

Anxious individuals have difficulty learning the causal statistics of aversive environments

Michael Browning¹, Timothy E Behrens¹, Gerhard Jocham^{1,2}, Jill X O'Reilly¹ & Sonia J Bishop^{1,3}

Statistical regularities in the causal structure of the environment enable us to predict the probable outcomes of our actions. Environments differ in the extent to which action-outcome contingencies are stable or volatile. Difficulty in being able to use this information to optimally update outcome predictions might contribute to the decision-making difficulties seen in anxiety. We tested this using an aversive learning task manipulating environmental volatility. Human participants low in trait anxiety matched updating of their outcome predictions to the volatility of the current environment, as predicted by a Bayesian model. Individuals with high trait anxiety showed less ability to adjust updating of outcome expectancies between stable and volatile environments. This was linked to reduced sensitivity of the pupil dilatory response to volatility, potentially indicative of altered norepinephrergic responsivity to changes in this aspect of environmental information.

From what to eat, which task to prioritize, whether to call a friend, go on a date or book a family holiday, the decisions we take are guided by what we expect the outcomes of those decisions to be. Sometimes it can seem extremely difficult to work out which course of action is least likely to result in a feared outcome, whether avoiding being the one made redundant or preventing an argument turning into a breakup. Individuals who suffer from anxiety focus disproportionately on the potential occurrence of future negative outcomes and whether or not they can be averted. Attempts to understand this have led to suggestions that anxiety may be linked to difficulty in estimating the probability or severity of future negative events or in combining these estimates when choosing between actions^{1–3}. Advances in computational modeling provide us with the tools to investigate which of the mechanisms involved in decision-making are disrupted in anxious individuals.

There are often times when the correct course of action seems unclear, and indeed our best choice of what to do leads unexpectedly to an aversive outcome. This may in turn lead to the next decision being harder still. Should one switch behavioral patterns or stay one's course? Anxious individuals are reported to be highly intolerant of situations characterized by uncertainty as to what will happen or which course of action should be followed^{4,5}. They are more likely to endorse finding such situations distressing, leading to a sense of immobilization^{4,5}. One possible explanation for this is that individuals prone to anxiety may have difficulty estimating outcome likelihood when there are sources of uncertainty complicating the action-outcome relationship.

In the context of decision-making models, a number of alternate forms of uncertainty are recognized^{6–8}. Here, we focus on two types. One source of uncertainty is produced by noise in the relationship between actions and outcomes, such as occurs if an action only leads to a given outcome on a proportion of the occasions it is performed.

A second source of uncertainty is produced when the underlying causal structure is non-stationary, or volatile, for example, when action-outcome contingencies switch and an action that was primarily associated with a given outcome becomes predominantly associated with another. If unexpected or 'surprising' outcomes are caused by noise, then current action choices are optimally determined by averaging over the outcomes of many previous actions. In contrast, if surprising outcomes are caused by a change in action-outcome contingencies in a volatile environment, then only the most recent events should be used to guide action choice⁶. In terms of formal reinforcement accounts of learning^{9,10}, a higher learning rate should be implemented when the environment is volatile than when it is stable. Computational studies of decision-making reveal that healthy volunteers do indeed adapt their learning rate in response to changes in environmental volatility^{8,11}. Specifically, the behavior of participants is consistent with faster updating of action-outcome contingencies in volatile than in stable environments. This change in learning behavior occurs even when changes in environmental volatility are not explicitly cued and has been shown to closely track an optimal Bayesian decision-making strategy^{8,11}.

Individuals prone to anxiety might either show a general deficit in updating action-outcome estimates following unexpected aversive outcomes or a specific problem with adjusting the speed of updating (that is, learning rate) to reflect the stability or volatility of the current environment. Recent Pavlovian fear conditioning findings suggest that trait-anxious individuals struggle to adjust fear downregulation to reflect changes in stimulus-stimulus contingencies between contexts¹². If individuals prone to anxiety have a particular difficulty in processing contingencies that change over time or between contexts, we might predict that they will show a deficit in using changes in environmental volatility to infer whether or not action-outcome contingencies have changed, and to alter their behavioral choices

¹Functional MRI of the Brain Centre, Nuffield Department of Clinical Neurosciences, John Radcliffe Hospital, Oxford, UK. ²Center for Behavioral Brain Sciences, University of Magdeburg, Magdeburg, Germany. ³Department of Psychology and Helen Wills Neuroscience Institute, University of California, Berkeley, California, USA. Correspondence should be addressed to S.J.B. (sbishop@berkeley.edu).

Received 3 August 2014; accepted 29 January 2015; published online 2 March 2015; doi:10.1038/nn.3961

accordingly. Such a deficit would mean difficulty in differentiating contexts or periods of time in which unexpected aversive outcomes should be ignored as chance events and those in which unexpected aversive outcomes are likely to signal a change in action-outcome contingencies, and as such should lead to a change in action. This could potentially explain why individuals at trait risk for anxiety show intolerance of uncertainty, hesitation and poor decision-making in real-life settings in which the relationship between alternate courses of actions and avoidance of future negative events is often uncertain.

Our primary hypothesis was therefore that, when attempting to avoid aversive outcomes, high trait-anxious individuals would be less able than low trait-anxious individuals to adjust their updating of action-outcome contingencies in response to changes in environmental volatility. We focused on individual differences in trait anxiety to investigate mechanisms that might convey vulnerability to anxiety disorders, without the confounding effects of psychotropic medication or chronic illness. We used an aversive learning task that consisted of two blocks: in one, the relationship between participants' actions and the outcomes of those actions was stable; in the other, this relationship was volatile, repeatedly changing over time. This was equivalent in design to a previously used task⁸, but with action choice to gain financial reward replaced by action choice to avoid bursts of electrical stimulation. The difference in participants' learning rate between the stable and volatile task blocks provided a measure of participants' ability to adapt their learning to changes in environmental volatility. We predicted that the difference in learning rate between stable and volatile task blocks would be reduced in high trait-anxious individuals.

Recent findings indicate that pupil dilation may track important changes in the causal statistics of the environment, including changes in environmental volatility^{13,14}. This adds to accounts arguing that pupil diameter may reflect, among other influences, activity of the locus coeruleus norepinephrine system^{15–17} and models proposing norepinephrinergic control over the learning of environmental uncertainty⁶. Given these findings, and established reports of altered norepinephrinergic function in anxiety^{18,19}, we also hypothesized that trait anxiety-related deficits in differential learning between volatile and stable task blocks would be accompanied by a reduced pupil dilatory response to environmental volatility.

Consistent with our predictions, high trait-anxious individuals showed a specific deficit in adjusting learning rate in response to changes in environmental volatility in our aversive learning task. This was associated with a reduced pupil response to trial-wise estimates of environmental volatility. These results provide evidence that

trait vulnerability to anxiety is associated with impoverished use of environmental statistics, especially that pertaining to environmental volatility, to determine the extent to which to update action-outcome contingencies when attempting to avoid aversive outcomes. This may represent a core deficit underlying impoverished decision-making in individuals at elevated risk of developing anxiety disorders.

RESULTS

We asked 31 participants (22 females) screened to have a range of scores on the Spielberger State-Trait Anxiety Inventory (STAI²⁰) trait subscale to perform an isoluminant version of a two-armed bandit learning task⁸ in which outcomes were moderately painful electrical shocks. Before participants completed the task, the intensity of the electrical shocks were calibrated so that the maximum level administered had a subjective pain level of 7 on a scale of 1 (minimal pain) to 10 (worst possible pain) (Online Methods and **Supplementary Fig. 1**).

On each trial, participants had to choose one of two shaped, isoluminant, gabor patches, either of which might result in the delivery of an electrical shock (**Fig. 1a**). In the stable task block, one of the two shapes consistently predicted shock delivery with a probability of 75%, and the other shape resulted in shock delivery on the remaining trials. In the volatile task block, the shape most predictive of shock delivery reversed on five occasions across the block of trials (**Fig. 1b**). On each trial, each shape was associated with a specific magnitude of electrical shock that would be delivered if the shape was chosen and if it resulted in shock administration on that trial. This magnitude was displayed in the center of each shape. Magnitude values were scaled between 1 and 99, where 1 equated to a subjective pain level of 1 and 99 to a subjective pain level of 7. The trial-wise magnitudes of the shocks for each shape were chosen from independent random distributions (Online Methods). To perform the task optimally, participants had to integrate the information about shock magnitude and shock probability, the latter needing to be inferred from the outcome of previous trials. The shape chosen and the time taken to select it was recorded. Pupil dilatory responses were also measured. As a result of equipment failure, behavioral and pupil data were lost from one participant (excluded) and pupil data from two further participants.

Effect of environmental volatility on learning rate

Learning rate α reflects the extent to which participants' choice behavior was guided by the outcomes of recent actions versus those further back in the individual's experience. At high learning rates, choice behavior is strongly controlled by the outcomes of recent actions.

Figure 1 Task structure. **(a)** Example trial. Participants had to choose one of two shaped Gabor patches. Each shape contained a two digit number that indicated the magnitude of electrical shock that might be received should that shape be chosen. Following option presentation, onset of a response cue indicated that participants could make their choice. After response, a variable interval was followed by outcome delivery. The shape associated with the electrical shock for that trial was displayed in the center of the screen for 2 s. If the participant had chosen this shape, an electrical shock of the indicated magnitude was delivered at the onset of the outcome period. **(b)** Outcome probabilities across the course of the task. The task comprised two blocks. In the stable block (shaded), one shape (for example, the circle) had a 75% probability of resulting in an electrical shock being delivered if it was chosen; the other shape (for example, a square) had a 25% probability of resulting in shock delivery. In the volatile block (unshaded), the probability that choice of a given shape would result in shock delivery switched every 20 trials between 80 and 20%. Participants were randomly assigned to complete the task with the stable block first (as shown) or with the volatile block first.

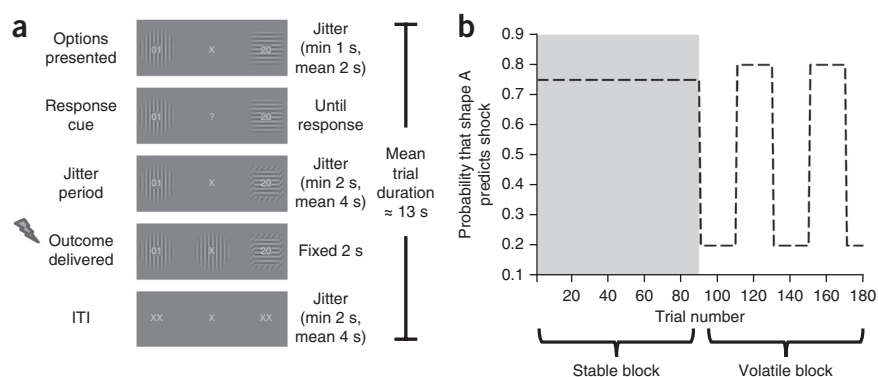
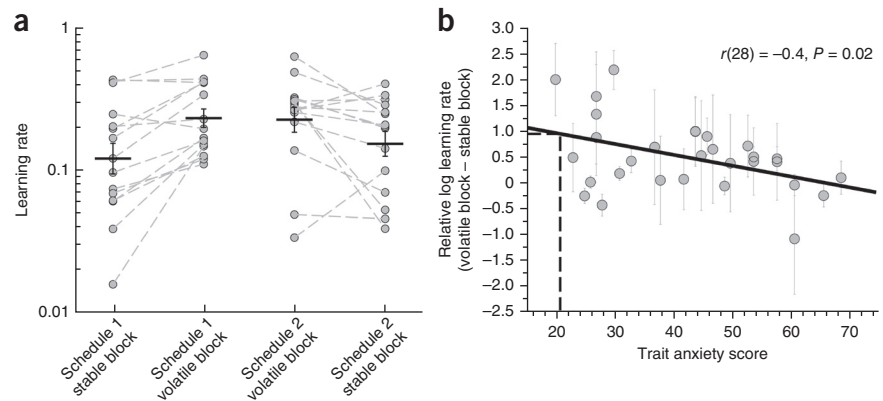


Figure 2 Estimates of participants' learning rates. (a) Participants' choices during the stable and volatile blocks of the aversive learning task were fitted with a Rescorla Wagner learning model in which learning rate was allowed to vary. Estimates of individual participants' learning rates are displayed (circles) separately for the stable and volatile blocks for the two task schedules (schedule 1 = stable task block first; schedule 2 = volatile task block first). A logarithmic scale is used. Black lines display mean (\pm s.e.m.) of participant learning rates, gray dashed lines link the learning rates in volatile and stable blocks for each participant. Participants showed higher learning rates in the volatile versus stable blocks regardless of the order in which they were completed, $F(1,28) = 16.3$, $P < 0.001$. (b) The relative log learning rate for the volatile versus the stable blocks (i.e. $\log(\text{LR in volatile block}) - \log(\text{LR in stable block})$) was negatively correlated with participant trait anxiety, $r(28) = -0.42$, $P = 0.02$. The black dashed line indicates the degree to which the model of an optimal Bayesian learner (as described previously⁸) adjusted its learning rate. As can be seen, low trait-anxious participants altered their learning rates to a similar degree to the Bayesian Learner, with high trait-anxious participants showing a reduced adaptation of learning rate between the volatile and stable blocks of the task. Error bars represent the s.d. of the estimated parameters from the behavioral model for each subject.



More precisely, the difference between the expected and actual outcome on a given trial has a large effect on change in outcome expectancy and subsequent choice behavior. In contrast, at low learning rates, surprising outcomes lead to little change in outcome expectancy and behavior. We estimated participants' learning rates in the stable and volatile task blocks by fitting a simple Rescorla Wagner learning model to their choices in each task block (Online Methods and **Supplementary Modeling Note**). We assessed whether participants, as a group, adapted their learning rate in response to the change in environmental volatility between the stable and volatile blocks. Consistent with prior findings for reward⁸, which we replicated (**Supplementary Fig. 2a**), participants' learning rates were higher in the volatile than in the stable block of our aversive learning task ($F(1,28) = 16.3$, $P < 0.001$; **Fig. 2a** and **Supplementary Fig. 3a**), regardless of the order in which the two blocks were completed ($F(1,28) = 1.0$, $P = 0.3$).

In addition to updating their expectancies, or 'beliefs', to reflect current action-outcome contingencies, participants needed to transform these beliefs into action choices. We formally described this process using a softmax action selector model that contains two free parameters: a risk preference term that controls the relative weight of outcome probability versus shock magnitude used in the calculation of expected value (that is, whether a given participant tends to prefer low-probability, high-magnitude shocks over high-probability, low-magnitude shocks), and an inverse decision temperature term, an estimate of the extent to which calculations of expected value influence choice⁸ (**Supplementary Modeling Note**). Across participants, there was no effect of block volatility, block order or their interaction on either of these parameters ($P > 0.1$). This indicates that the effect of the task volatility manipulation was confined to participants' learning rates, consistent with expectations based on previous findings⁸.

Trait anxiety is linked to reduced adaptation of learning rate

Consistent with our predictions, elevated trait anxiety was associated with a reduced change in learning rate between blocks ($r(28) = -0.42$, $P = 0.02$; **Fig. 2b**). An effect of anxiety was only seen for change in learning rate between stable and volatile blocks; there was no significant relationship between anxiety and mean learning rate across blocks ($r(28) = 0.1$, $P = 0.6$) or with learning rate in each block (volatile block, $r(28) = -0.1$, $P = 0.6$; stable block, $r(28) = 0.26$, $P = 0.16$). In other words, low trait-anxious individuals were able to adjust their learning appropriately between blocks, learning (that is, updating

outcome estimates) fast when the world was fast-changing and slowly when the world was stable. By contrast, high trait-anxious individuals were not able to learn in this flexible manner, instead learning similarly in both blocks. It is notable that mean learning rate was not modulated by trait anxiety. This suggests that high trait-anxious individuals did succeed in updating outcome expectancies following surprising outcomes, but that they were unable to modulate this on the basis of the volatility of the current environment.

The optimal adjustment of learning rate between stable and volatile conditions has previously been formally described in a Bayesian framework for a reward-based version of the task used here⁸. In this prior study, participants' behavior closely tracked that of an ideal Bayesian learner. We used the same Bayesian formalization to provide an estimate of the information available to participants during the aversive learning task and the degree to which learning rate should be adjusted between blocks (Online Methods and **Supplementary Modeling Note**). Participants with low levels of trait anxiety altered their learning rate between the stable and volatile blocks to an equivalent degree to the ideal Bayesian learner⁸. As trait anxiety levels increased, participants diverged increasingly from the optimal change in learning rate between stable and volatile blocks described by the Bayesian learner model (**Fig. 2b**), showing reduced adaptation of learning rate. Additional modeling revealed that, even when multiple parameters, including a decay term, were allowed to compete for influence over a dynamic learning rate, trait anxiety was uniquely associated with a reduction in the influence of environmental volatility (**Supplementary Modeling Note** and **Supplementary Fig. 4**).

