average learning rate when adjusting beliefs according to the outcome observed on that trial. We tested this idea directly, as follows.

First, we examined the relationship between pupil change, pupil average, and learning rate for individual subjects. We used a regression model to describe learning rate (z -scored per subject) in terms of pupil change and pupil average. On average, this linear regression computed per subject yielded a positive coefficient for pupil change (mean=0.108 z_{LR}/z_{PC} , $p<0.05$ for H_0 : mean=0) and a smaller, not statistically significant, positive coefficient for pupil average (mean=0.085 z_{LR}/z_{PA} , $p=0.13$; Fig. 7A).

Second, we used a simple, weighted sum of pupil change and pupil average to assess their combined predictive power across subjects. Using weights equal to the mean value of the per–subject regression coefficients from the previous analysis (Fig. 7A), the weighted sum was moderately predictive of learning rate across all subjects ($r=0.067, p<0.001$). However, this analysis did not take into account a systematic, negative dependence of the sum of these per–subject coefficients (which is related to the overall ability of the weighted sum to account for learning rate) on subjective hazard rate predicted by pupil dynamics (Fig. 7B). Subjects with low pupil–predicted hazard rates had pupil responses that were good predictors of learning rate. Subjects with increasingly high pupil–predicted hazard rates had pupil responses that were increasingly less predictive, and in some cases negatively predictive, of learning rate.

Third, we used a more complicated linear model that also included across–subject differences in pupil dynamics that related to subjective hazard rates, which markedly improved our overall ability to use pupil metrics to predict learning rates. This model had three terms: 1) the sum of pupil change and pupil average computed per trial, weighted according to average regression coefficients in Fig. 7A; 2) the pupil–predicted hazard rate,

computed per subject (see Fig 6C); and 3) the multiplicative interaction between these two variables. Using this model, pupil measurements could effectively predict learning rates for all data from all subjects ($r=0.38$, $p<0.001$). These predictions accounted for variations in learning rates both across (Fig. 6B) and within (Fig. 7C) subjects.

Task-independent pupil manipulation altered behavior

To examine whether the correlations between pupil measures and learning behavior might reflect an underlying causal process, we used an arousal manipulation that affected pupil diameter and measured its effects on learning behavior. In particular, we occasionally and without warning switched the auditory cue that preceded fixation. Subjects were told that these auditory–cue switches were unrelated to the task and they therefore should ignore the specific sounds. Nevertheless, this manipulation led to increases in both pupil average and pupil change on trials in which the fixation cue was switched (Fig 8A; t -test for H_0 : mean effect size=0, $p<0.001$ for both pupil average and pupil change). Thus, we caused consistent changes in the pupil measures that were correlated with the computational variables needed to solve the task.

This manipulation caused systematic changes in task performance that depended on baseline pupil diameter (Fig. 8B). For trials with relatively small baseline diameter (i.e., less than its per–subject median value), individual subjects tended to use larger learning rates on auditory–switch trials than otherwise (Fig 8B abscissa; mean across subjects=0.113, t -test for H_0 : mean=0, $p<0.01$). For trials with relatively large baseline diameter, subjects used slightly smaller learning rates on auditory–switch trials than otherwise, although this trend was not statistically significant (Fig 8B ordinate; mean=-0.037, $p=0.35$). The average difference in the size of these effects from small– versus large–diameter trials was >0 , implying that the effects of this manipulation depended on baseline pupil diameter (Fig 8B diagonal; paired t -test, $p<0.001$). These effects did not result from systematic differences in task conditions for switch versus non–switch trials, because the same three analyses yielded no effects when applied to learning rates computed by our reduced Bayesian model ($p>0.5$).

This dependence on baseline pupil diameter is suggestive of the Yerkes–Dodson “inverted U” relationship between arousal and learning. According to that idea, learning is highest for moderate levels of arousal and lowest for either overly high or overly low levels of arousal²⁰. Our subjects appeared to be consistently engaged during task performance, implying that we were probably not sampling overly low or high arousal states. Nevertheless, in a narrower range and assuming a correspondence between arousal state and baseline pupil diameter, we found that the relationships between learning behavior and our arousal manipulation were qualitatively consistent with an “inverted U.” In particular, auditory–switch trials tended to correspond to the largest increases in learning rate when baseline pupil diameter was relatively low (steepest ascent in the “inverted U”) and the largest decreases in learning rate when baseline pupil diameter was relatively high (steepest descent in the “inverted U”; Fig. 8C, open circles).

This “inverted U” relationship was also apparent in our previous pupil measurements, in two ways. First, across subjects, those with larger average pupil diameters during outcome viewing tended to use learning rates that were less, or even negatively, predicted by

fluctuations in pupil metrics relative to other subjects (Fig 7B). Second, subjects that had lower pupil–predicted hazard rates used learning rates that were positively correlated with pupil metrics when their baseline pupil diameter was low but negatively correlated when their baseline pupil diameter was high (Fig 8C, filled circles). Thus, results from both our pupil-manipulation and pupil-measurement experiments were consistent with an important role for the arousal system in the rational regulation of learning.

Discussion

We examined the relationship between pupil diameter, which is related to arousal and autonomic state, and learning rate, which describes the extent to which new information is used to adjust existing cognitive beliefs. Consistent with previous work^{2,22,23}, we found that human subjects performing a predictive–inference task were most heavily influenced by outcomes that occurred shortly after a change point in the outcome–generating process. One possible mechanism for this effect is a dynamic regulation of the relative influence of incoming information on cortical processing³. Insights into the computations required for such a regulator are provided by a reduced model that approximates the ideal observer for the task, describes subject behavior, and bases learning rates on two parameters that we found to be represented in pupil measurements: change–point probability and relative uncertainty.

In our model, change–point probability depends on the absolute value of the most recent prediction error and drives increased learning after surprisingly large errors. We found that change–point probability was positively correlated with changes in pupil diameter. This relationship is consistent with early pupillometry studies that showed an inverse relationship between stimulus–evoked pupil responses and stimulus probability, as well as more recent work interpreting outcome–locked pupil responses in terms of the surprise associated with errors in judging uncertainty, called the risk prediction error^{23–25}. We also found that pupil change was not always directly related to change–point probability, with particularly large pupil changes on trials with exactly zero error that might have been surprisingly rewarding and/or reflected an association with an atypical consequence (i.e., no possibility of updating the next prediction).

Relative uncertainty, the second parameter in our model, represents uncertainty about the true underlying mean and drives learning from outcomes that occur after a change point. We found that relative uncertainty was correlated with average pupil diameter. We also found that changes in another form of uncertainty that should not drive learning (i.e., changes in the standard deviation of the generative process in our task) did not lead to similar effects on pupil diameter. These results are complementary to a recent finding that pupil diameter tends to increase during exploratory decisions that occur during periods of uncertainty about the best available option⁶. These findings suggest that pupil–linked arousal systems encode an uncertainty signal that facilitates both learning and information–seeking behaviors.