As outlined earlier, when learning rate is high, recent trial outcomes inform behavioral choice to a greater extent. Consistent with elevated trait anxiety being associated with reduced adaptation of learning rate, high trait-anxious individuals showed less adjustment between volatile and stable blocks in their use of recent trial outcomes to inform choices than low trait-anxious individuals (**Supplementary Fig. 3c,d**).

An additional illustration of the consequences of non-optimal learning for actual choice behavior is provided by consideration of performance on those trials on which choice is hardest. Following previous work⁸, this was defined as those trials with a difference in expected value between the two alternatives of 5 or less. Trait anxiety was significantly positively associated with the number of shocks received on these trials ($r(28) = 0.37$, $P = 0.02$, one-tailed;

Supplementary Fig. 5a). Furthermore, across participants, the number of shocks received on these ‘difficult’ trials was negatively correlated with the extent of adaptation of learning rate between stable and volatile blocks ($r(28) = -0.31$, $P = 0.05$, one-tailed; **Supplementary Fig. 5b**).

For completeness, we also examined whether estimates of the risk preference and decision temperature parameters from the softmax action selector model varied as a function of participant trait anxiety levels. There was no significant relationship between trait anxiety and these parameter estimates across or within blocks, or as a function of block type ($P > 0.1$; **Supplementary Fig. 6**). Furthermore, the relationship between trait anxiety and change in learning rate between blocks remained significant when controlling for either change in decision temperature ($r_{\text{partial}}(27) = -0.38$, $P = 0.04$) or risk preference ($r_{\text{partial}}(27) = -0.41$, $P = 0.03$).

Responses to surprising outcomes are not modulated by trait anxiety

Adaptation to environmental volatility requires processing of the surprise value or unexpectedness of the chosen action’s outcome and assessment as to whether or not there has been a change in action–outcome contingencies. This may be an implicit rather than an explicit process. One possibility is that trait-anxious individuals’ difficulty in adapting to changes in environmental volatility in our aversive learning task might arise secondarily to a more general deficit in processing the surprise value of outcomes. As outlined above, this seems unlikely given that there was no effect of trait anxiety on mean learning rate across task blocks. To further investigate this, we examined whether trait anxiety modulated choice reaction time as a function of the surprise value of the previous trial’s outcome. In the reinforcement learning literature, it has been reported that, when an outcome is surprising, participants typically show slowed decision-making on the subsequent trial²¹. Consistent with this, we observed that, across participants, surprising outcomes were indeed associated with slowed choice reaction times on the subsequent trial (that is, across participants, the beta weight of outcome surprise on trial n on individual participant’s reaction times on trial $n + 1$ differed significantly from zero ($t(29) = 2.3$, $P = 0.03$; Online Methods). However, the effect of surprise on participants’ reaction times was not modulated by participant trait anxiety ($r(28) = 0.1$, $P = 0.5$). This finding is consistent with high trait-anxious participants being specifically insensitive to changes in environmental volatility rather than simply unable to process surprising outcomes. In these analyses, trial-wise outcome surprise was calculated, using the Bayesian learner, by taking the negative logarithm of the conditional probability of the outcome observed given the prior estimated probability of (that is, ‘belief’ in) that outcome²² (Online Methods and **Supplementary Modeling Note**). An alternate version of these analyses using a non-Bayesian measure of surprise produced similar results (**Supplementary Fig. 7a**).

Pupil dilation tracks both volatility and surprise

Recent accounts have described an increase in pupil dilation in response to both environmental volatility and outcome surprise^{13,14,23}. Prior to examining differences between high and low trait-anxious individuals, we sought to establish, at a cross-group level, the pupil dilatory response to volatility and outcome surprise in our current study.

Trial-wise pupil responses to choice outcome were baseline corrected (Online Methods). The post-outcome period was sampled using 6,000 1-ms bins. Regression analyses were conducted for each of these bins to examine the extent to which trial-wise estimates of

volatility and outcome surprise derived from the Bayesian learner predicted pupil dilation (controlling for outcome, that is, shock or no shock, shock magnitude, and expected value of the chosen and not chosen shapes). These regression analyses produced two time series of beta-weights: one for volatility and one for surprise. These time series were down-sampled to give beta estimates of the effects of surprise and volatility on pupil dilation for sequential 1-s time bins across the outcome period.

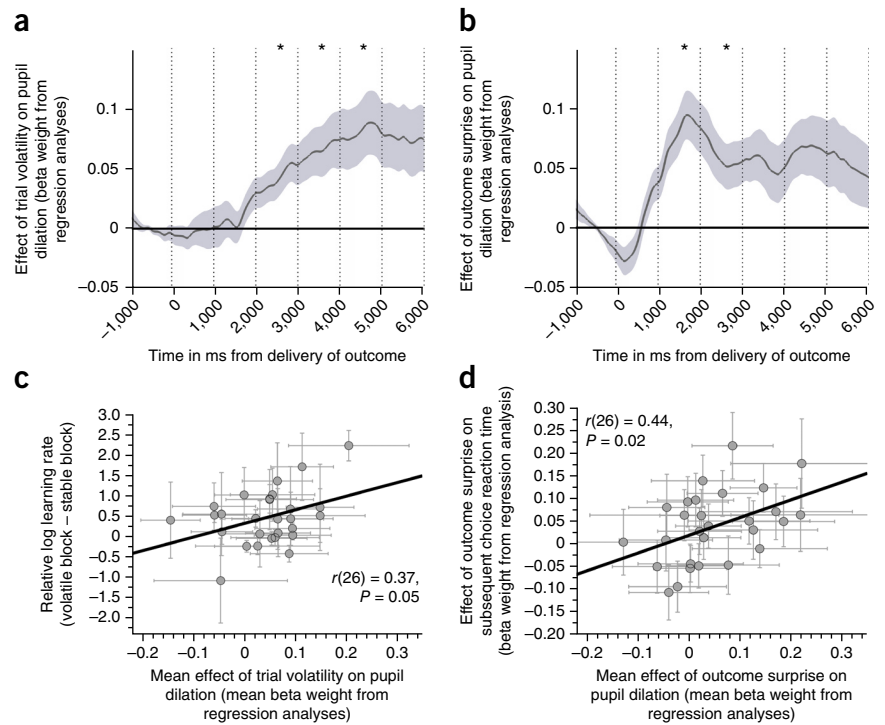
Group-level analyses revealed that trial volatility was significantly associated with an increase in pupil diameter following outcome delivery ($F(1,26) = 9.8$, $P = 0.004$). Bonferroni-corrected one-sample t tests, performed for each time bin, indicated that this effect was significant from 2–5 s post outcome (**Fig. 3a**). Outcome surprise was also positively associated with an increase in pupil diameter ($F(1,26) = 9.2$, $P = 0.005$). The effect of surprise was observed slightly earlier than that of volatility, being significant from 1–3 s post outcome ($P < 0.05$, Bonferroni corrected; **Fig. 3b**). For each individual, we also calculated single trial-wise summary measures of pupil responsivity to volatility and outcome surprise. These were estimated using the mean beta-weight of the surprise and volatility regressors across the whole 6-s post-outcome period. Correlational analyses using these summary measures revealed that participants’ mean pupil response to trial-wise volatility predicted the degree to which they adjusted their learning rate between the volatile and stable blocks of the task ($r(26) = 0.37$, $P = 0.05$; **Fig. 3c**). In addition, participants’ mean pupil response to outcome surprise predicted the extent to which they showed choice reaction time slowing as a function of the unexpectedness of the previous trial’s outcome ($r(26) = 0.44$, $P = 0.02$; **Fig. 3d**). These pupil dilation parameters may reflect changes in activity in the locus coeruleus norepinephrine system, although additional dopaminergic or cholinergic influences cannot be ruled out. The current findings are consistent with the existence of a functional relationship between the alterations in neurotransmission that underlie these pupilometry changes and the mechanisms that enable environmental statistics to be used to guide learning about the causal structure of the environment^{13–16}.

Trait anxiety modulates the pupil response to volatility

As reported above, we observed that high trait-anxious individuals showed impoverished adjustment of learning rate between the stable and volatile blocks of the aversive learning task employed here, but did not show impoverished adjustment of choice reaction times following surprising outcomes. This is consistent with a specific deficit in the use of environmental volatility to adjust action–outcome updating and, through this, to guide decision-making. We were therefore interested in whether high trait-anxious participants would show a reduced pupil dilatory response to environmental volatility alone. This was indeed what we observed. Elevated trait anxiety was associated with a decreased mean pupil response to volatility ($r(26) = -0.51$, $P = 0.005$; **Fig. 4a,b** and **Supplementary Fig. 8c,d**). In contrast, there was no modulatory effect of trait anxiety on the mean pupil response to surprise ($r(26) = -0.16$, $P = 0.4$; **Fig. 4c,d** and **Supplementary Fig. 8e,f**). Together with the absence of any significant relationship between trait anxiety and the behavioral response to surprise, these findings are consistent with trait anxiety being specifically associated with an impoverished ability to process, and adapt to, changes in environmental volatility.

To rule out the possibility that the effects of volatility on pupil dilation post outcome, or the modulation of this by anxiety, were a result of changes in pupil reactivity with time on task, we reran the pupil analyses reported above with two additional control regressors

Figure 3 Post-outcome pupil dilation tracks both environmental volatility and outcome surprise. **(a,b)** Time courses for the effect of trial-wise estimates of volatility **(a)** and surprise **(b)** on pupil dilation following presentation of the outcome. The graphs show the mean across participants ($n = 28$) of the beta weights obtained by regressing post outcome pupil dilation against trial-wise estimates of environmental volatility and outcome surprise. Post-outcome pupil dilation was greater for trials in which environmental volatility was high ($F(1,26) = 9.8$, $P = 0.004$) and the outcome was surprising ($F(1,26) = 9.2$, $P = 0.005$). Asterisks indicate 1-s time bins in which the effect of volatility or surprise on pupil dilation post-outcome differed significantly from zero (Bonferroni corrected for multiple comparisons, corrected P values < 0.05 , two tailed). The effect of trial-wise volatility was longer lasting and had a later onset than that of outcome surprise. **(c)** The degree to which an individual's pupil tracked volatility (calculated as the mean beta weight across the 6-s post-outcome period) predicted change in learning rate between volatile and stable blocks ($r(26) = 0.37$, $P = 0.05$). **(d)** The degree to which an individual's pupil tracked surprise predicted the extent of surprise-related choice reaction time slowing on the subsequent trial ($r(26) = 0.44$, $P = 0.02$). Shaded regions in **a** and **b** represent the s.e.m. Error bars in **c** and **d** represent the s.d. of the regression coefficients (beta weights) from the pupil analysis and the parameter estimates from the behavioral model for each subject.



that represented the trial number in each block. Following the addition of these control regressors, the post-outcome pupil response to volatility ($F(1,26) = 4.3$, $P = 0.049$) and surprise ($F(1,26) = 8.6$, $P = 0.007$) remained significant and the relationship between trait anxiety and the mean post outcome pupil response to volatility became, if anything, stronger ($r(26) = -0.52$, $P = 0.004$; **Supplementary Fig. 9**). Similarly, inclusion of regressors representing the interaction of volatility by outcome (shock versus no shock; **Supplementary Fig. 10a**),

surprise by outcome (**Supplementary Fig. 10b**), and the interaction between volatility and surprise (**Supplementary Fig. 11**) did not alter the relationship between trait anxiety and the post-outcome pupil response to volatility, which remained robust in all cases ($r(26) \leq -0.49$, $P \leq 0.007$). Notably, the effect of volatility on post-outcome pupil dilation was not modulated by outcome type (**Supplementary Fig. 10a**), nor was the effect of surprise (**Supplementary Fig. 10b**).

We further addressed the possibility that high anxious individuals' pupillary response does in fact track environmental volatility, but in a manner less like the Bayesian learner

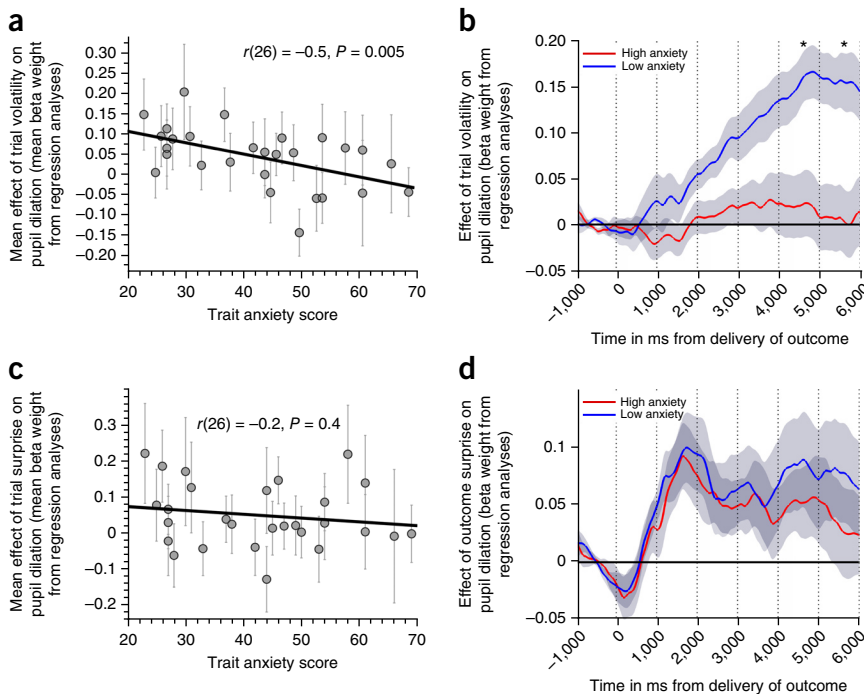


Figure 4 The relationship between trait anxiety and post-outcome pupil dilation as a function of trial-wise estimates of volatility and surprise. **(a)** The degree to which participants' pupil dilation, post-outcome, tracked environmental volatility was negatively related to trait anxiety ($r(26) = -0.51$, $P = 0.005$). **(b)** Using a median split on participants' trait anxiety scores, low anxious participants ($n = 15$) showed a clear pupil response to environmental volatility, whereas high anxious participants ($n = 13$) did not (asterisks indicate 1-s time bins in which Bonferroni-corrected t tests differed between the groups at $P < 0.05$ corrected, two-tailed). **(c)** Pupil response to outcome surprise was not related to individual differences in trait anxiety ($P = 0.4$). **(d)** This is illustrated using a median split on trait anxiety. Error bars in **a** and **c** represent the s.d. of the regression coefficients (beta weights) from the pupil analysis for each subject. Shaded regions in **b** and **d** represent the s.e.m.

than low anxious participants, by conducting additional analyses using non-Bayesian estimates of environmental volatility and outcome surprise (**Supplementary Fig. 7b,c**). These additional analyses replicated the main findings of an inverse relationship between trait anxiety and the effect of volatility, but not surprise, on the post outcome pupil response.

Finally, we investigated the relationship between adaptation of learning between task blocks and baseline, or tonic, pupil dilation. Electrophysiological studies have revealed that the locus coeruleus norepinephrine system displays two modes of activity: periods of reduced tonic activity associated with increased phasic responses to salient stimuli and periods of increased tonic activity associated with reduced phasic responses¹⁷. Consistent with this observation and the suggested influence of norepinephrinergic activity on pupil dilation, a previous pupillometry study reported that reduced baseline pupil size (argued to reflect, at least in part, tonic norepinephrinergic activity) was associated with increased phasic pupil responses¹⁶. As reported above (**Fig. 3**), we observed a positive association between environmental volatility and the magnitude of the phasic post-outcome pupil response. We therefore tested whether there was also evidence for a reduction in tonic baseline pupil size during volatile periods and whether this predicted participant behavior, and varied as a function of anxiety.

Following previously described methodology¹⁶, we obtained a measure of the degree to which tonic pupil size varied as a function of trial-wise volatility by analyzing non-baseline corrected pupil data from the 1-s time bin before the outcome was presented (that is, the baseline period). There was no relationship between mean baseline pupil size across the task and trait anxiety ($r(26) = 0.1$, $P = 0.6$). Across participants, periods of increased volatility were associated with reduced baseline pupil size ($t(27) = 2.5$, $P = 0.02$)¹⁶. However, differences between participants in the magnitude of baseline pupil size suppression were not significantly correlated with differences in learning rate between volatile and stable task blocks ($r(26) = -0.14$, $P = 0.5$), nor did it vary significantly as a function of trait anxiety ($r(26) = 0.08$, $P = 0.7$). Thus, only the phasic outcome-yoked component of the pupillary response predicted participant behavior and correlated with trait anxiety.

DISCUSSION

In real life, a given action rarely results in the same outcome 100% of the time. Thus, it is important to be able to gauge how likely our actions are to produce good or bad outcomes, and to be able to know when to update our estimates of action-outcome contingencies in response to new experiences. Knowing whether to write off an unexpected aversive outcome as a chance event depends on whether we are in a quickly changing environment, where surprising outcomes should lead us to revise estimated contingencies rapidly, or a stable one, where unexpected outcomes are more likely to reflect noise in action-outcome relationships and we should revise our estimates more slowly⁶. Our results suggest that people are generally able to rationally adapt their learning about aversive outcomes on the basis of whether action-outcome associations are volatile or stable. However, individuals with high trait anxiety show an impoverished ability to do this. This difficulty in using information about the stability of action-outcome contingencies to correctly judge whether or not to repeat an action that has led to an unexpected aversive outcome may well lead high trait-anxious individuals to engage in poor decision-making. It might also result in aversive outcomes being experienced as less predictable and less avoidable. This could in turn lead to an increase in anxiety-related symptomatology, and potentially even be involved in the onset or maintenance of anxiety disorders.

Non-luminance-related changes in pupil size are argued to reflect, among other influences, the activity of central arousal systems, including the locus coeruleus–norepinephrine system^{15–17}. Activity of this system is held to be closely linked to the processing of environmental volatility and the use of this to guide learning^{6,14}. Consistent with this, we observed that participants showed a significantly greater increase in pupil dilation following trial outcomes in which environmental volatility was high. Furthermore, the degree to which participants' post-outcome pupil dilatory response tracked environmental volatility was significantly correlated with the degree to which they adjusted their learning rate between stable and volatile blocks. Of greatest pertinence to our study, trait anxiety modulated this pupil response to volatility, with high trait-anxious participants showing a smaller effect of volatility on post-outcome pupil dilation. Additional studies involving pharmacological manipulation are required to determine whether this reflects an anxiety-related deficit in norepinephrinergic control of the mechanisms that enable environmental volatility to guide action-outcome updating versus the effect of an upstream influence acting on the locus coeruleus–norepinephrine system.

Notably, elevated trait anxiety was not associated with an altered pupil response to outcome surprise. Trait anxiety also did not modulate behavioral responses to surprising outcomes, that is, there was no influence of trait anxiety on choice reaction time slowing on trials following surprising outcomes. These findings suggest that the reduced influence of environmental volatility on action-outcome contingency learning in high trait-anxious participants is unlikely to merely be a secondary consequence of a more general insensitivity to surprising outcomes. Rather, individuals with high levels of trait anxiety appear to respond normally to the experience of unexpected aversive outcomes, but are unable to utilize higher order (that is, block dependent) statistical information present in the distribution of these outcomes to modulate the updating of action-outcome contingency estimates. These findings also indicate that any anxiety-related deficit in the norepinephrinergic control of learning mechanisms would have to be fairly specific in nature.

A number of open questions remain. The first concerns whether anxiety is specifically linked to difficulty in using environmental volatility to update action-outcome estimates when possible outcomes are negatively valenced. A reward-based version of the task used here revealed no significant relationship between trait anxiety and use of the environmental volatility to update learning ($P > 0.1$, **Supplementary Fig. 2b**). However, although this raises the possibility that anxiety is associated with a specific difficulty in aversive learning, it is difficult to draw conclusive inferences from this null result, especially as the effect of anxiety on the change in learning rate between stable and volatile task blocks was in the same direction in both tasks and the difference in the strength of this relationship between tasks was not significant ($P > 0.1$, **Supplementary Fig. 2b**). In addition, the reward task used a secondary reinforcer (financial reward), whereas our aversive learning task adopted a primary reinforcer (electrical stimulation). Thus, the specificity of the anxiety-related deficit that we observed to aversive versus reward-based learning and to learning in the context of primary versus secondary reinforcers remains to be established.

As mentioned earlier, there is also an interesting parallel between the findings reported here and those from a recent classical fear conditioning study, where high trait-anxious individuals showed impoverished ability to differentially regulate conditioned fear responses across contexts that differed in stimulus-stimulus contingencies¹². In this prior study, high trait-anxious participants also appeared to be less able to adapt to the associative statistics of different contexts.

Future work should address whether there is indeed a common deficit in responding to changes in aversive stimulus-stimulus and action-outcome contingencies in high trait-anxious individuals, and whether this has a causal role in the anxiety experienced by these individuals as well as their elevated risk for developing full blown anxiety disorders. In addition, an important question is whether the anxiety-related deficit in contingency learning reported here, and potentially entailed in prior findings¹², is unique to anxiety or reflects a common deficit shared with risk for depression. This question reflects increasing awareness of the comorbidity between anxiety and depressive disorders and the need to identify which of those biological mechanisms identified as potentially conferring vulnerability to disease are unique to anxiety or depression versus common to both²⁴.

In conclusion, our findings reveal that trait vulnerability to anxiety is associated with a deficit in the use of higher order statistics about the causal structure of adverse environments to guide decision-making. High trait-anxious individuals did not differ from low trait-anxious individuals in their mean learning rate or in their behavioral or pupillary response to surprising adverse outcomes. Nor did they show altered preferences for minimizing shock probability versus shock magnitude. Instead, their pattern of decision-making was indicative of a selective difficulty with differentially updating action-outcome contingencies as a function of whether the current environment was stable or volatile. Our pupillometry data also confirmed a specific insensitivity to environmental volatility and raised the possibility that this might reflect impoverished modulation by environmental volatility of activity in the locus coeruleus-norepinephrine system.

In everyday life, determining whether, given the current context, an unexpected negative outcome is probably a chance event or something likely to occur again if the action that led to it is repeated may be essential to personal relationships and work place judgments. A deficit in this aspect of learning may have an important maintaining, or even etiological, role in the anxiety experienced by high trait-anxious individuals. We have taken a step toward elucidating this deficit, and hope to have illustrated how computational models can be integrated with behavioral and pupillometry analyses to begin to identify the mechanisms underlying disrupted decision making in high trait-anxious individuals.

METHODS

Methods and any associated references are available in the [online version of the paper](#).

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

The authors would like to thank S. Hicks and C. Kennard for use of the eye-tracking system. This research was supported by European Research Council grant GA 260932 and US National Institutes of Health grant R01MH091848.

AUTHORS CONTRIBUTIONS

M.B. and S.J.B. wrote the manuscript. M.B., S.J.B. and T.E.B. designed the task. M.B. collected the data. All of the authors contributed to data analysis. T.E.B. developed the Bayesian model. All of the authors commented on the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

- Butler, G. & Mathews, A. Cognitive processes in anxiety. *Adv. Behav. Res. Ther.* **5**, 51–62 (1983).
- Hartley, C.A. & Phelps, E.A. Anxiety and decision-making. *Biol. Psychiatry* **72**, 113–118 (2012).
- Paulus, M.P. & Yu, A.J. Emotion and decision-making: affect-driven belief systems in anxiety and depression. *Trends Cogn. Sci.* **16**, 476–483 (2012).
- Birrell, J., Meares, K., Wilkinson, A. & Freeston, M. Toward a definition of intolerance of uncertainty: a review of factor analytical studies of the Intolerance of Uncertainty Scale. *Clin. Psychol. Rev.* **31**, 1198–1208 (2011).
- Carleton, R.N., Norton, M.A. & Asmundson, G.J. Fearing the unknown: a short version of the Intolerance of Uncertainty Scale. *J. Anxiety Disord.* **21**, 105–117 (2007).
- Yu, A.J. & Dayan, P. Uncertainty, neuromodulation, and attention. *Neuron* **46**, 681–692 (2005).
- 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).
- Behrens, T.E., Woolrich, M.W., Walton, M.E. & Rushworth, M.F. Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–1221 (2007).
- Sutton, R. & Barto, A.G. *Reinforcement Learning* (MIT Press, 1998).
- Rescorla, R.A. & Wagner, A.R. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. in *Classical Conditioning II: Current Research and Theory* (eds. A.H. Black & W.F. Prokasy) 64–99 (Appleton-Century-Crofts, 1972).
- Behrens, T.E., Hunt, L.T., Woolrich, M.W. & Rushworth, M.F. Associative learning of social value. *Nature* **456**, 245–249 (2008).
- Indovina, I., Robbins, T.W., Nunez-Elizalde, A.O., Dunn, B.D. & Bishop, S.J. Fear-conditioning mechanisms associated with trait vulnerability to anxiety in humans. *Neuron* **69**, 563–571 (2011).
- Preuschoff, K., 't Hart, B.M. & Einhauser, W. Pupil dilation signals surprise: evidence for noradrenaline's role in decision making. *Front. Neurosci.* **5**, 115 (2011).
- Nassar, M.R. *et al.* Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* **15**, 1040–1046 (2012).
- Eldar, E., Cohen, J.D. & Niv, Y. The effects of neural gain on attention and learning. *Nat. Neurosci.* **16**, 1146–1153 (2013).
- Gilzenrat, M.S., Nieuwenhuis, S., Jepma, M. & Cohen, J.D. Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cogn. Affect. Behav. Neurosci.* **10**, 252–269 (2010).
- 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).
- Gorman, J.M., Hirschfeld, R.M. & Ninan, P.T. New developments in the neurobiological basis of anxiety disorders. *Psychopharmacol. Bull.* **36** (suppl. 2), 49–67 (2002).
- Charney, D.S. & Drevets, W.C. Neurobiological basis of anxiety disorders. in *Neuropsychopharmacology: the Fifth Generation of Progress* (eds. Davis, K.L., Charney, D.S., Coyle, J.T. & Nemeroff, C.) 901–930 (American College of Neuropsychopharmacology, 2002).
- Spielberger, C.D., Gorsuch, R.L. & Lushene, R.D. *Manual for the State-Trait Anxiety Inventory (STAI)* (Consulting Psychologists Press, 1983).
- Bestmann, S. *et al.* Influence of uncertainty and surprise on human corticospinal excitability during preparation for action. *Curr. Biol.* **18**, 775–780 (2008).
- MacKay, D.J.C. *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, 2003).
- 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**, 3660–3669 (2013).
- Casey, B.J. *et al.* DSM-5 and RDoC: progress in psychiatry research? *Nat. Rev. Neurosci.* **14**, 810–814 (2013).

ONLINE METHODS

Participants. 31 (22 female) English speaking, healthy participants, aged between 18 and 40 years (mean = 23.7, s.d. = 5.6), were recruited from the local community. One was a replacement for a prior participant where data was lost. The number of participants recruited was based on estimates of sample sizes from prior studies using similar tasks or studying effects of trait anxiety^{8,12,23}. Pre-screening on the trait subscale of the Spielberger State-Trait Anxiety Inventory (trait-STAI)²⁰ was conducted to achieve an approximately even number of participants with scores in the ranges 20–30, 30–40, 40–50 and >50. Exclusion criteria included current receipt of psychoactive medication or psychological therapy, neurological illness or suicidal symptomatology. Data from one participant was lost, as noted above, due to equipment failure. Analyzable eye tracking data was obtained on all but two of the remaining participants.

General procedure. The study was approved by the Oxford University Research Ethics Committee. Following provision of written informed consent, participants completed two learning tasks; the novel aversive learning task and a parallel reward learning task. The latter was included to establish replication of cross-participant effects of the volatility manipulation in a reward context (as previously reported⁸), the current focus of interest being the extension of this to learning in an aversive context. The reward task is not reported in detail here; we note that cross-participant effects of volatility upon learning were indeed replicated but no significant effects of trait anxiety were observed (Supplementary Fig. 2). The order of tasks was counterbalanced across participants. There was no effect of task order.

The tasks were presented on a cathode ray monitor connected to a computer running Presentation software version 15.1 (Neurobehavioral Systems). Participants' heads were stabilized using a head and chin rest placed 70 cm from the screen. An Eyelink 1000 system (SR Research) was mounted on the headrest and collected eye position and pupil size data from the right eye at a rate of 1,000 Hz. Electrical stimulation was delivered as trains of 2-ms pulses using a DS7AH constant current electrical stimulator (Digitimer). This was controlled by the stimulus presentation computer and connected to the volar surface of participants' non-dominant forearm using a bipolar electrode.

The aversive learning task. The aversive learning task (Fig. 1) was adapted from a structurally equivalent reward learning task previously successfully used to examine the impact of changes in environmental volatility upon rate of learning^{8,11}. It comprises two blocks of 90 trials in which participants choose between two stimuli that are probabilistically associated with receipt of an electrical shock (Fig. 1). In one block, choice-outcome contingencies are stable (shape A results in a shock 75% of the time, shape B 25% of the time). In the other block, choice-outcome contingencies are volatile, switching every 20 trials. The magnitude of shock received if a given shape is chosen and does result in a shock on that trial is specified separately for each shape and varies from trial to trial (more details are given below). These key details are important for the following reasons: (1) trial-wise manipulation of magnitude of shock leads to trial-wise changes in the relative expected value of each shape and enables all 90 trials of each block to contribute to estimation of the learning rate within that block (rather than just 'start' (of block) trials and 'switch' trials (those following changes in action-outcome contingency within the volatile block)); (2) this in turn allows us to maximize trial number within blocks as opposed to the number of separate blocks—this is critical to establish the extent to which participants are able to modulate their learning rate as a function of the volatility of action-outcome contingencies, as volatility can only be estimated across a large number of trials (see the **Supplementary Modeling Note** for an illustration of the rate at which information on volatility may be accrued).

Task details. On each trial, participants were presented with a fixation cross flanked by two shapes which were offset by approximately 7° visual angle. Participants' task was to choose one of the two shapes, one of which would result in the delivery of an electrical shock. In the stable block, one of the two shapes predicted the occurrence of the shock with a probability of 75%, the other shape resulting in shock delivery on the remaining trials. In the volatile block, the shape that most predicted the shock switched every 20 trials (with one shape predicting it with a probability of 80%, the other with a probability of 20%). The magnitude of shock that would be delivered if administered was specific to each

shape and displayed as a two digit number (between 01 and 99) in the shape's center. On each trial, the magnitude value for each shape was chosen from two separate random distributions. Following a variable period of time (drawn from a Poisson distribution, min = 1 s, mean = 2 s), the central fixation cross changed to a question mark, indicating that participants could respond. Participants had up to 4 s to choose one of the two shapes. After the participant responded, there was a jittered interval (min 2 s, mean 4 s) before outcome presentation. In the outcome phase of all trials, the shape resulting in the shock was shown at fixation, between the two options, for two seconds (Fig. 1). If participants had chosen that shape, a shock was delivered at the beginning of this period (the intensity of the shock being determined by the magnitude associated with the chosen shape). Following this, the shapes were replaced by crosses (a single cross at fixation, and two double crosses either side) and there was a further jittered interval (min 2 s, mean 4 s) before the next trial began. Participants completed 180 trials, 90 in each block. The two task blocks were completed sequentially with no break between them and with no explicit cuing as to the division of the task into two distinct blocks. Equal numbers of participants were randomly assigned to complete either the stable or volatile block first.

Stimuli. Visual stimuli were constructed to minimize variations in luminance during the course of the task. All stimuli were shaped Gabor patches, with an area of ~16° squared visual angle (that is, approximately 4° × 4°, matched on area across shape types), the same mean luminance as the background (83.9 ± 0.5 Cd m⁻²), and a spatial period of 0.5° visual angle. Each participant completed the task using either circle and square or triangle and cross shaped gabor patches.