We also found strong individual differences in task behavior that could be captured by fitting a prior expectation about the rate of change points (hazard rate) to behavioral data. We found that subjects who were fit by higher hazard–rate models tended to have larger

pupil dilations during the outcome-viewing period. This physiological difference arose early in the viewing period, consistent with the idea that these individual differences reflected a prior expectation about the source of the upcoming error.

We used these relationships between pupil metrics and change-point probability, relative uncertainty, and the hazard-rate prior to predict the extent to which subjects were influenced by each new outcome. We also manipulated pupil diameter using a task-irrelevant auditory manipulation that resulted in changes in task performance that were consistent with our measured relationships between pupil metrics and key task variables. These results provide new insights into the specific computations that are reflected in pupil diameter and establish their causal role in belief updating.

These computations likely involve, at least in part, neural activity in the locus coeruleus. One intriguing possibility is that the two key variables from our model are encoded by two distinct modes of locus coeruleus activation⁵: change-point probability, reflected in pupil change, is encoded by phasic activation of the locus coeruleus, whereas relative uncertainty, reflected in pupil average, is encoded by tonic activation of the locus coeruleus. Although direct confirmation is still needed, this idea is supported by several lines of evidence, including: 1) a compelling example of simultaneous measurements of locus coeruleus activity and pupil diameter in a monkey that are closely correlated⁵; 2) similar modulations of pupil diameter and locus coeruleus activity under certain task conditions, such as changes in utility in that affect behavioral engagement^{6,7}; and 3) a proposed anatomical substrate involving common activation from the nucleus paragigantocellularis, which contributes to both locus coeruleus and sympathetic nervous system function^{4,26}. The consequence of locus coeruleus involvement would be the task-related release of norepinephrine throughout the nervous system. Consistent with our results, norepinephrine release is thought to permit or facilitate changes in behavior that follow unexpected changes in the environment and learning in general, possibly by modulating experience-dependent neural plasticity^{3,27–32}.

More generally, our results are consistent with the idea that brain areas that regulate the influence of newly arriving information on existing beliefs are also strongly linked to arousal and autonomic function^{1,3,6,7,25,33,34}. These areas likely include not just the locus coeruleus but also the anterior cingulate cortex (ACC), which has strong reciprocal connections with the locus coeruleus and whose activity encodes several signals closely related to change-point probability, including unsigned prediction errors and learning rates^{1,5,21,35}. This arousal system appears to govern not simply overall alertness or other non-specific factors that might affect overall task performance, but rather a computationally sophisticated process that rationally regulates the influence of new sensory information in a dynamic environment. These computations take into account both ongoing processing of task-relevant variables like change-point probability and relative uncertainty and state variables including prior expectations about the rate of change. These factors are combined in a manner that is consistent with the Yerkes-Dodson “inverted U” relationship between arousal level and learning rate (Fig. 8C)²⁰.

In summary, our work suggests a relationship between arousal state and learning rate that is likely a result of a coordinated learning-arousal network including the locus coeruleus and

ACC. The representation of normative learning variables in this network suggests that subtle changes in arousal might reflect rational regulation of the influence of new information on ongoing inferences about a dynamic world.

Methods

Predictive–inference task

Human subject protocols were approved by the University of Pennsylvania Internal Review Board. Thirty subjects (19 female, 11 male; age range = 19–29 years) participated in the primary study and an independent sample of 29 subjects (17 female, 12 male; age range = 19–25 years) participated in the arousal manipulation study after providing informed consent. Both studies used a predictive–inference task that required subjects to predict each subsequent number to be presented in a series². For each trial t , a single integer (X_t) was presented that was a rounded pick sampled independently and identically from a Gaussian distribution whose mean (μ_t) changed at unsigned change points and whose standard deviation (σ_t) was fixed to either 5 or 10 within each block of 200 trials. Change points occurred with a probability of zero for the first three trials following a change point and 0.1 for all trials thereafter.

To facilitate measurements of non–luminance–mediated effects on pupil diameter, we used a different visual display and task timing than in our previous study². Subjects were shown a numeric representation of their current prediction at a central location on a CRT monitor. Background screen pixels were a checkerboard of light and dark pixels (mean \pm STD luminance in a circle with radius 6.5 cm = 0.457 ± 0.010 cd/m²). Numbers were drawn in an intermediate gray color (0.445 ± 0.005 cd/m²). When viewed passively by a control group of four subjects outside of the context of the predictive–inference task, no individual stimulus (number) had a significant effect on average pupil diameter or evoked changes in pupil diameter (t -test for H_0 : equal means between each stimulus and all others, $p>0.3$ for all stimuli after correcting for multiple comparisons), nor did the number of digits contained within the stimulus affect either pupil variable ($p>0.4$).

For each trial, the subject indicated his or her updated prediction using a video gamepad. Each prediction was constrained to be between the previous prediction and the most recent outcome, thus limiting learning rates to between zero and one. After the new prediction was chosen, the numeric representation of this prediction disappeared, an auditory cue was played, and a numeric representation of the new outcome was shown. Subjects were instructed to fixate centrally for 2 s at this point; failure to do so (within a square window, 9° per side) resulted in a tone indicating a fixation error. After 2 s the new outcome disappeared, the prediction re–appeared, and an auditory cue was played to indicate that the prediction should be updated. Fourteen subjects also participated in a control version of the task in which the prediction was reset after viewing the new outcome to reflect an update equivalent to a learning rate of 0.5. For this task, the same motor output (in terms of number or duration of button presses) was required to use a learning rate of either zero or one on each trial.

Subjects were told that the numbers were generated from a noisy process and that several discreet change points would occur over the course of the task. They were instructed to make a prediction on each trial (B_t) such that the average error made on all predictions, $\langle |B_t - X_t| \rangle$, would be minimized. Payout depended on how well they achieved this goal, as described previously².

The pupil-manipulation task was identical to primary version of the task, except that the auditory cue played at the beginning of fixation was occasionally switched to another sound from a library of 31 sound effects downloaded from an online library. Sounds were 0.09–1.4 s in duration (mean \pm STD = 0.72 \pm 0.42 s) and played at 56–70 dB (A-weighted; mean \pm STD = 62.5 \pm 3.9 dB). Switch trials occurred at random, with a probability of 0.1 on the 9 trials following a switch, 0.8 thereafter. On switch trials, the given sound was played, on average, 7 dB louder than otherwise. Seven of 29 subjects completing the pupil-manipulation task were excluded from further analyses because of an excessive number of fixation errors (blinks or lost fixation on >40 percent of trials).