Calibration of electrical shocks. Prior to completion of the learning task, a calibration procedure was conducted to equate subjective pain across participants, as far as possible. During calibration, pain was reported by participants using a 10 point scale on which 1 was defined as "minimal pain", 10 as "worst possible pain" and 7 as the worst pain that the participant could tolerate receiving up to 20 times during the task. The amplitude of a single 2-ms period of electrical stimulation was increased from zero until it produced a sensation rated as 1/10. The amplitude of a single 2-ms pulse was then kept at this level with the intensity of the shock gradually increased up to a subjective intensity of 7/10 by increasing the number of 2-ms pulses delivered in a train. This resulted in a single pulse producing 1/10 pain and a maximum number of pulses, delivered in a train, which produced 7/10 pain. Across participants, the maximum number of pulses used ranged from 8 to 48, with no relationship between this measure and trait-anxiety ($r(28) = 0.18, P = 0.3$). Participants subsequently completed 14 trials during which the intensity of electrical shock was randomly varied by changing the number of pulses delivered in a train between 1 and the number required to produce subjective pain of 7/10. Participants' subjective pain ratings of these different levels of shock were fitted to a sigmoid curve. This was used to determine the number of pulses given for each magnitude shock level (1–99) during the task (that is, a magnitude of '1' would receive a single 2-ms pulse, a magnitude of '99' would receive the number of pulses required to produce a subjective pain rating of '7', all other magnitude level / number of pulse conversions were calculated using the sigmoidal fit (Supplementary Fig. 1)).

Pupil dilation preprocessing. Pupil diameter measurements were collected using the Eyelink 1000 system as described above, at a sampling rate of 1,000 Hz. This data was cleaned with blinks removed using the Eyelink system's built in filter. Missing data points were linearly interpolated (trials in which more than 50% of the eyetracking data was interpolated were not used in subsequent analyses, mean = 16% of trials). The resulting trace was subjected to a low pass Butterworth filter with a cutoff of 3.75 Hz and then z transformed across the session¹⁴. Data were extracted from each trial using a window based on the presentation of the outcome. This included a 1-s baseline period before the presentation of the outcome, and a 6-s period following outcome presentation (this comprised the 2-s period when the shape associated with the shock was shown and the subsequent 4 s following removal of outcome information). Baseline correction was achieved by subtracting the mean pupil size during the baseline period from each time point post outcome.

An additional analysis of tonic baseline pupil diameter was also conducted. The pupil data submitted to this analysis was preprocessed in an identical manner to that described above, except for the absence of baseline correction¹⁶.

Decision choice data: estimating change in learning rate. The primary measure of interest was change in learning rate between stable and volatile task blocks. This was calculated by fitting a simple learning model, as previously described⁸, to participants' choice data. Details of the model, which consisted of a Rescorla-Wagner predictor¹⁰ coupled to a softmax based action selector, and the process by which the model parameters were fitted to participant choice data are provided in the **Supplementary Modeling Note**. In this note, we also provide evaluation of the model's performance (ability to account for participant choice) and formal comparison of its fit with that of alternative models.

Calculating trial-wise estimates of surprise and volatility, using a Bayesian learner, for use in reaction time (RT) and pupillometry analyses. Findings from previous decision-making studies suggest that human participants approximate an optimal Bayesian learner in their use of environmental statistics to implicitly guide decision choice. In particular, their choice behavior is consistent with use of trial-wise Bayesian estimates of environmental volatility to determine the extent to which to update estimates of outcome likelihood and to in turn use this to guide the action they take^{8,11}. Given this, we used a Bayesian learner (following⁸) to calculate trial-wise estimates of volatility (**Supplementary Modeling Note**). We sought to establish whether these estimates would predict trial-wise changes in the pupil response to outcome presentation. This prediction was informed by recent findings¹⁴ and made in light of the suggestion that the pupil dilatory response indexes activity in the locus coeruleus norepinephrine system^{15–17}, which has been argued to be sensitive to environmental uncertainty⁶.

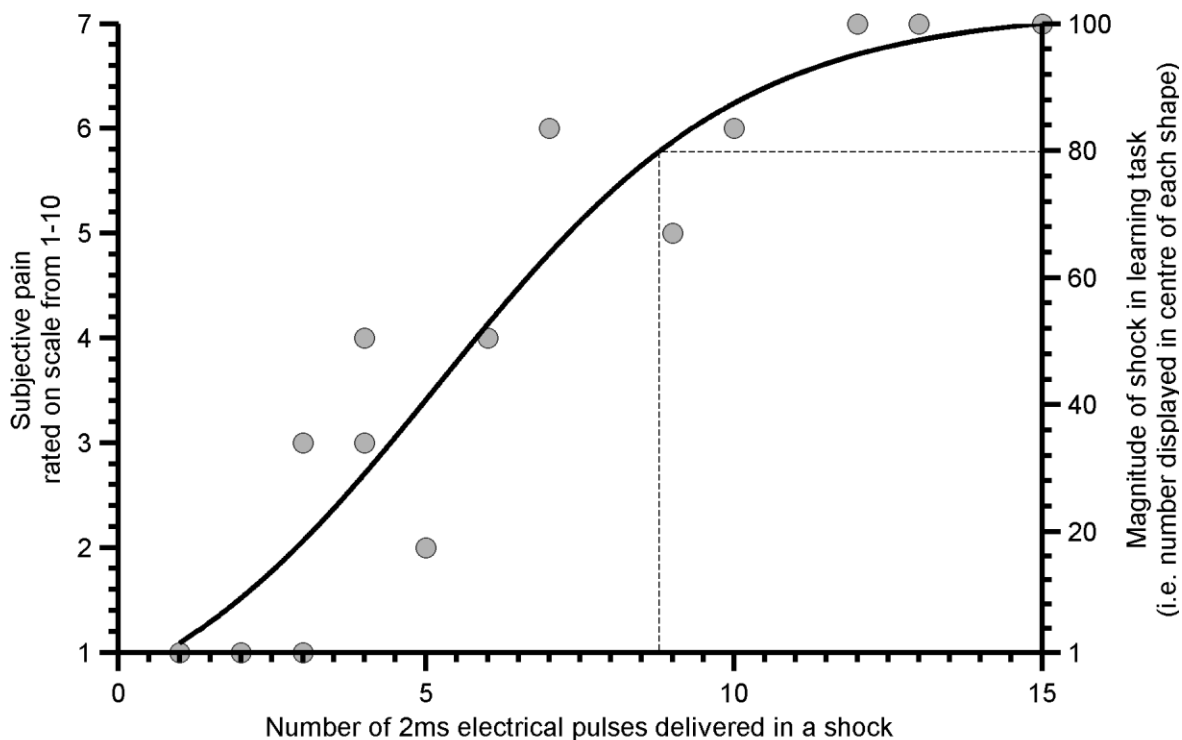
The extent to which an outcome is 'surprising' or unexpected is effectively an unsigned prediction error and has been linked to slowing on trials following those where the outcome is highly surprising²¹. The Bayesian learner was also used to calculate trial-wise estimates of surprise. These were estimated as the negative logarithm of the conditional probability of the outcome observed given the prior estimated probability of the outcome (that is, the Shannon information of the outcome²²; **Supplementary Modeling Note**). These trial-wise estimates of surprise were entered into regression analyses of pupil response to outcome presentation and also into regression analyses examining decision choice reaction time as a function of outcome surprise on the prior trial.

Decision choice reaction time analyses. For the reaction time analyses, data were excluded for trials where response times were so short (<200 ms) that they were unlikely to reflect a genuine choice (mean = 2% of trials). Reaction time data from all remaining trials were log transformed before being entered into regression analyses. The trial-wise measure of surprise on each trial n was regressed against the reaction time for trial $n+1$ on a participant-wise basis. Additional regressors of no interest were included to control for other task parameters that could influence reaction time. These comprised outcome (0, 1, where 1 = shock delivered), shock

magnitude (1–99), the pascalian (expected) values of the chosen and unchosen stimulus, and the Bayesian estimate of the volatility of the current trial. These participant-wise regression analyses output a beta weight for each participant which indicated the extent to which their choice reaction time on trial $n + 1$ was influenced by how surprising the previous outcome (on trial n) had been. These beta estimates were entered into a second-level one sample t test (to examine effects of outcome surprise on subsequent trial choice reaction time across participants) and a second-level Pearson product moment correlational analysis (to examine the modulation by trait anxiety of the effect of outcome surprise upon subsequent trial reaction time). For these and all other analyses reported in the paper, before parametric tests were conducted, the Kolmogorov-Smirnov test was used to ensure the data was distributed normally. All statistical tests reported were two-tailed unless specified otherwise.

Pupil dilation analyses. Preprocessing of the pupil data is described above, together with procedures for baseline correction. Preliminary checks were conducted to ensure that the main analyses were unlikely to be confounded by anxiety-related differences in fixation or saccades. These checks revealed that there was no significant relationship between trait anxiety and the number of saccades (P values > 0.25), or the time spent in fixation (P values > 0.35) in any of the 1-s time bins of the outcome period. Further, there was no relationship between these parameters and the degree to which participants altered their learning rate between volatile and stable task blocks (number of saccades: P values > 0.18; time spent in fixation: P values > 0.12).

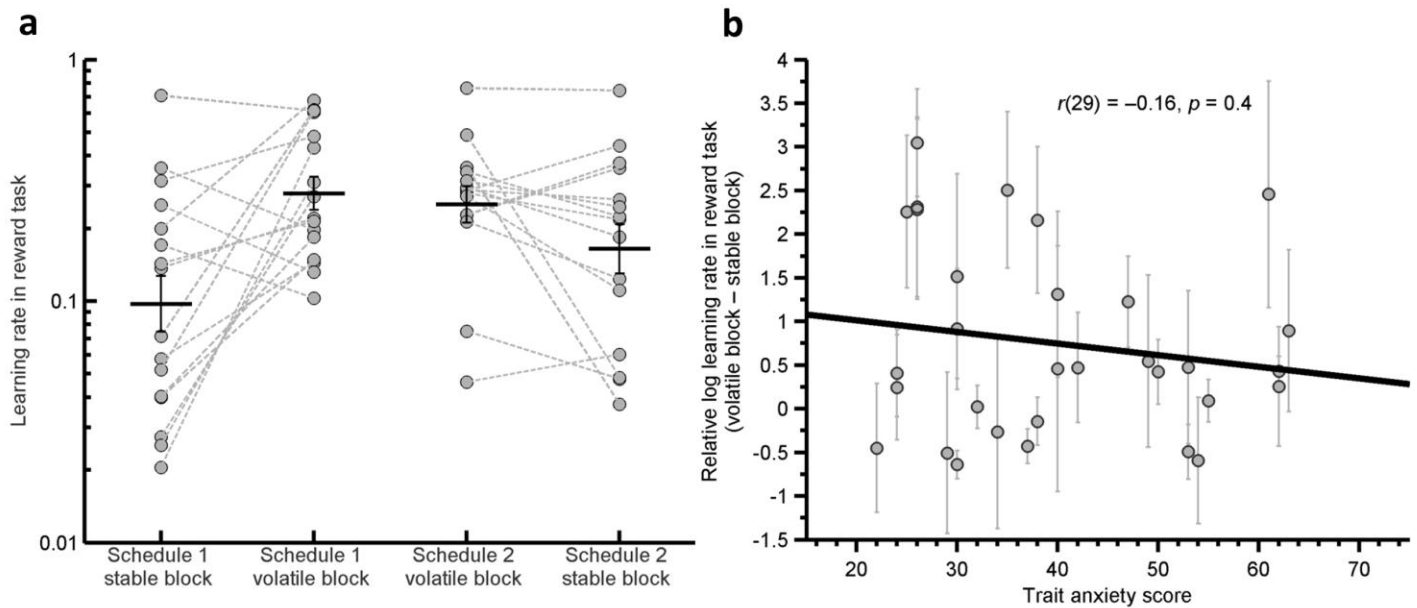
Regression analyses were conducted to examine the effects of trial-wise volatility and outcome surprise on pupil dilation following outcome presentation. The post-outcome period was sampled using 6,000 1-ms bins. Regression analyses were conducted for each of these bins, with Bayesian estimates of trial-wise volatility and surprise entered as regressors of interest; outcome (0, 1, where 1 = shock delivered), shock magnitude (1–99), and the pascalian values of the chosen and unchosen stimulus were entered as control regressors. The resultant timeseries of beta-weights for volatility and surprise were down-sampled to give beta estimates of the effects of surprise and volatility on pupil dilation for six sequential one second time bins across the outcome period. These beta values were entered as the dependent measure into analyses of variance (ANOVAs), one for surprise and one for volatility. Time bin was entered as a within-subject factor and block order (stable versus volatile block first) as a between subject variable. Cross time bin summary measures of pupil responsivity to trial volatility and outcome surprise were also calculated. These were used for correlational analyses against both behavioral indices of interest and our between subject measure of trait anxiety. Additional graphical representation of the effects of trait anxiety was achieved by dividing participants into two groups using a median split on trait anxiety scores, (**Fig. 4**, and **Supplementary Figs. 8 and 9**).



Supplementary Figure 1

Calibration of Electrical Shocks.

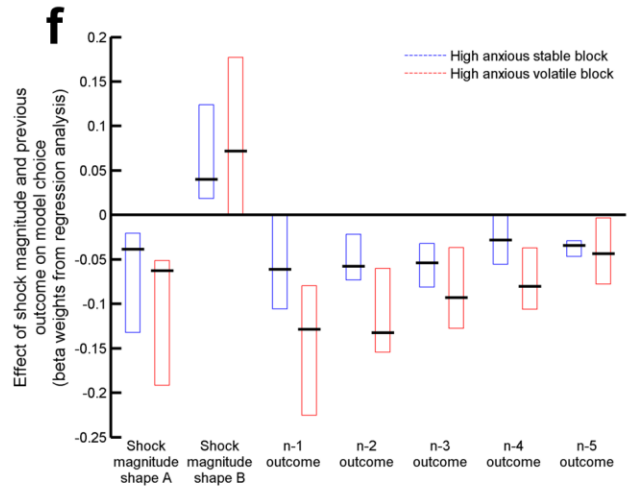
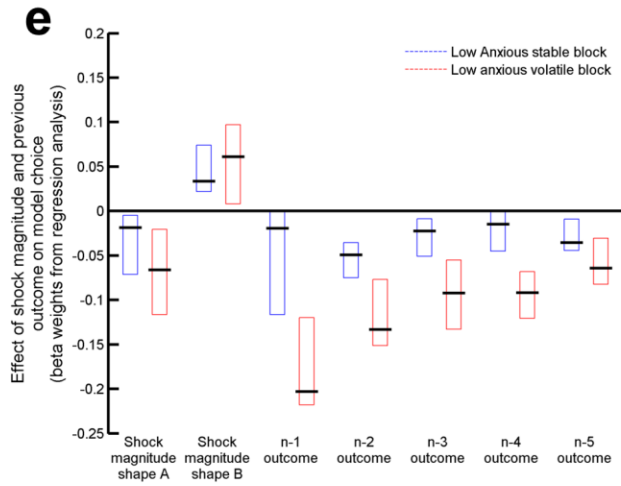
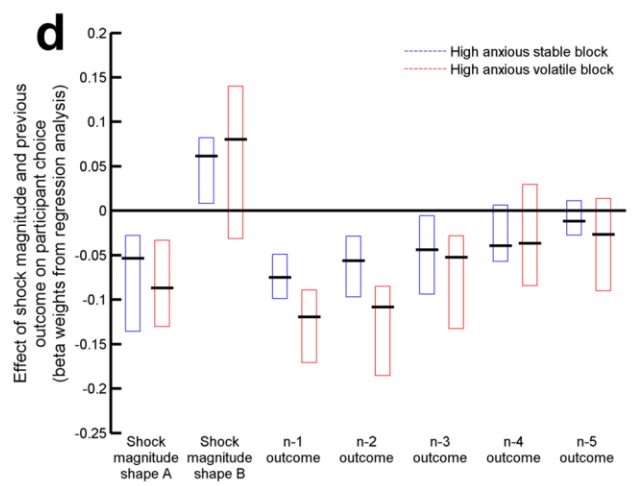
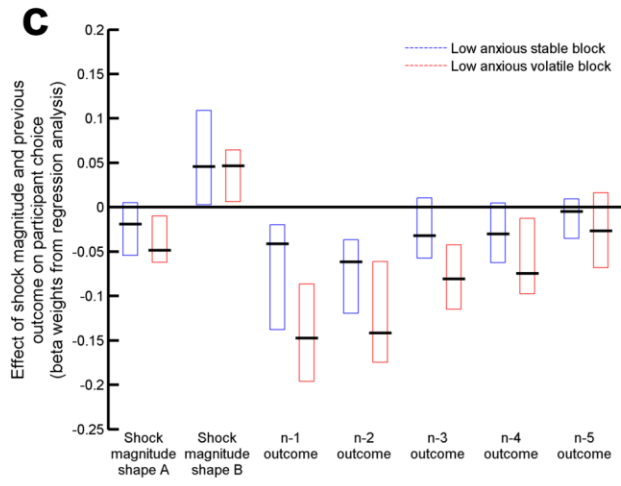
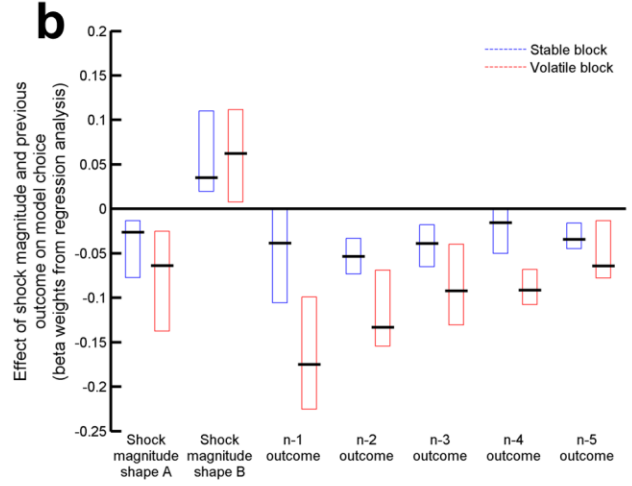
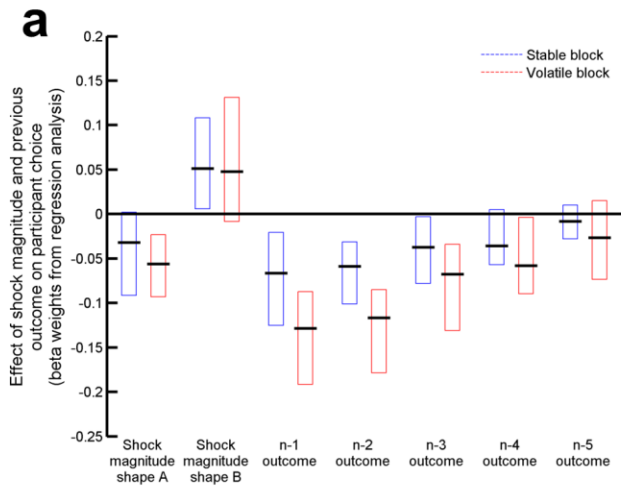
Before each participant completed the aversive learning task, subjective pain scores (left hand axis) were collected for different objective strengths of electrical shock (x axis), see online methods. A 10 point scale was used for subjective pain where 1=minimal pain, 10=worst possible pain. No shock that exceeded a subjective level of 7 was applied. A sigmoid curve fitted to this data was used to equate the subjective level of pain experienced during the learning task across participants. The graph shows data points (circles) and the fitted sigmoid (black line) from one participant. In the task, the shock magnitude associated with each shape was varied on a trial-to-trial basis from 1-99, where 1 equated to a subjective pain level of 1, and 99 to a subjective pain level of 7. Magnitude values between these extremes were calculated using the sigmoidal fit. This is illustrated here by the dashed line, which shows how the objective strength of the electrical stimulus (x axis) corresponding to a desired 'magnitude' of 80 (right hand axis) would be calculated for this participant.