Pupil-diameter measurements

Pupil diameter was sampled at 120 Hz and recorded throughout the task using an infrared video eye-tracker (ASL, Inc.). Blinks were identified using a custom blink filter based on pupil diameter and vertical and horizontal eye position, then removed by linear interpolation of values measured just before and after each identified blink. Blink-filtered diameter was low-pass filtered using a Butterworth filter with a cutoff frequency of 3.75 Hz. These filtered measurements were then z -scored within each session.

All analyses excluded trials in which blinks or fixation errors during outcome viewing were detected online (these events were followed by a beep to remind the subject to minimize their occurrence). The first 20 trials from each block were also excluded to avoid possible changes in average luminance at block boundaries. Pupil average was computed for each trial by taking the mean of all 240 z -scored pupil measurements from the 2 s-long outcome-viewing period of the trial. Pupil change was computed for each trial by subtracting the average pupil measurement from early in the outcome-viewing period (0–1 s after outcome presentation) from the average pupil measurement from late in the outcome-viewing period (1–2 s after outcome presentation). Trials that included blinks that were detected offline (but not online) were used to compute pupil average by interpolating values from just before and just after the blink. These trials were not used to compute pupil change, which was much more sensitive to the timing of blinks.

Reduced Bayesian model

Optimal performance on the predictive-inference requires inferring the probability distribution over possible outcomes on the next timestep, given all previous data and the process by which those data were generated: $p(X_{t+1}|X_{1:t})$. Because the relationship between the data on the next timestep is independent of all previous data conditioned on the mean of the current distribution (μ), the solution can be formulated in terms of μ :

$$p(X_{t+1}|X_{1:t}) = \sum_{\mu_t} p(X_{t+1}|\mu_t)p(\mu_t|X_{1:t}) \quad [1]$$

and the probability distribution over possible means given previous data can be inverted according to Bayes' rule:

$$p(\mu_t|X_{1:t}) = \frac{p(X_{1:t}|\mu_t)p(\mu_t)}{p(X_{1:t})} \quad [2]$$

Although computationally tractable solutions to this problem exist, these solutions specify learning rates that are complicated functions of either the probability distribution over all possible means¹ or over all possible "runs" of non-change-point trials¹⁹. To simplify the algorithm, the reduced model computes the posterior probability distribution over possible means as described above but maintains only the first two moments of this distribution. This assumption massively reduces the number of required computations but has minimal effects on performance². An added advantage of this model is that it can be formulated as a delta rule:

$$\begin{aligned} B_{t+1} &= B_t + \alpha_t \times \delta_t \\ \delta_t &= x_t - B_t \end{aligned} \quad [3]$$

where B is the belief about the mean of the underlying distribution; α is the learning rate; and δ is the prediction error, which is the difference between the actual and predicted outcome. The learning rate depends on two variables that are updated on each trial:

$$\alpha_t = \tau_t + (1 - \tau_t)\Omega_t \quad [4]$$

where change-point probability (Ω) reflects the probability that μ_t is not equal to μ_{t-1} , and relative uncertainty (τ) reflects the variance on the predictive distribution in μ (i.e., uncertainty about the location of the mean) divided by the variance on the predictive distribution in X (i.e., total uncertainty about the location of the next outcome).

Performance of the reduced Bayesian model also depends on an expectation about the prior probability on change points, or the hazard rate. Specifically, hazard rate directly influences the computation of change-point probability on each trial:

$$\Omega_t = \frac{U(X_t|0, 300)H}{U(X_t|0, 300)H + \mathcal{N}(X_t|B_t, \sigma_t^2)(1 - H)} \quad [5]$$

Where U and N represent uniform and normal distributions, respectively; H is the hazard rate; B_t is the model's prediction on trial t ; and σ^2 is the total variance on the predictive distribution, which is discussed below. We incorporated hazard rate into the model in two ways: 1) using the true generative hazard rate for trials in which a change point did not recently occur (0.1) or 2) by fitting the model to behavior by minimizing the total squared difference between subject and model predictions using a constrained search algorithm (fmincon in MATLAB) with hazard rate as a free parameter.

The total variance on the predictive distribution in the model comes from two sources:

$$\sigma_t^2 = N^2 + \frac{\tau_t N^2}{1 - \tau_t} \quad [6]$$

The first source is the standard deviation on the outcome-generating distribution (N). The second source is uncertainty about the mean of that distribution and depends on both N and relative uncertainty (τ). Here we set N to be the actual experimental standard deviation, but we update τ after each outcome according to the variance on the predictive distribution over possible means:

$$\tau_{t+1} = \frac{N^2 \Omega_t + (1 - \Omega_t)(\tau_t N^2) + \Omega_t(1 - \Omega_t)(\tau_t + B_t(1 - \tau_t) - X_t)}{N^2 \Omega_t + (1 - \Omega_t)(\tau_t N^2) + \Omega_t(1 - \Omega_t)(\tau_t + B_t(1 - \tau_t) - X_t) + N^2} \quad [7]$$

such that if a change point occurs, relative uncertainty is reset to 0.5 (first term in numerator); if a change point does not occur, relative uncertainty is reduced (second term in numerator); and if the model is uncertain about whether a change point occurred, relative uncertainty is increased to reflect this uncertainty (third term in numerator).

Statistical analyses

Trial-by-trial values of pupil average and pupil change were each z-scored for the full session (z_{PA} and z_{PC} , respectively) and then fit with a linear regression model using four parameters: 1) change-point probability, computed by the reduced Bayesian model for each trial; 2) relative uncertainty, computed by the reduced Bayesian model for each trial; 3) noise, the standard deviation of the outcome-generating distribution; and 4) a binary vector specifying whether or not the subject correctly predicted the outcome on that trial. We also used a larger model that, in addition to the above four parameters, included: the average horizontal and vertical eye position and the change in horizontal and vertical eye position measured during the outcome-viewing period; the subject's prediction and the computer-generated outcome from the current trial; the pupil change measured on the previous trial; and the trial number within the block and within the session.

Pupil-predicted hazard rates were derived from pupil measurements and the reduced Bayesian model as follows. First, we inferred the subjective hazard rate used by each subject by fitting his or her behavioral data to the reduced Bayesian model with hazard rate (H) as the only free parameter. Next, we fit a linear regression model explaining H in terms of pupil measurements. That model had two terms, computed per subject: 1) the mean value of pupil average, and 2) an index of pupil dynamics. The index was computed as the mean value of the dot product of trial-by-trial pupil measurements and the mean-subtracted curve shown in Fig. 6B. Finally, we used the coefficients from a linear fit that excluded the data from an individual subject to combine the mean pupil average and pupil-dynamics index (from the excluded subject) into a pupil-predicted hazard rate for that subject.