Supplementary Figure 2

Effects of volatility upon participants' learning rates in a structurally equivalent reward task.

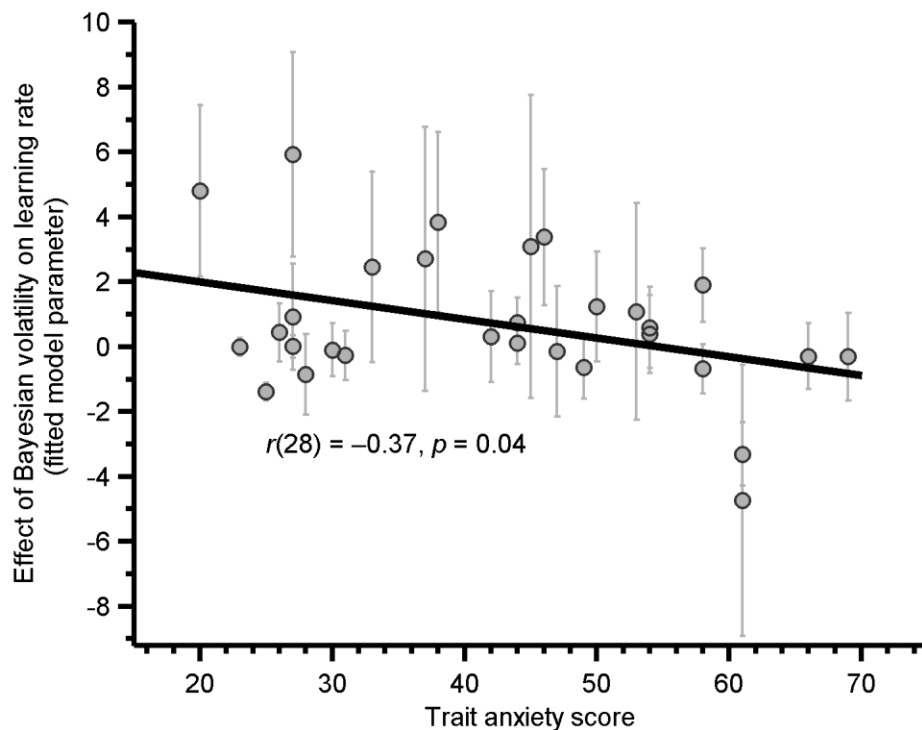
In addition to the aversive learning task, reported within the main manuscript, participants also completed a structurally equivalent reward learning task. Here the two shapes chosen between had different probabilities of leading to winning points, with the magnitude of reward points being unique for each shape and varying randomly between 1 and 99 across trials. Following the task the points participants had won were converted into a monetary reward (either £5 or £10). **(a)** Participants' choices during the stable and volatile blocks of the reward task were fitted with the same Rescorla Wagner learning model used in the aversive task. Estimates of individual participants' learning rates are displayed (circles), using a logarithmic scale, separately for the stable and volatile blocks for the two task schedules (Schedule 1 = stable task block first, $n = 16$, Schedule 2 = volatile task block first, $n = 15$). Black lines display mean (\pm SEM) of participants' estimated learning rates, grey dotted lines link the learning rates in volatile and stable blocks for each participant. Participants showed higher learning rates in the volatile versus stable blocks regardless of the order in which they were completed, $F(1,29) = 15.3, p = 0.001$, replicating previous results⁸. **(b)** The relative log learning rate for the volatile versus the stable blocks (i.e. $\log(\text{LR in volatile block}) - \log(\text{LR in stable block})$) was not significantly correlated with participant trait anxiety, $r(29) = -0.16, p = 0.4$. It should be noted that the difference in the effect of anxiety on learning rate for this reward learning task and the aversive learning task reported in the main manuscript was also not significant, $F(1,28) = 0.57, p = 0.46$. This hence limits the conclusions that can be drawn regarding the specificity of the anxiety-related deficit in adjusting learning rate to cases where outcomes are aversive as opposed to reward-related. Error bars represent the standard deviation of the estimated parameters from the behavioral model for each participant.



Supplementary Figure 3

Influence of potential shock magnitude and prior outcome history on actual and simulated behavior choice as a function of block volatility (a,b) and participant anxiety (c-f).

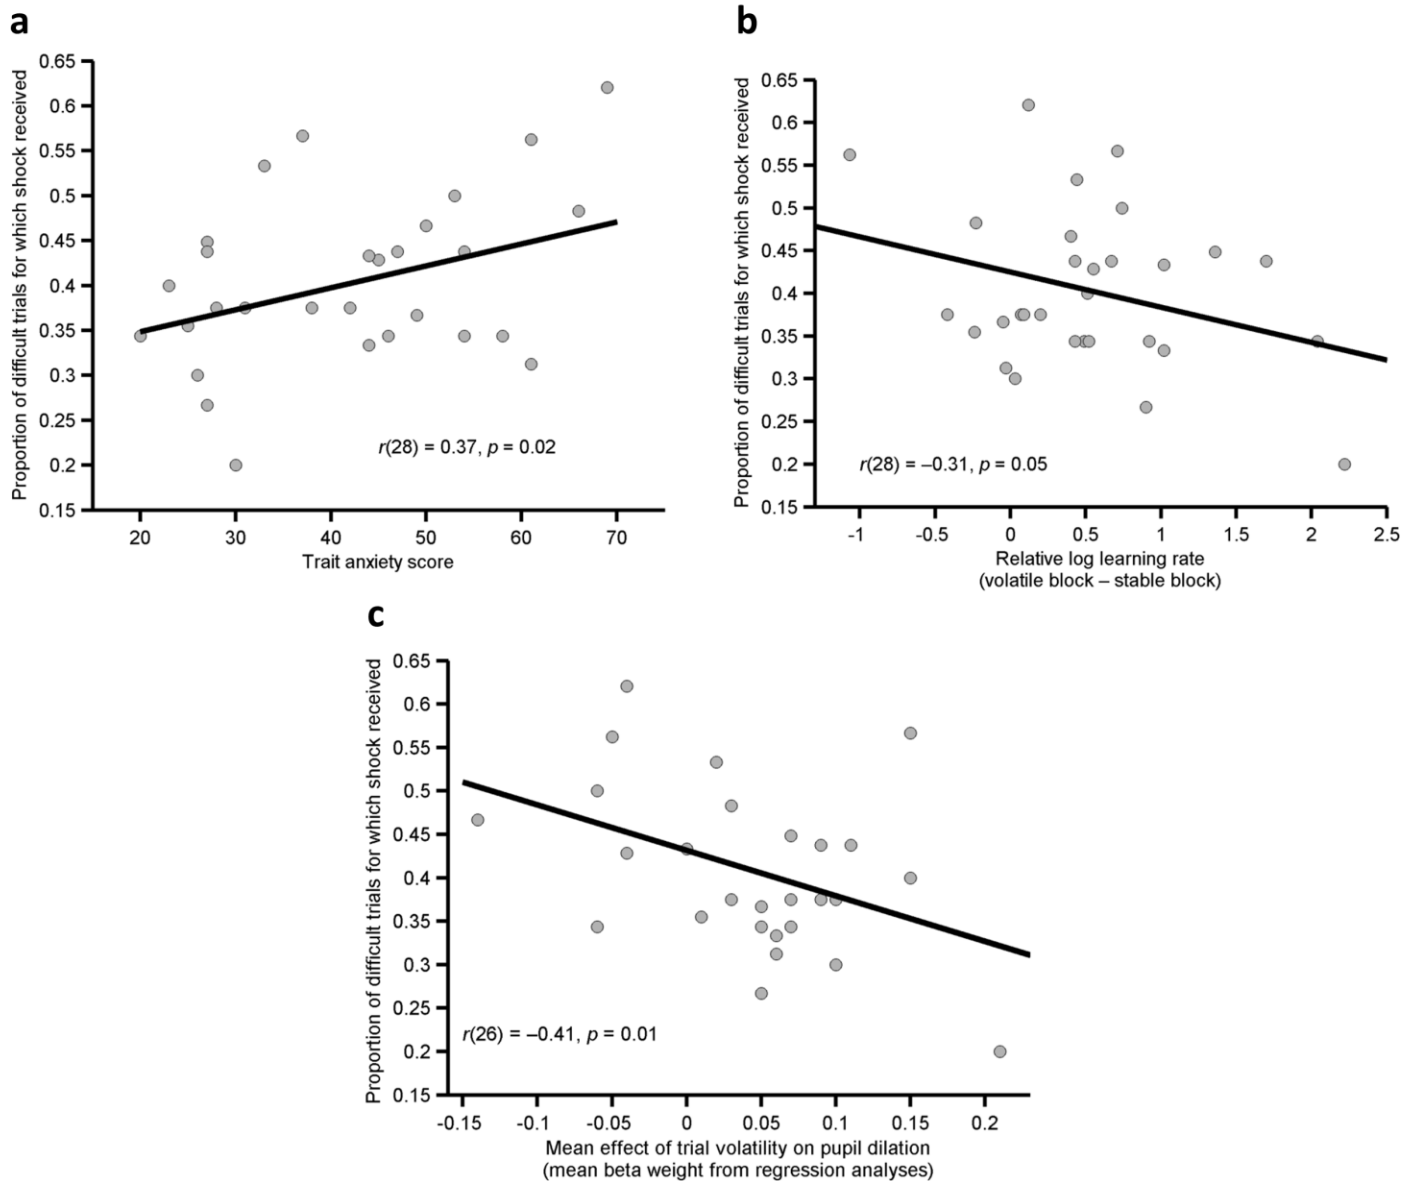
Regression analyses were conducted on a participant by participant basis to examine the extent to which current shock magnitude linked to each shape and the outcomes of the five previous trials predicted choice of shape A versus shape B (e.g. circle versus square) on a given trial. The dependent variable was participant choice (**a,c,d**) or model choice (**b,e,f**) on trial 'n' (coded as 1 if the subject/model chose shape A (e.g. circle) and 0 if shape B (e.g. square)). For the model choice regressions, the coupled Rescorla Wagner (RW) predictor and softmax action selector learning model was used to simulate choice behavior on a trial to trial basis, using the blockwise estimates of the free parameters (for learning rate, decision temperature and inverse temperature) previously calculated for each participant (see online methods). In each regression analysis, predictor variables comprised shock magnitude for shape A, shock magnitude for shape B, and outcome for trial n-1 to n-5 (coded as 1 if shape A was associated with the shock). As the distribution of parameter estimates across participants was non-normal, median values (horizontal black lines) and interquartile ranges (boxes) are presented. **(a)** Influence of the predictor variables on subject choice, estimated separately for trials in the stable (blue boxes) and volatile (red boxes) blocks. The impact of previous outcomes decreases across time and - consistent with the higher learning rate used by participants in volatile blocks - the impact of prior outcomes on shape choice is greater for more recent trial outcomes (n-1, n-2) in the volatile versus the stable block. **(b)** A parallel analysis was run on the model-derived choices. The similar pattern in **(b)** as for **(a)** is as expected given the good fit of the RW & softmax action selector model to the behavior data (see Supplementary Modeling Note). Plots **c** and **d** illustrate how the influence of prior outcomes on behavioral choice as a function of block volatility differs between low and high trait anxious individuals (as defined using a median split, for ease of graphical illustration). As expected given the inverse relationship between trait anxiety and change in learning rate between stable and volatile blocks (**Fig 2b**), a greater difference between blocks in the effect of recent prior outcomes on choice behavior is observed for low (**c**) than for high (**d**) anxious participants. Plots **e** and **f** illustrate that a parallel pattern is observed when using simulated data from the RW and softmax action selector model. In summary, panels **a**, **c** and **d** illustrate, in a model-free manner, the key behavior of participants assessed by the learning rate analyses reported within the main manuscript. Additionally, the plots shown in panels **b**, **e** and **f** indicate that the coupled RW predictor and softmax action selector is able to reliably capture these behaviors.



Supplementary Figure 4

Relationship between trait anxiety and Bayesian volatility when multiple parameters, including a decay function, are allowed to compete for influence over a dynamic learning rate.