Pupil-predicted learning rates were computed according to the relationships between pupil metrics and model parameters. Linear fits to the relationship between pupil average and relative uncertainty were computed for each subject, and these fits were used estimate

relative uncertainty for each trial-by-trial measurement of pupil diameter. Linear fits to the relationship between pupil change and change-point probability were computed for each subject, and these fits were used to estimate change-point probability for each trial-by-trial measurement of pupil change. To compute predicted learning rates, the two predicted model quantities were combined according to Eq. 4. We also used a more complex linear model that took into account pupil-predicted hazard rates; see text for details.

Arousal-induced learning effects for the inverted-U analyses were computed separately for sound-manipulation and non-manipulation sessions. For sound-manipulation sessions, learning rates were fit to a cumulative Weibull as a function of error magnitude for each subject and noise condition, to account for the relationship shown in Fig. 4A. Residuals from this fit, which reflected error-independent variability in learning rate, were z-scored per subject. Initial pupil diameter, as measured by the average diameter during the first 100 ms of the outcome phase, was also z-scored per subject. Data were binned across subjects according to the initial diameter z-score. The effect of the sound manipulation was computed as a signed d' describing the difference in the z-scored residual learning rates used on auditory shift versus non-auditory shift trials. For non-manipulation sessions, the relationship between pupil metrics and learning rate was characterized only for subjects with low pupil-predicted hazard rates (<0.6). Subjects with high pupil-predicted hazard rates tended to have small or negative relationships between pupil metrics and learning rate and thus were omitted from this analysis. Arousal effect size was computed as the correlation coefficient between the weighted sum of pupil metrics and learning rate, each z-scored per subject (positive/negative values indicate that learning rates tended to increase/decrease as pupil effects increased) for equally sized bins of baseline pupil diameter (z-scored per subject).

Acknowledgments

We thank Jon Cohen, Sascha du lac, Long Ding, Yin Li, and Joy Nassar for helpful comments. Supported by EY015260, the McKnight Endowment Fund for Neuroscience, the Burroughs-Wellcome Fund, the Sloan Foundation, and NIH Training Grant in Computational Neuroscience T90 DA22763.

References

1. Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat. Neurosci.* 2007; 10:1214–1221. [PubMed: 17676057]
2. Nassar MR, Wilson RC, Heasly B, Gold JI. An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* 2010; 30:12366–12378. [PubMed: 20844132]
3. Yu AJ, Dayan P. Uncertainty, neuromodulation, and attention. *Neuron.* 2005; 46:681–692. [PubMed: 15944135]
4. Nieuwenhuis S, De Geus EJ, Aston-Jones G. The anatomical and functional relationship between the P3 and autonomic components of the orienting response. *Psychophysiology.* 2010
5. Aston-Jones G, Cohen JD. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* 2005; 28:403–450. [PubMed: 16022602]
6. Jepma M, Nieuwenhuis S. Pupil diameter predicts changes in the exploration-exploitation trade-off: evidence for the adaptive gain theory. *J. Cogn. Neurosci.* 2011; 23:1587–1596. [PubMed: 20666595]

7. Gilzenrat MS, Nieuwenhuis S, Jepma M, Cohen JD. Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cogn. Affect. Behav. Neurosci.* 2010; 10:252–269. [PubMed: 20498349]
8. Krugman HE. Some applications of pupil measurement. *Journal. Of. Marketing. Research.* 1964; 1:15–19.
9. Granholm E, Steinhauer SR. Pupillometric measures of cognitive and emotional processes. *Int. J. Psychophysiol.* 2004; 52:1–6. [PubMed: 15003368]
10. Schmidt, HS.; Fortin, LD. Electronic pupillography in disorders of arousal. In: Fortin, LD.; Schmidt, HS.; Guilleminault, editors. *Sleeping and waking disorders: Indication and technique.* Menlo Park, CA: Addison–Wesley; 1982. p. 127-143.
11. Kahneman D, Beatty J. Pupil diameter and load on memory. *Science.* 1966; 154:1583–1585. [PubMed: 5924930]
12. Richer F, Beatty J. Contrasting effects of response uncertainty on the task–evoked pupillary response and reaction time. *Psychophysiology.* 1987; 24:258–262. [PubMed: 3602280]
13. Hakerem G, Sutton S, Zubin J. Pupillary reactions to light in schizophrenic patients and normals. *Ann. N. Y. Acad. Sci.* 1964; 105:820–831. [PubMed: 14212837]
14. Einhäuser W, Stout J, Koch C, Carter O. Pupil dilation reflects perceptual selection and predicts subsequent stability in perceptual rivalry. *Proc. Natl. Acad. Sci. U. S. A.* 2008; 105:1704–1709. [PubMed: 18250340]
15. Van Olst, EH.; Heemstra, ML.; Ten Kortenaar, T. Stimulus significance and the orienting reaction. In: Kimmel, H.; Olst, EH.; van Orlebeke, JF., editors. *The orienting reflex in humans.* Hillsdale, NJ: Erlbaum; 1979. p. 521-547.
16. Sutton, RS.; Barto, AG. *Reinforcement learning: An introduction.* Cambridge, MA: MIT Press; 1998.
17. Adams, RP.; MacKay, DJC. *Bayesian Online Changepoint Detection.* University. Of. Cambridge. Technical. Report. 2007.
18. Fearnhead P, Liu Z. On–line inference for multiple changepoint problems. *Journal. Of. The. Royal. Statistical. Society.: Series. B. (Statistical. Methodology).* 2007; 69:589–605.
19. Wilson RC, Nassar MR, Gold JJ. Bayesian online learning of the hazard rate in change–point problems. *Neural. Comput.* 2010; 22:2452–2476. [PubMed: 20569174]
20. Yerkes RM, Dodson JD. The relation of strength of stimulus to rapidity of habit–formation. *Journal. Of. Comparative. Neurology. And. Psychology.* 1908; 18:459–482.
21. Krugel LK, Biele G, Mohr PN, Li SC, Heekeren HR. Genetic variation in dopaminergic neuromodulation influences the ability to rapidly and flexibly adapt decisions. *Proc. Natl. Acad. Sci. U. S. A.* 2009; 106:17951–17956. [PubMed: 19822738]
22. Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat. Neurosci.* 2007; 10:1214–1221. [PubMed: 17676057]
23. Raisig S, Welke T, Hagenhof H, van der Meer E. I spy with my little eye: detection of temporal violations in event sequences and the pupillary response. *Int. J. Psychophysiol.* 2010; 76:1–8. [PubMed: 20093148]
24. Friedman D, Hakerem G, Sutton S, Fleiss JL. Effect of stimulus uncertainty on the pupillary dilation response and the vertex evoked potential. *Electroencephalogr. Clin. Neurophysiol.* 1973; 34:475–484. [PubMed: 4121320]
25. Preuschoff K, 't Hart BM, Einhäuser W. Pupil Dilation Signals Surprise: Evidence for Noradrenaline's Role in Decision Making. *Front. Neurosci.* 2011; 5:115. [PubMed: 21994487]
26. Aston–Jones G, Ennis M, Pieribone VA, Nickell WT, Shipley MT. The brain nucleus locus coeruleus: restricted afferent control of a broad efferent network. *Science.* 1986; 234:734–737. [PubMed: 3775363]
27. Sara SJ, Vankov A, Hervé A. Locus coeruleus–evoked responses in behaving rats: a clue to the role of noradrenaline in memory. *Brain. Res. Bull.* 1994; 35:457–465. [PubMed: 7859103]
28. Aston–Jones G, Rajkowski J, Kubiak P. Conditioned responses of monkey locus coeruleus neurons anticipate acquisition of discriminative behavior in a vigilance task. *Neuroscience.* 1997; 80:697–715. [PubMed: 9276487]

29. Tully K, Bolshakov VY. Emotional enhancement of memory: how norepinephrine enables synaptic plasticity. *Mol. Brain*. 2010; 3:15. [PubMed: 20465834]
30. Harley CW. A role for norepinephrine in arousal, emotion and learning?: limbic modulation by norepinephrine and the Kety hypothesis. *Prog. Neuropsychopharmacol. Biol. Psychiatry*. 1987; 11:419–458. [PubMed: 3321150]
31. Corbetta M, Patel G, Shulman GL. The reorienting system of the human brain: from environment to theory of mind. *Neuron*. 2008; 58:306–324. [PubMed: 18466742]
32. Boudet S, Sara SJ. Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends. Neurosci*. 2005; 28:574–582. [PubMed: 16165227]
33. Critchley HD, Mathias CJ, Dolan RJ. Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron*. 2001; 29:537–545. [PubMed: 11239442]
34. Critchley HD. Neural mechanisms of autonomic, affective, and cognitive integration. *J. Comp. Neurol*. 2005; 493:154–166. [PubMed: 16254997]
35. Matsumoto M, Matsumoto K, Abe H, Tanaka K. Medial prefrontal cell activity signaling prediction errors of action values. *Nat. Neurosci*. 2007; 10:647–656. [PubMed: 17450137]

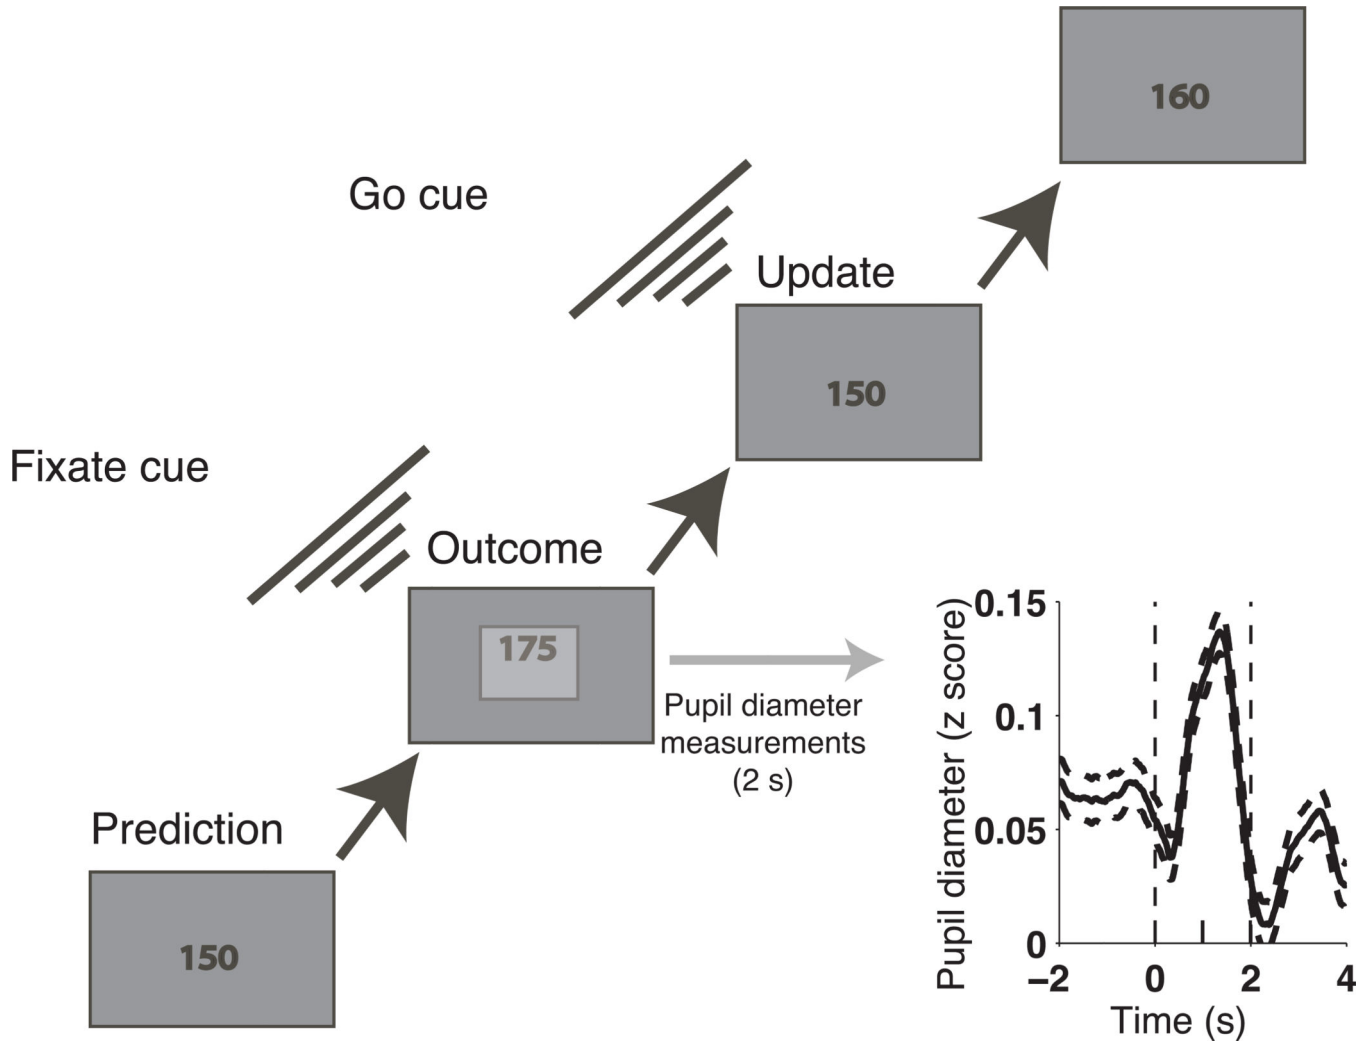


Figure 1. Predictive–inference task sequence and pupillometry. Learning rate was computed by dividing the difference in the prediction from one trial to the next by the difference between the current outcome and the current prediction. Inset: mean±SEM pupil diameter, averaged across z-scores computed per subject, aligned to outcome presentation (time=0). Pupil average was computed for each trial as the mean pupil diameter, z-scored by subject, across the entire 2-s fixation window (vertical dashed lines). Pupil change was computed for each trial as the difference in mean diameter, z-scored by subject, measured late (time=1–2s) versus early (time=0–1s) during fixation.