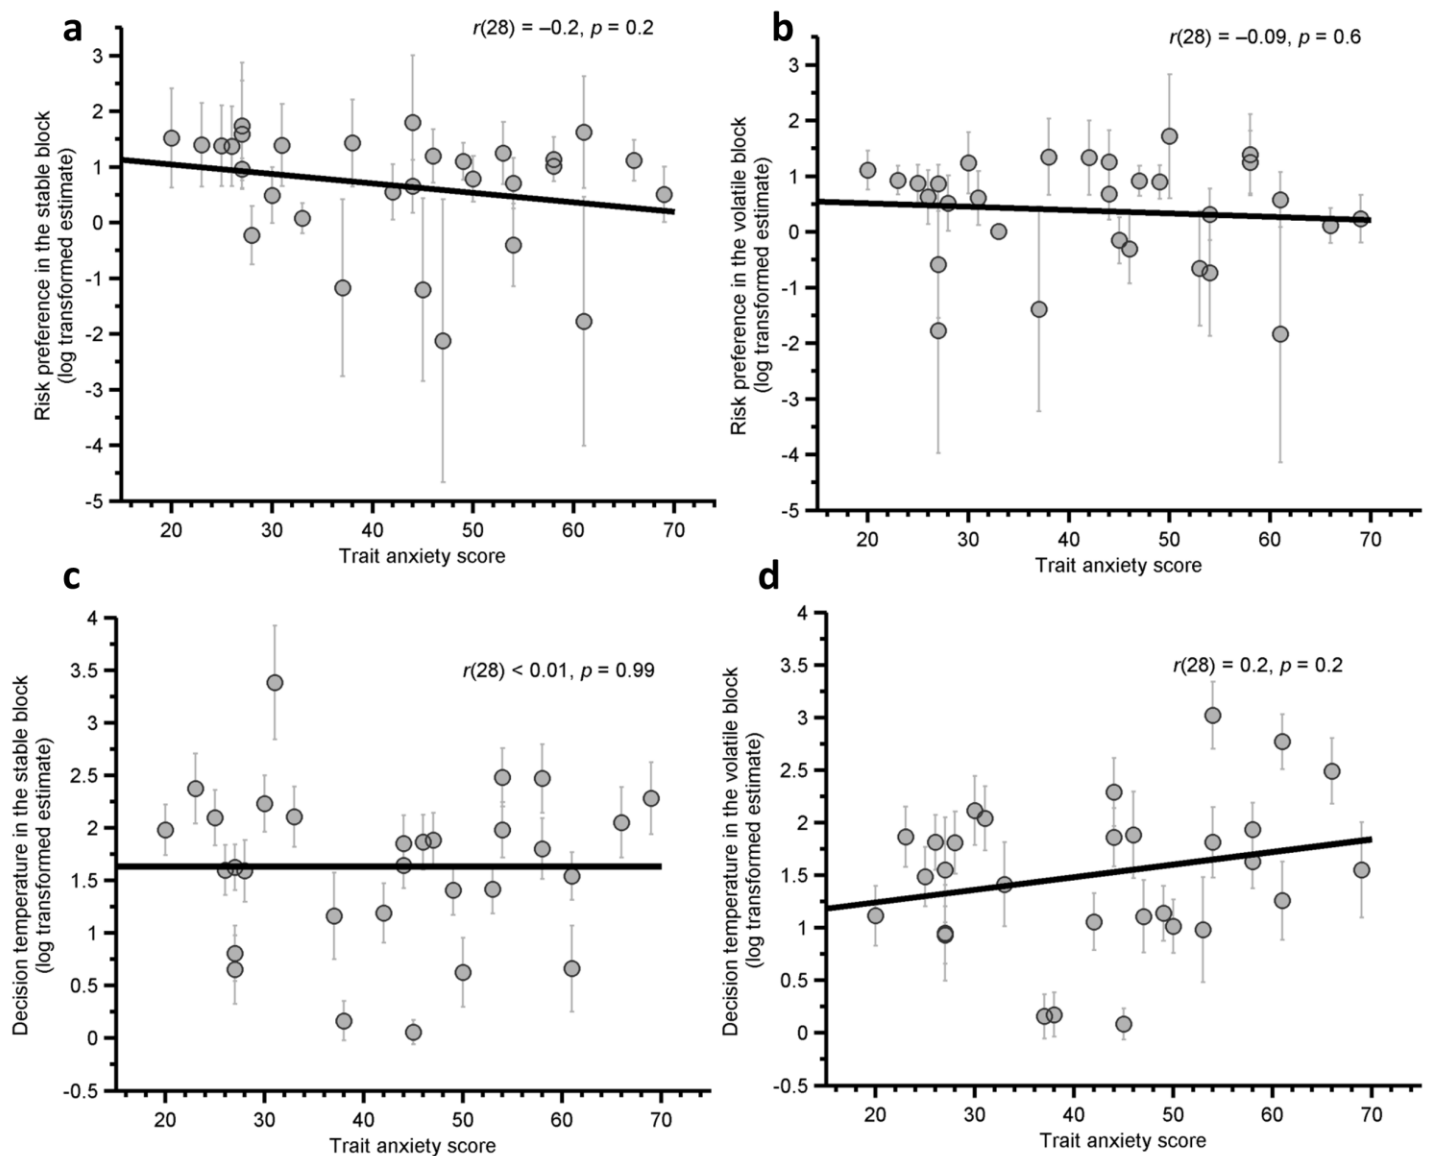
As described in the Supplementary Modeling Note, a more complex model, designed to capture a variety of potential influences on learning rate, was additionally fitted to choice data at the level of the individual participant. On each trial, learning rate was estimated as a weighted sum of three terms: the mean learning rate across trials (α_{mean}), the trial-wise demeaned Bayesian volatility (subject specific weight = ω_{Bayes}) and the trial-wise demeaned exponential decay function (subject specific weight = ω_{Exp}). These parameters along with the slope of the exponential function (λ) and the decision temperature and risk preference terms from the selector model were fitted to participant behavior across the entire task. The resulting parameter estimates were then correlated, across participants, against trait anxiety. Of the six parameters fitted to participants' behavior, only the weight of the volatility term (ω_{Bayes}) was significantly correlated with trait anxiety, illustrated above $r(28) = -0.37, p = 0.04$, with none of the other parameters showing a significant relationship with trait anxiety, $p > 0.2$. Consistent with the findings of the main paper, this result indicates that, even when a variety of factors were allowed to influence a dynamic learning rate, trait anxiety was uniquely associated with reduced influence of environmental volatility upon learning rate. Error bars represent the standard deviation of the parameter estimates for each subject.



Supplementary Figure 5

Shock receipt on difficult trials (where the two options were close in expected value) was positively correlated with trait anxiety (**a**), and negatively correlated with both change in learning rate between stable and volatile blocks (**b**) and modulation of post-outcome pupil dilation by trial volatility (**c**).

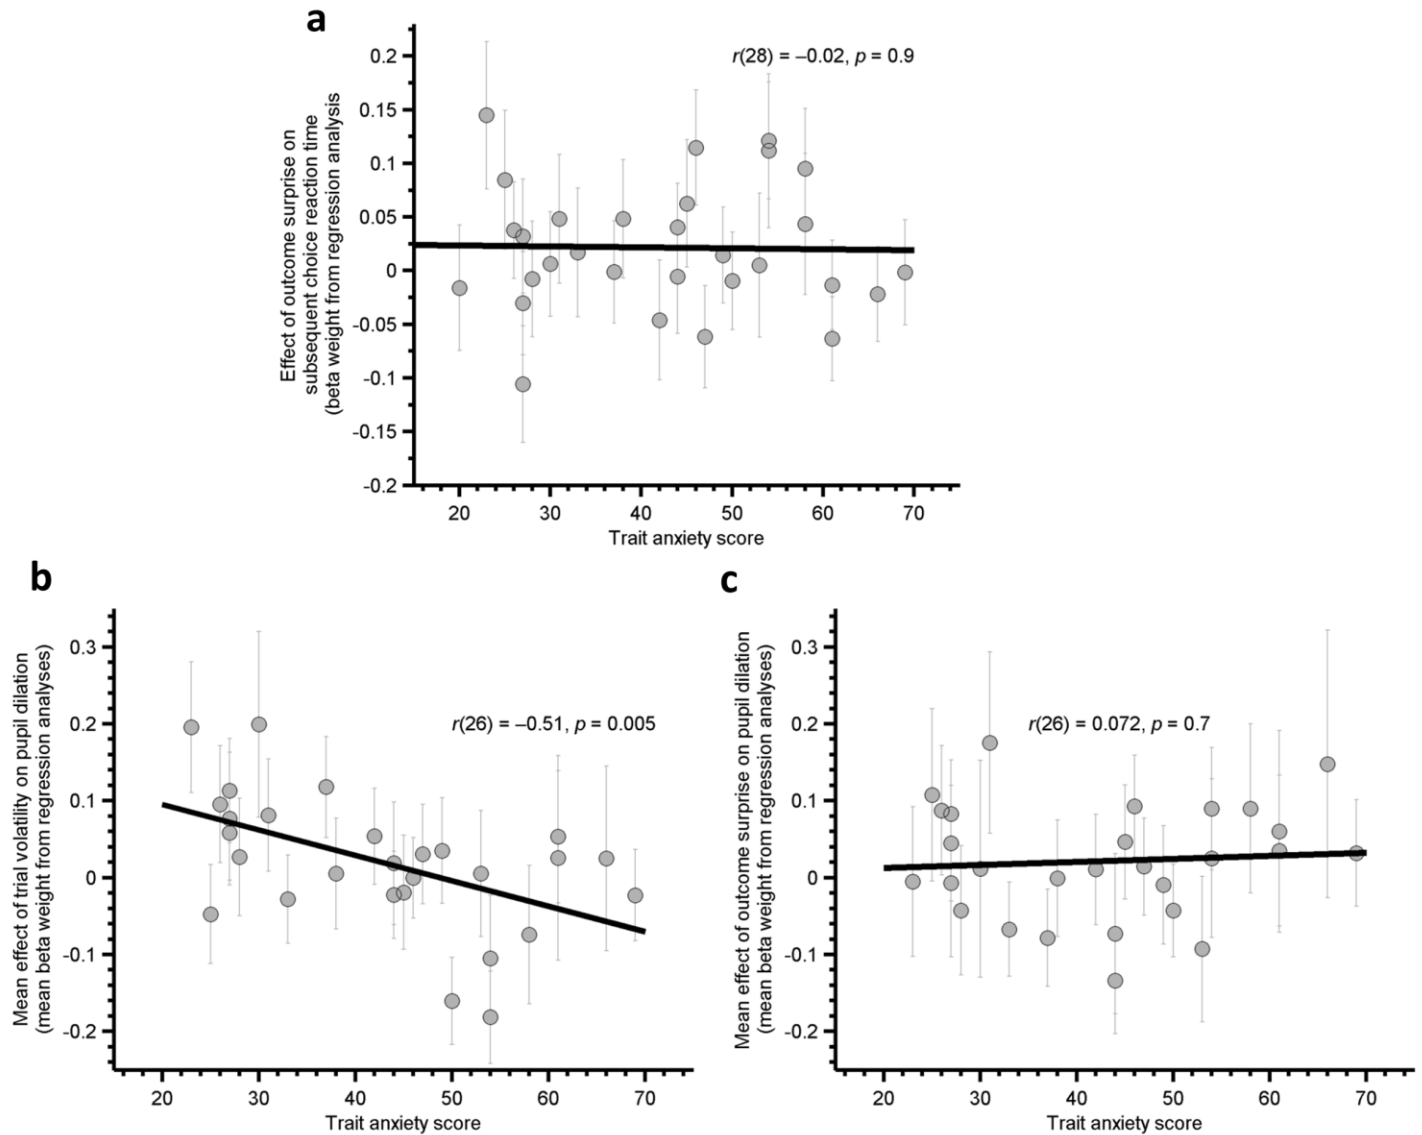
The advantage of flexible adaptation of learning rate to changes in environmental volatility is that it allows learning to occur at an optimal rate for a given environment. As subjects with higher levels of trait anxiety were less able to adjust their learning between stable and volatile blocks, they would be predicted to more often incorrectly judge the action most likely to result in shock. This in turn would be most likely to have an observable effect on trials where the two options were close in expected value. Following Behrens et al⁸, we defined 'difficult trials' as those in which the Pascalian value of the two shapes differed by less than 5. (**a**) Trait anxiety was positively correlated with the proportion of 'difficult' trials on which shock was received. (**b**) We further confirmed that this 'proportion of difficult trials where shock was received' index of performance was negatively related to adjustment of learning rate between stable and volatile task blocks, $r(28) = -0.31, p = 0.05$, 1-tailed (**b**), as well as to the pupil response to trial volatility, $r(26) = -0.41, p = 0.01$, 1-tailed (**c**).



Supplementary Figure 6

No significant association was observed between Trait Anxiety and Risk Preference or Inverse Decision Temperature.

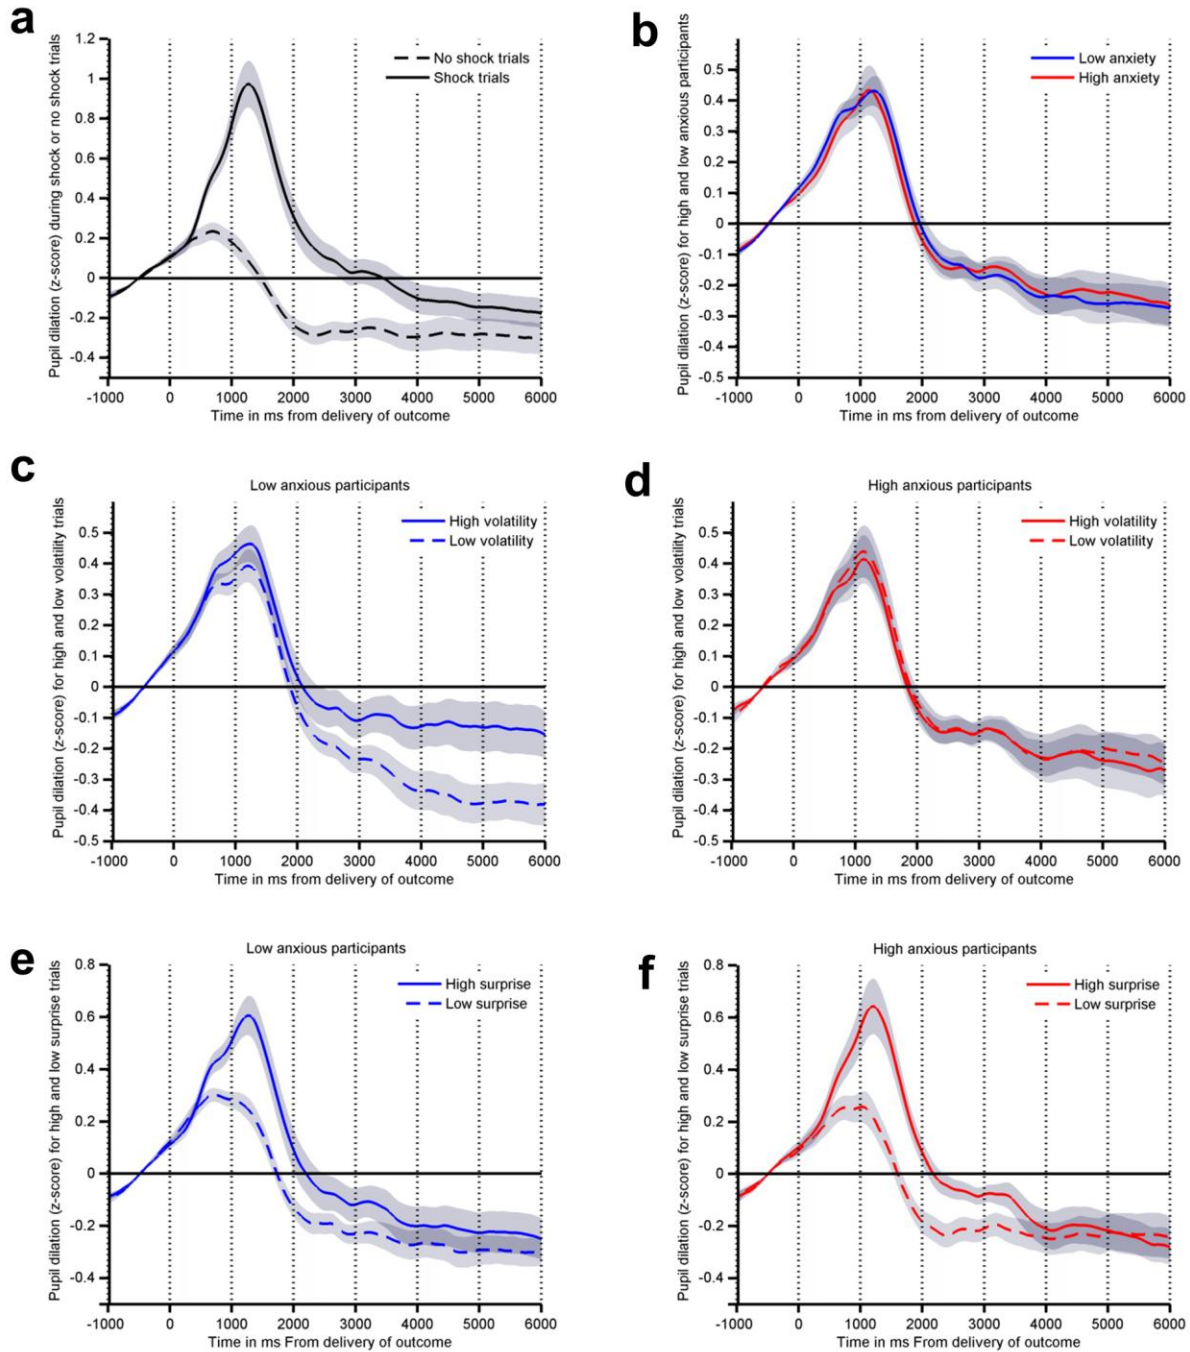
The softmax action selector contained two free parameters: risk preference (the tendency to minimize shock probability versus shock magnitude) and inverse decision temperature (the degree to which a participant used expected values, i.e. estimated shock probability \times shock magnitude, for each shape to guide choices). These parameters were estimated from participants' choices in the stable and volatile blocks of the aversive learning task (see Online Methods). Panels **a-d** show the relationship between the log transformed parameter estimates and trait anxiety. There was no significant relationship between trait anxiety and risk preference as calculated for either the stable (**a**; $r(28) = -0.23, p = 0.2$) or volatile (**b**; $r(28) = -0.09, p = 0.6$) task blocks. Trait anxiety was also not associated with the mean value of this parameter across blocks or with the difference score between blocks (p s $> .1$). There was also no relationship between trait anxiety and inverse decision temperature as estimated for the stable (**c**; $r(28) < 0.01, p = 0.99$) or volatile (**d**; $r(28) = 0.24, p = 0.2$) blocks. Trait anxiety was also not associated with mean decision temperature across blocks, or with the difference in decision temperature between blocks (p s $> .1$). Error bars represent the standard deviation of the parameter estimates for each subject.