Figure 2.

Task performance. A, Learning rates were highest after subjects made larger errors, scaled by noise (as indicated). Points and errorbars are mean \pm SEM from all subjects. B, Learning rates were highest on change-point trials and decayed thereafter, similarly for both noise conditions. Points and errorbars are mean \pm SEM from all subjects. C, Learning-rate distributions across all trials from each of the 30 subjects (abscissa), sorted by median learning rate. Horizontal line, box, and whiskers indicate median, 25th/75th percentiles, and 5th/95th percentiles, respectively.

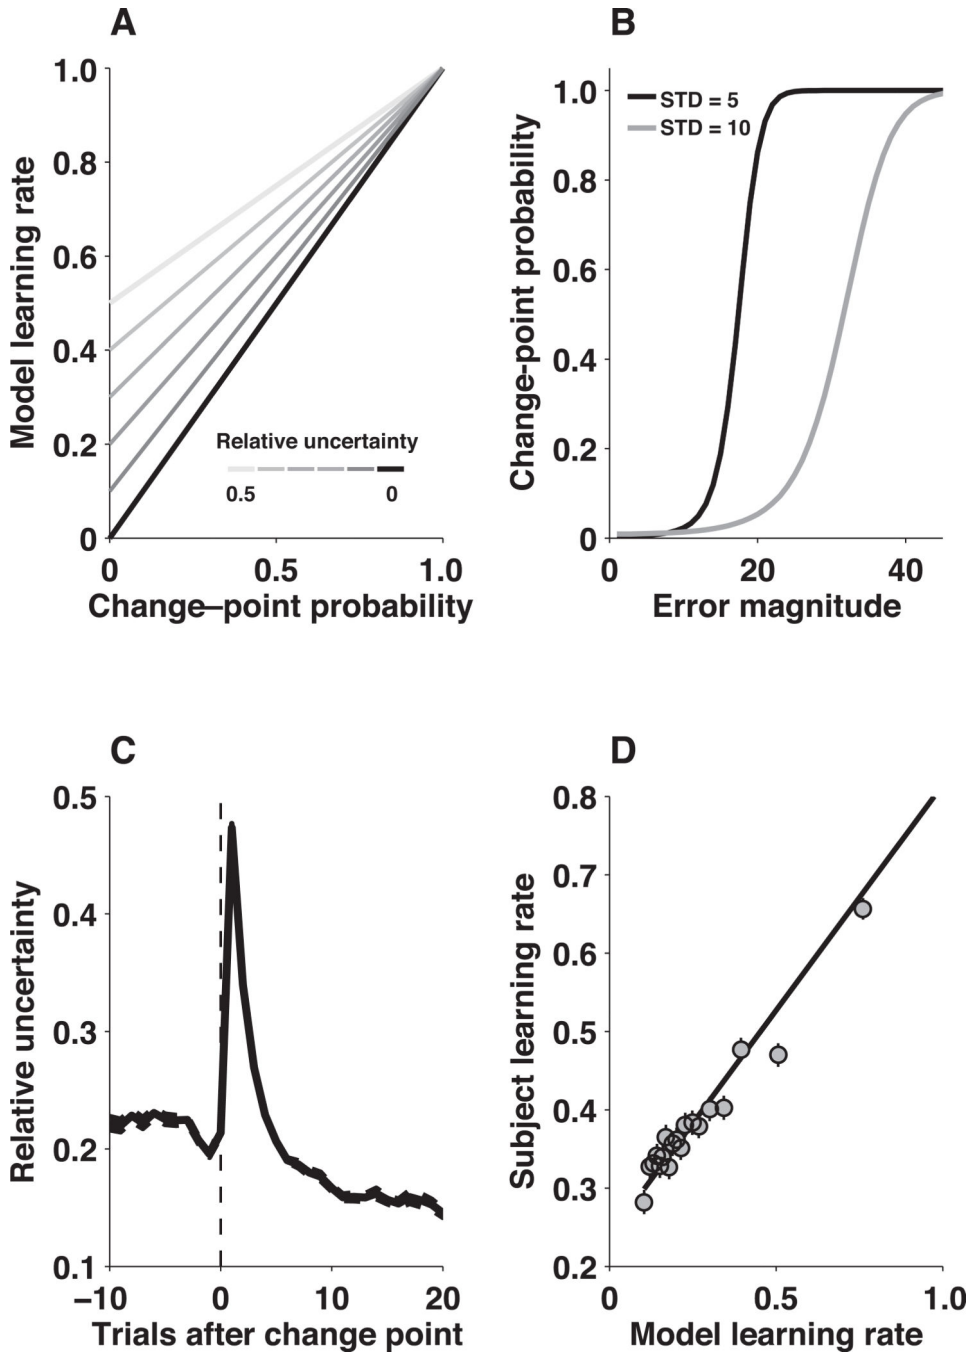


Figure 3. Reduced Bayesian model. A, Learning rate as a function of change-point probability (abscissa) and relative uncertainty (line shading), as computed by the model. B, Change-point probability computed by the model as a function of error magnitude (abscissa) for the two different noise conditions, as indicated, computed for a given relative uncertainty (equal to 0.02 for this figure). C, Mean±SEM relative uncertainty computed by the model aligned to change points from all sequences experienced by the subjects for the two different noise conditions. D, Trial-by-trial comparison of subject and model learning rates. Model

learning rates were computed using the same sequence of outcomes experienced by each subject. Points and error bars are mean \pm SEM data from all subjects grouped into 20 five-percentile bins according to the corresponding model learning rate. The solid line is a linear fit to the unbinned data ($r=0.33$, $p<0.001$).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