Supplementary Figure 7

Analyses of reaction time and pupil dilation data using non-Bayesian estimates of trial volatility and outcome surprise

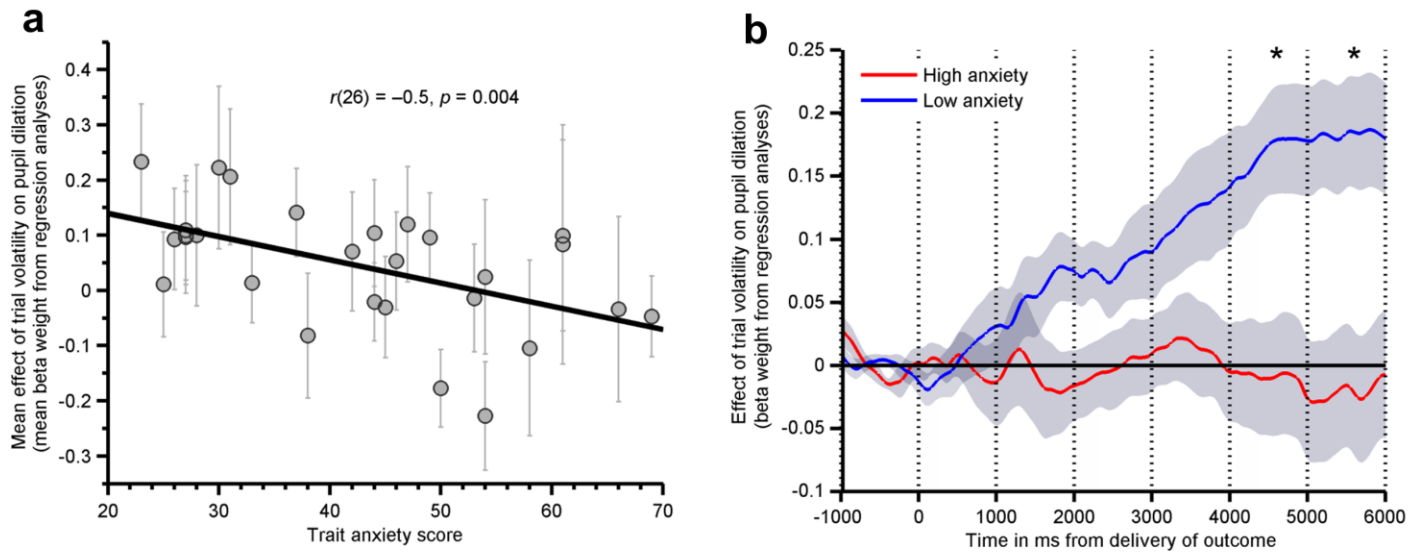
Analyses of the pupil dilation and reaction time data reported in this paper utilized estimates of trial volatility and outcome surprise derived from a Bayesian learner (see Supplementary Modeling Note). Here, we replicated these analyses using non-Bayesian estimates of volatility and surprise. Volatility was simply coded as '1' for all trials in the volatile block and '0' for all trials in the stable block. In addition, a non-Bayesian surprise regressor was created by coding all trials in which the less predictive stimulus was associated with the shock as being high surprise ('1') and the other trials as low surprise ('0'). These measures of volatility and surprise were entered into the regression models predicting reaction time slowing (**a**) and pupil response (**b,c**) together with the other, control, predictor variables previously used. As in the main analyses, the beta estimates for these regressors were correlated against trait anxiety. The results of these additional analyses replicated those reported in the paper: trait anxiety did not influence the degree to which participants slowed their response following a surprising outcome, $r(28) = -0.02, p = 0.9$ (**a**). Further trait anxiety was associated with a significantly reduced pupil response during high volatility trials, $r(26) = -0.51, p = 0.005$ (**b**), but showed no significant relationship with the pupil response to surprising outcomes, $r(26) = 0.07, p = 0.7$ (**c**). In summary, the relationship reported between trait anxiety and effects of volatility and outcome surprise held regardless of the manner in which these indices were calculated. Error bars represent the standard deviations of the parameter estimates from the behavioral model (**a**) and the regression coefficients (beta weights) from the pupil analysis (**b, c**) for each subject.



Supplementary Figure 8

Baseline corrected, z-transformed eyetracker traces from the aversive learning task.

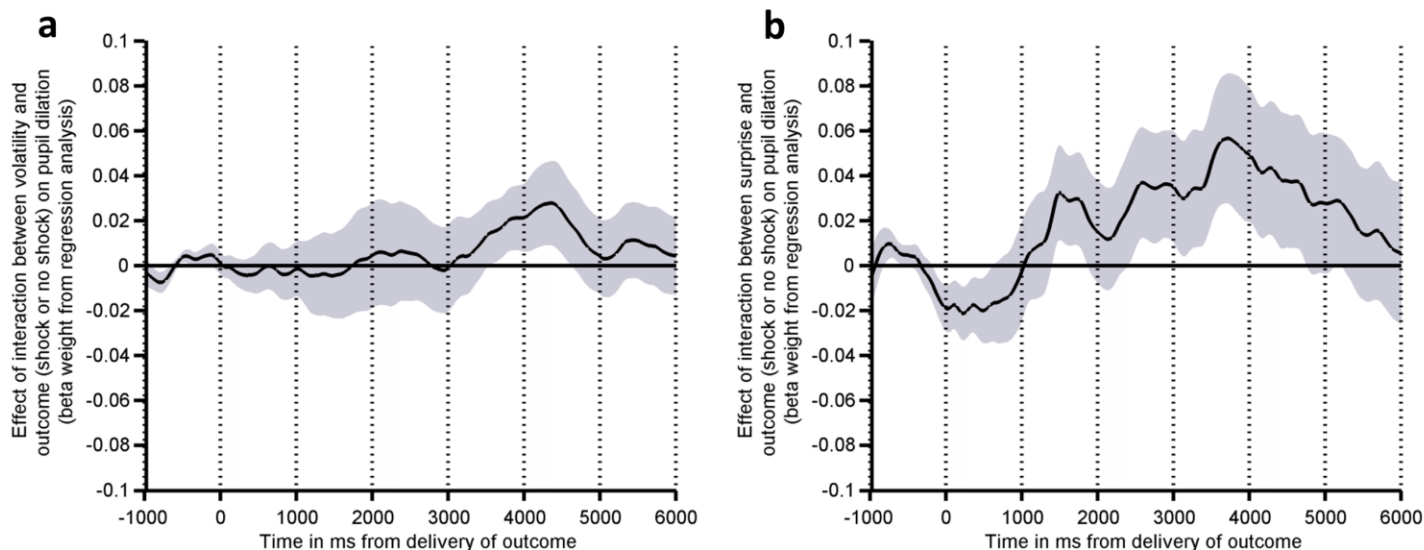
(a). Mean pupil dilation across participants separated into trials in which shock was administered or no shock was administered. (b). Mean pupil dilation across trials shown separately for high and low trait anxious participants. Effect of volatility on pupil dilation shown separately for low (c) and high (d) trait anxious participants. Effect of surprise on pupil dilation shown separately for low (e) and high (f) trait anxious participants. NB for illustration, participants have been separated into high and low anxious groups based on a median split on trait anxiety scores. Similarly trials were classified as high/low volatility or surprise based on a median split on these variables. Lines show mean value (across participants) of z-transformed, baseline corrected pupil diameter. Shaded regions represent SEM.



Supplementary Figure 9

The relationship between trait anxiety and post outcome pupil dilation as a function of trial-wise estimates of volatility, controlling for trial number within each block.

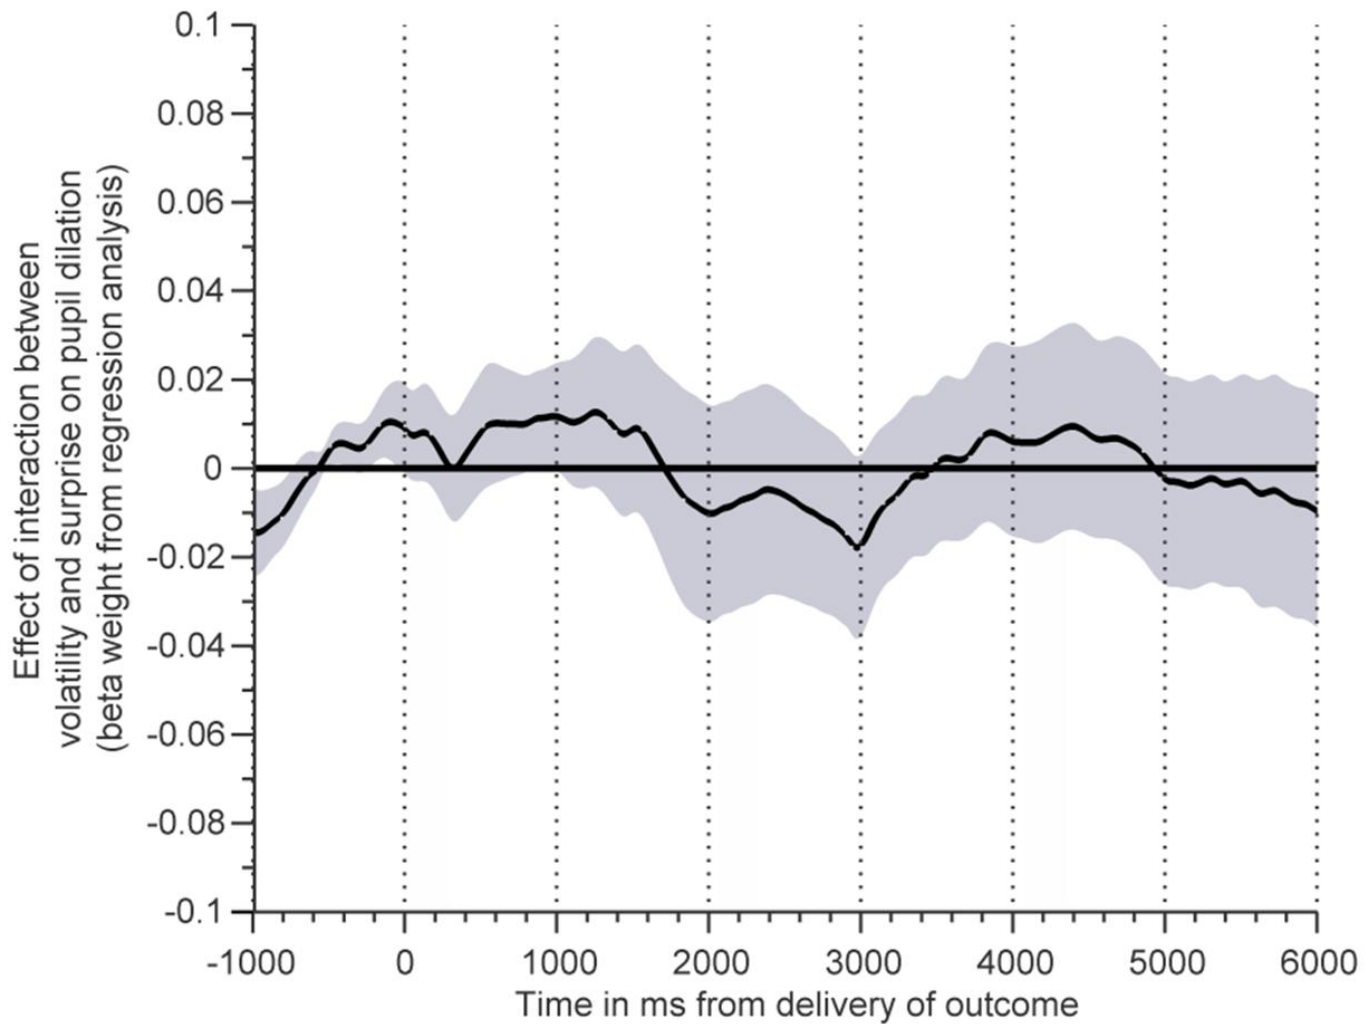
This figure illustrates the same analyses as those described in **Figure 4** of the main manuscript with the exception that regressors coding for the trial number within each task block have been included in the participant-wise regression analyses of the effect of volatility on post outcome pupil dilation. The results illustrated in **Figure 4** remain robust with the inclusion of these extra control parameters. Specifically, a strong negative correlation between trait anxiety and the effect of volatility upon pupil dilation post outcome is still observed (**a**), $r(26) = -0.5, p = 0.004$. As before, this is driven by low but not high trait anxious participants showing a modulatory effect of volatility on the post outcome pupil response (**b**), a median split on trait anxiety levels is used here as in **Figure 4**. Asterisks indicate 1s time bins in which Bonferonni corrected t-tests differed between the groups at $p < .05$ corrected, 2-tailed. Error bars in panel **a** represent the standard deviations of the regression coefficients (beta weights) from the pupil analysis for each subject. Shaded regions in panel **b** represent the standard error of the mean.



Supplementary Figure 10

Results from additional pupil analyses investigating the interaction between outcome (shock received versus no shock received), and volatility (a) and surprise (b).

The delivery of a shock causes a large dilation of participants' pupils (**Figure S8a**). In order to determine whether effects of volatility and surprise were modulated by outcome (shock received versus no shock received), we reran the pupil regression analyses described in the main manuscript with additional regressors encoding the interaction between volatility and outcome and the interaction between surprise and outcome. Neither interaction term had a significant influence on pupil dilation (t-tests of mean regressor values against 0, $p > 0.1$ for all time bins). Consistent with this, adding these terms to the model reduced the goodness of fit of the model by an adjusted R^2 of 0.0002. Shaded areas represent the standard error of the mean.



Supplementary Figure 11

Interaction of estimated trialwise volatility and surprise on pupil dilation post outcome.

The pupil analysis was rerun with the inclusion of a surprise x volatility interaction term. There was no significant interaction of surprise x volatility upon pupil dilation in any of the 1s post-outcome time bins (t-test of mean regressor value against 0, $p > 0.3$ for all time bins). While the post-outcome main effect of surprise on pupil dilation precedes that of volatility (**Figure 3**), trialwise volatility changes very slowly (see Supplementary Modeling Note) – i.e. the volatility of a given trial n and that of the following trial $n+1$ is highly correlated ($r > .9$). Hence this regression analysis should capture any effect of volatility on trial n on the response to surprise in trial $n+1$. Shaded areas represent standard error of the mean.

Supplementary Modeling note

Decision choice data: estimating change in learning rate.

The primary measure of interest in this paper was change in learning rate between stable and volatile task blocks. This was calculated by fitting a simple learning model⁸ to participants' choice data which consisted of a Rescorla-Wagner predictor¹⁰ coupled to a softmax based action selector. The predictor updated its estimate of outcome probabilities using the equation:

$$r_{(i+1)} = r_{(i)} + \alpha \varepsilon_{(i)}$$

Here $r_{(i+1)}$ is the estimated outcome probability for the $i + 1^{\text{st}}$ trial, $r_{(i)}$ is the estimated outcome probability for the i^{th} trial, α is the learning rate and $\varepsilon_{(i)}$ is the prediction error on the i^{th} trial.