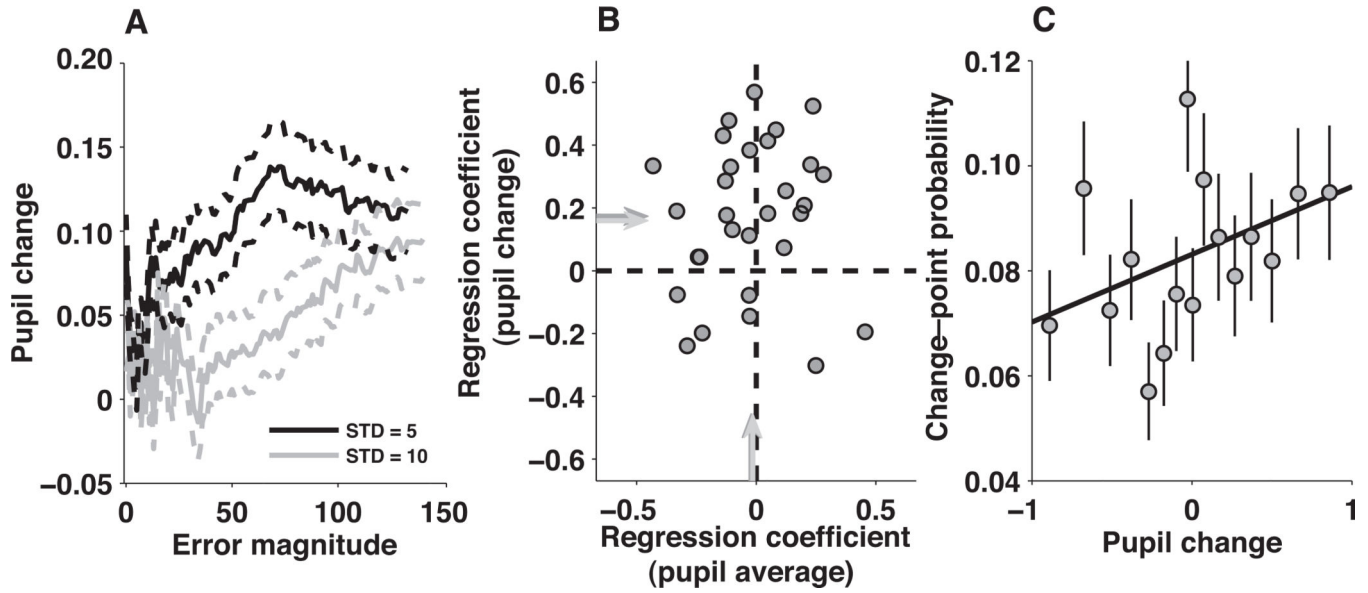


Figure 4.

Relationship between pupil change and change-point probability. A, Mean \pm SEM pupil change from all trials and all subjects for running bins of 150 trials, binned according to the absolute prediction error and sorted by noise, as indicated. B, Regression coefficients describing the linear relationship between change point-probability (p_{CH}) and z -scored pupil change (z_{PC} , ordinate) versus the regression coefficients describing the linear relationship between p_{CH} and z -scored pupil average (z_{PA} , abscissa). Points are regression coefficients computed for each subject individually, using the four-parameter regression model. Arrows indicate mean values from this model (dark, equal to $0.174 z_{PC}/p_{CH}$, t -test for H_0 : mean=0, $p<0.001$ for the ordinate, $-0.022 z_{PA}/p_{CH}$, $p=0.58$ for the abscissa) or from the full model (light, equal to $0.148 z_{PC}/p_{CH}$, $p<0.001$ for the ordinate, $-0.014 z_{PA}/p_{CH}$, $p=0.70$ for the abscissa). Dark arrows are partially occluded by light ones. C, Change-point probability from the reduced Bayesian model versus pupil change. Points and error bars are mean \pm SEM data from all subjects grouped into 20 five-percentile bins. The solid line is a linear fit to the unbinned data (slope = $0.012 p_{CH}/z_{PC}$, $p<0.001$ for H_0 : slope=0).

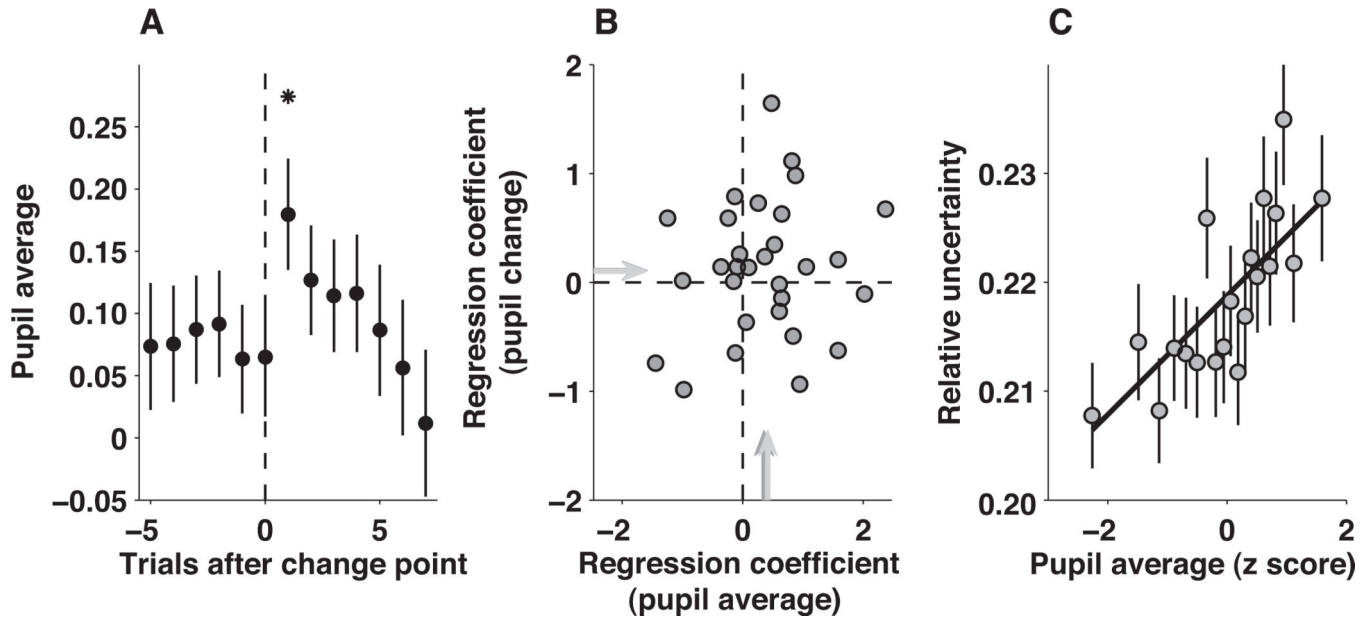


Figure 5.