The selector transforms these predictions into action probabilities as follows. First, it estimates the negative value or 'aversiveness' of the two options. This is mathematically equivalent to estimated value for reward versions of the bandit task but reflects the fact that the outcome concerned (shock administration) is negatively valenced. Hence a high probability large magnitude shock is of high 'negative' value. The equation below assumes, for ease of reference, that the circle and square stimuli were used in the task:

$$g_{\text{circle}(i+1)} = F(r_{(i+1)}, \gamma) f_{\text{circle}(i+1)}$$

$$g_{\text{square}(i+1)} = F(1 - r_{(i+1)}, \gamma) f_{\text{square}(i+1)}$$

Here $g_{\text{circle}(i+1)}$ and $g_{\text{square}(i+1)}$ are the estimated negative values of the stimuli on the $i + 1^{\text{st}}$ trial, $f_{\text{circle}(i+1)}$ and $f_{\text{square}(i+1)}$ are the known shock magnitudes for the two stimuli, and $F(r, \gamma)$ is a linear transform within the bounds of 0 and 1:

$$F(r, \gamma) = \max[\min[(\gamma(r - 0.5) + 0.5), 1], 0]$$

The risk preference parameter γ allows the model to place greater weight on outcome magnitude ($\gamma < 1$) or outcome probability ($\gamma > 1$) when calculating the expected value. Effectively, this allows

the model to flexibly capture the extent to which a given participant prefers to minimize the probability or the magnitude of potential shocks.

Given, as outlined above, that the estimated values are negative (the probability x magnitude of an aversive outcome), the final action probabilities were generated using the following probability distribution:

$$P_{(choice=circle)} = \frac{1}{1 + \exp(-\beta(g_{(square)} - g_{(circle)})}$$

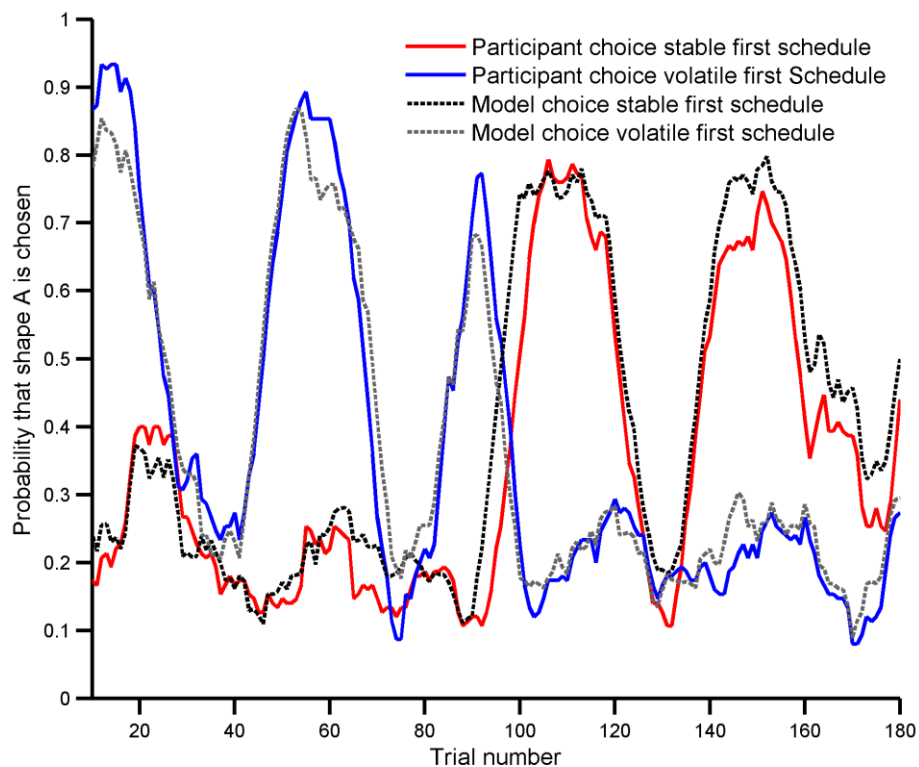
Here β is the inverse decision temperature and controls the degree to which the expected values are used in determining the shape chosen.

In total, the combined predictor and selector model has 3 free parameters (α , β , γ). These parameters were estimated separately for the stable and volatile blocks, for each participant, as in prior modeling of the structurally equivalent reward task⁸. Parameter estimates were obtained by calculating a joint posterior probability density function over all three parameter values, given the choice behavior of the participant. The values of the parameters were discretized in the joint function using the following limits: learning rate was represented using 30 points spaced equally from $\log(0.01)$ to $\log(1)$, risk preference using 15 points ranging from $\log(0.1)$ to $\log(10)$ and inverse temperature using 20 points ranging from $\log(1)$ to $\log(100)$. Individual parameter estimates (and the standard deviation (SD) of those estimates) were then defined as the expected value (and SD) of the marginal probability density function over the given parameter, obtained by direct integration⁸. We note that the SD of individual level parameter estimates is represented by error bars on the scatter plots in **Figures 2-4** and **Supplementary Figures S2, S6 and S9**. For completeness, we examined whether learning rates for the volatile and stable blocks differed in their estimation noise (i.e. the SD of the parameter estimates), and whether this varied as a function of trait anxiety. Of note, unlike the estimated parameters themselves, the SDs of the estimates are highly positively skewed requiring non-parametric statistical tests to be adopted in these analyses. Using a related

samples sign test, across participants, there was no significant difference in the standard deviation of the learning rate parameter estimates for the stable versus volatile blocks ($p=0.1$). Further, there was no relationship between trait anxiety and the standard deviation of the learning rate parameter estimate for either the stable (Spearman's $\rho = 0.01$, $p = 0.9$) or volatile (Spearman's $\rho = 0.03$, $p = 0.9$) task blocks.

When fitting the model parameters to participant choice, the first 10 trials of each block were omitted to allow comparison with the eyetracking data (here, initial trials were excluded to avoid luminance related changes in pupil size at the beginning of the blocks¹⁴). Additionally, trials in which no response was made were excluded (mean = 0.5% of trials). As the three free parameters are multiplicative (i.e. they act to multiply value or probability terms in the equations presented above), their logarithms were used when testing their relationship, across participants, with trait anxiety. Similarly change in the parameters between blocks was calculated as the difference of their log values.

An illustration of the ability of the model to recapitulate participant choice is provided in the figure below, which presents learning curves derived from participant choice on top of those derived from the model.



Learning Curves Illustrating Participant and Model Choice During the Learning Task. The plot shown here overlays the choice behavior of participants who completed the “stable first” task schedule (red line) or the “volatile first” task schedule (blue line) with the output (i.e. probability of choosing shape A) of the behavioral (Rescorla Wagner & Softmax Action Selector) model (dotted grey/black lines). In order to more clearly illustrate the effects of learning on participant choice and model output, the effect of shock magnitude (which varies randomly across trials) has been reduced by smoothing both participant choices and model predictions using a running average of 10 trials (i.e. this reduces the trial by trial variations attributable to changes in shock magnitude while preserving the lower frequency effects of learning). The model output was produced by running the behavioral model once for each participant, using the individually fitted parameter values, on the task trial sequences. As can be seen the model closely describes participant behavior during the task for both schedules.

To more formally test the model’s predictions of participant behavior, we discretized the model’s predictions (i.e., on a given trial, if it gave >50% likelihood that a participant would choose option A this was labeled as correct prediction of choice if the participant choose A). This inevitably leads to a relatively coarse measure of model performance, as trials where the model gives 90% likelihood of option A being chosen and 51% likelihood of option A being chosen are treated the same. Nevertheless, this binarized index still demonstrated that the model correctly predicted 81% of

choices across all trials and 76% of choices on ‘difficult trials’ (defined as in Behrens et al.⁸ as those where the expected value of the two options differed by less than 5). Additional comparison of this model against alternatives (two simplified models using only information about outcome probability not magnitude, and vice versa, and one keeping β and γ fixed across blocks) are given in the table below.

Model	Number of parameters	Negative Log Likelihood (lower is better)	AIC (lower is better)	Pseudo-R ² (higher is better)	Adjusted Pseudo-R ² (higher is better)
A					
(i) Full Model as used in main manuscript. (Uses prior outcomes and shock magnitude to guide choice)	6 (α, β, γ per block)	2217.6	4795.3	0.382	0.332
(ii) Prior outcome history model. (Uses only previous outcomes to guide choice)	4 (α, β per block)	2520.3	5280.7	0.300	0.267
(iii) Outcome value model (Uses only shock magnitudes associated with the two shapes to guide choice)	2 (β per block)	3516.8	7153.7	0.016	-0.0007
B					
Full Model with Single Risk Preference (γ) and Inverse Temperature (β) across both blocks	4 (α per block, single β, γ across blocks)	2305.9	4851.8	0.36	0.326

Model fits to participants’ choice behavior. (A). The full behavioral model (i), as used for the analyses reported in the main manuscript and in prior published studies of reward-learning⁸, assumes that participants combine information about outcome probability and magnitude when making their choices. An alternate possibility is that subjects predominantly used only one of these two sources of information. Hence, we examined the fit to participants’ behavior of the full model (i) relative to two simplified version of this model: (ii) a model which ignores shock magnitude information and (iii) a model which uses shock magnitude information but does not adapt its expected outcome probability on the basis of previous outcomes. We report two measures of model fit (negative log likelihood and the Pseudo-R² measure) as well as model fit statistics which penalize additional parameters (the Akaike Information Criteria, AIC, and the adjusted Pseudo-R² measure). As can be seen, the full model outperforms both the more limited models on all measures. (B) The

same measures of model fit are given for a simplified version of the full model which uses 2 learning rate parameters (one for the stable task block and one for the volatile task block) but only one risk preference and one inverse temperature parameter across the whole task. This simpler model performs marginally less well than the full model reported in the main manuscript (even when penalizing the additional parameters in the full model) and also has the disadvantage that it assumes that all of the effect of block volatility loads onto the learning rate parameter (rather than the risk preference or inverse temperature parameters). The full model allows these parameters to vary between blocks and thus does not rely on this assumption, also allowing for empirical examination of the relationship between anxiety and each of these parameter estimates. Note. The Pseudo-R² and adjusted Pseudo-R² measures were calculated as recommended by McFadden (McFadden, D., 1974, "Conditional Logit Analysis of Qualitative Choice Behavior" in *Frontiers in Econometrics*, ed. Zarembka, P., 105–142, Academic Press). α = learning rate; γ = risk preference, β = inverse decision temperature.

The Bayesian learner.

The Bayesian Learner has previously been described in detail by Behrens and colleagues⁸. Briefly, the learner estimates the likelihood of shock delivery following a given choice on trial $i+1$ given the outcome of the choice made on the current trial (i) and the previously estimated probability of that outcome. It takes advantage of the Markovian assumption that the full history of estimated outcome probabilities does not need to be used, but that this can be approximated by using the estimate of outcome probability from trial i alone. Under this assumption, the probability that a given choice (and by subtraction from 1, the probability for the other choice) will result in shock delivery on the next trial, i.e (r_{i+1}), is represented using a beta distribution with mean r_i and a width parameter V . V is defined as $\exp(\nu)$, where ν is the estimated volatility of the environment.

$$p(r_{i+1}|r_i, \nu) \sim \beta(r_i, V)$$

In other words, the estimated volatility ν is the log of the variance parameter (V) of this beta distribution. A second parameter used is k , which reflects uncertainty in the current estimate of environmental volatility – this is high in environments that rapidly transit between stable and volatile periods. Formally, $K = \exp(k)$ and is the width parameter of a normal distribution, centered on ν_i , which represents the probability distribution of volatility on the following trial ($i+1$):

$$p(\nu_{i+1}|\nu_i, k) \sim N(\nu_i, K)$$

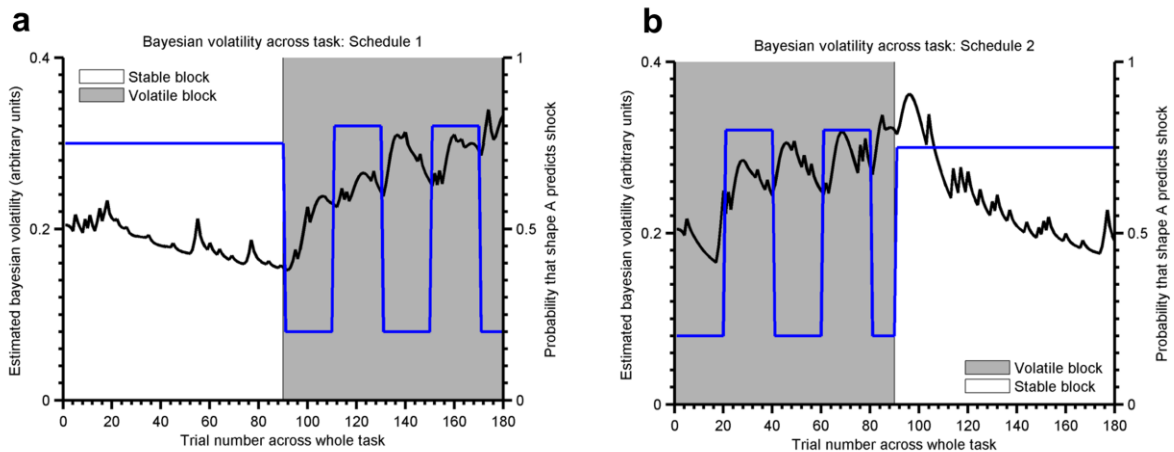
The joint probability of r_{i+1} , v_{i+1} and k is estimated following the outcome (y) of each trial as:

$$p(r_{i+1}, v_{i+1}, k | y_{\leq i+1}) \propto p(y_{i+1} | r_{i+1}) \int \left[\int p(r_i, v_i, k | y_{\leq i}) p(v_{i+1} | v_i, k) dv_i \right] p(r_{i+1} | r_i, v_{i+1}) dr_i$$

Trialwise estimates of the individual parameters (r_{i+1} , v_{i+1} , k) are obtained by marginalization of the joint probability function. In addition, trial-wise estimates of surprise are calculated by taking the negative logarithm of the conditional probability of the outcome on trial i , given the expectation of the Bayesian learner²², i.e:

$$surprise_i = -\log(p(y_i | E(r_i)))$$

The figure below illustrates the trialwise estimates of volatility obtained from the Bayesian Learner across the aversive learning task.



Environmental volatility v_i as estimated by the Bayesian learner on a trial-wise basis across the stable and volatile task blocks. Panels (a) and (b) depict the estimates that correspond to presentation of the stable block prior to, or following, the volatile block, respectively. Estimated volatility (black lines, left hand axes), can be seen to change gradually across each block of 90 trials as the learner accrues information about the (in)stability of task contingencies. This gradual accrual of information means that the initial trials of each block provide least information about environmental volatility, and a large number of trials are needed to assess the extent to which participants can alter their learning rate in response to volatility. The underlying outcome probabilities of the task are included as blue lines (right hand axes).

Modelling alternative potential influences on learning rate

As described above, participants' learning was estimated within each task block using a Rescorla-Wagner predictor¹⁰ coupled to a softmax based action selector. This enabled us to estimate changes in learning rate between the volatile and stable blocks of the task. There are, however, other potential influences on learning rate that are not captured by this model and which may conceivably be influenced by trait anxiety. In particular, learning rates often start high and decay exponentially during the course of learning (see, for example, Schiller, D. et al., 2008, "From fear to safety and back: reversal of fear in the human brain." *The Journal of neuroscience*: 28, 11517-11525). We therefore tested whether, using a more complex learning model in which a variety of factors compete for influence over learning rate, trait anxiety would be specifically associated with volatility related effects on participants' learning rate.

The learning rate (α) on trial (i) of the alternative model was calculated as:

$$\alpha_{(i)} = \alpha_{mean} + \omega_{Bayes}(v_i - \bar{v}) + \omega_{Exp}(e^{-\lambda i} - \overline{e^{-\lambda}})$$

Here v_i is the Bayesian volatility on trial i , \bar{v} represents the expected (mean) value of volatility across the task, $e^{-\lambda i}$ is a decaying exponential function and $\overline{e^{-\lambda}}$ is the expected value of this function. The learning rate on a given trial is a weighted sum of three terms: the mean learning rate across trials (α_{mean}), the trial-wise demeaned Bayesian volatility (subject specific weight = ω_{Bayes}) and the trial-wise demeaned exponential decay function (subject specific weight = ω_{Exp}). These parameters along with the slope of the exponential function (λ) and the decision temperature and risk preference terms from the selector model were fitted to participant behavior across the entire task (parameter fits were calculated by the same Bayesian procedure as used in the main study analysis). The results of this analysis are presented in **Figure S4**.