Relationship between pupil diameter and relative uncertainty. A, Mean \pm SEM pupil average from all subjects as a function of trials relative to task change points. Asterisk indicates trials differing significantly from all other trials (permutation test for H_0 : equal means after correction for multiple comparisons, $p < 0.05$). B, Regression coefficients describing the relationship between relative uncertainty (RU) and z -scored pupil change (z_{PC} , ordinate) versus the regression coefficients describing the relationship between RU and z -scored pupil average (z_{PA} , abscissa). Points are regression coefficients computed for each subject individually, using the four-parameter regression model. Arrows indicate mean values from this model (dark, equal to $0.135 z_{PC}/RU$, t -test for H_0 : mean=0, $p=0.28$ for the ordinate, $0.35 z_{PA}/RU$, $p < 0.05$ for the abscissa) or from the full model (light, equal to $0.127 z_{PC}/RU$, $p=0.24$ for the ordinate, $0.40 z_{PA}/RU$, $p < 0.01$ for the abscissa). Dark arrows are partially occluded by light ones. C, Relative uncertainty from the reduced Bayesian model versus pupil average. Points and error bars are mean \pm SEM data from all subjects grouped into 20 five-percentile bins. The solid line is a linear regression to unbinned data (slope = $0.0055 RU/z_{PA}$, $p < 0.001$ for H_0 : slope=0).

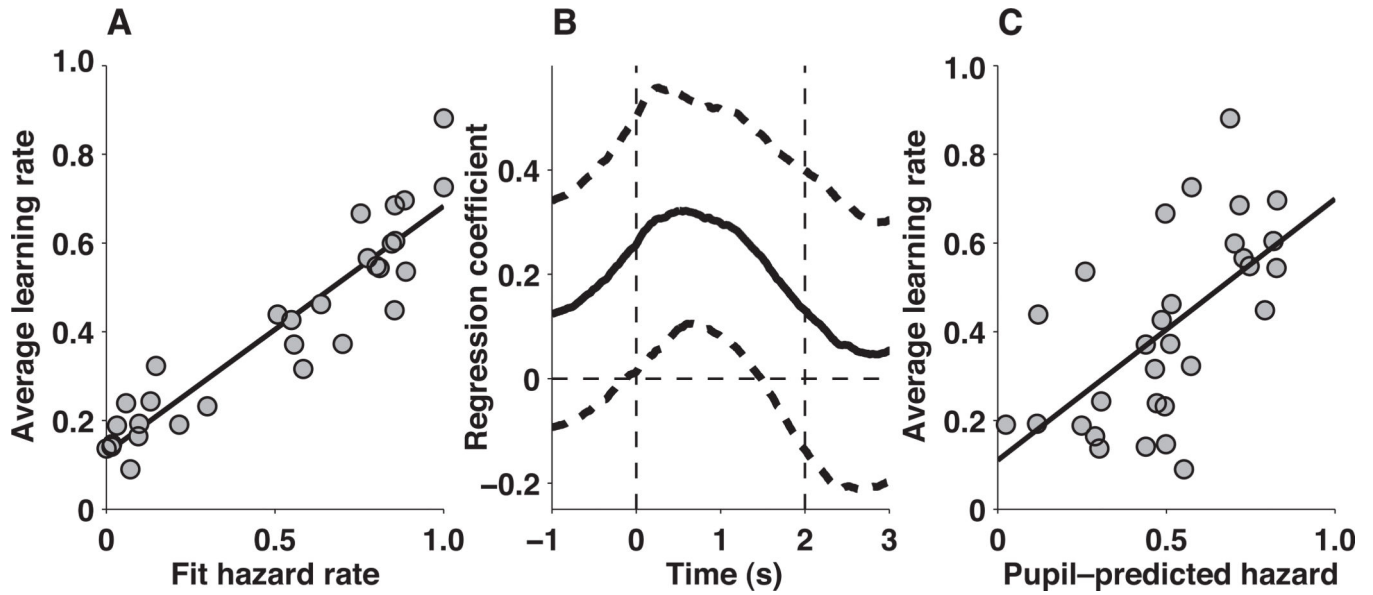


Figure 6.

Individual differences in learning rate, hazard rate, and pupil diameter. A, Mean learning rate per subject versus the hazard rate of the reduced Bayesian model that best fit that subject's performance (points). The solid line is a linear fit ($r=0.93$, $p<0.001$). B, Regression coefficients describing the relationship between fit hazard rates and bin-by-bin pupil measurements across subjects, computed in sliding 8.3-ms bins and aligned to outcome presentation (time=0). Dotted lines indicate 95% confidence intervals. C, Relationship between pupil-predicted hazard rate and average learning rate for each subject (points). Pupil-predicted hazard rates were computed using a linear regression model that included both shape and magnitude of the average pupil response for each subject (see Methods). The solid line is a linear fit ($r=0.59$, $p<0.001$).




Figure 7.

Pupil metrics predict learning rate. A, Regression coefficients describing the linear, trial-by-trial relationships between pupil change and the subsequent learning rate (ordinate) and between pupil average and the subsequent learning rate (abscissa). Points are regression coefficients computed for each subject individually, using a four-parameter regression model that also included trial number and block number as covariates. B, The relationship between learning rate and pupil parameters depended on the subject's baseline pupil response. For each subject, the sum of the regression coefficients from panel A are plotted as a function of the pupil-predicted hazard rate from Figure 6C. The line is a linear fit ($r = -0.059$, $p < 0.001$). C, Predicted versus actual learning rate. Both values are z-scored per subject. Data from all subjects are grouped into 20 equally sized bins of predicted learning rate. The line is a linear fit to the unbinned data (Slope = $0.052 \text{ zActual/zPredicted}$, $p < 0.001$ for H_0 : slope=0).




Figure 8.

Effects of the pupil manipulation. A, Evoked changes in pupil diameter. For each subject, pupil average (ordinate) and pupil change (abscissa) were z-scored across all trials. Each point represents the difference in the mean z-scores for auditory switch versus non-switch trials for an individual subject. Positive values indicate larger values on switch trials. B, Evoked changes in learning behavior. For each subject, learning rate was z-scored across all trials and fit to a cumulative Weibull as a function of error magnitude for each noise condition, to account for the relationship shown in Fig. 4A. Each point represents the difference in the mean value of the residuals from these fits for auditory switch versus non-switch trials for an individual subject, separated by trials in which the initial pupil diameter was smaller (ordinate) or larger (abscissa) than its median value. Positive values indicate larger learning rates on auditory switch trials. C, A possible relationship between learning and arousal based on an “inverted U” (light gray, modeled as Gaussian). A given change in learning for a given a change in pupil metrics (ordinate), plotted as a function of baseline pupil diameter (abscissa), is shown for: 1) the hypothesized Gaussian (its derivative is shown in dark gray), 2) the measured effects of the auditory manipulation (open points), and 3) the measured relationship between pupil metrics and learning rate during non-manipulation sessions. See Methods for details.