

Humans adapt rationally to approximate estimates of uncertainty

Erdem Pulcu , Michael Browning

Department of Psychiatry, University of Oxford, Oxford, UK • Oxford Health NHS Trust, Oxford, UK

 https://en.wikipedia.org/wiki/Open_access Copyright information

Reviewed Preprint

v1 • January 13, 2025

Not revised

eLife Assessment

This study is **important**, advancing our understanding of how humans adapt to uncertainty in dynamic environments by investigating the interplay between two types of uncertainty-volatility (systematic changes in outcomes) and noise (random variability in outcomes). Using an innovative experimental task, reinforcement learning (RL) models, and Bayesian Observer Models (BOM), the authors demonstrate that humans exhibit approximate rationality, often misattributing noise as volatility and adopting suboptimal learning rates in noisy conditions. The evidence is **compelling**, supported by a well-designed experimental task that independently manipulates noise and volatility, robust behavioral data, and computational modeling; the inclusion of BOM lesioning and physiological validation through pupillometry provides a nuanced understanding of suboptimal human learning. While the study could benefit from expanding the model space (e.g., by including latent state models) and offering greater clarity in task instructions and raw behavioral data, these limitations do not undermine the strength of the findings.

<https://doi.org/10.7554/eLife.103734.1.sa3>

Abstract

Summary

Efficient learning requires estimation of, and adaptation to, different forms of uncertainty. If uncertainty is caused by randomness in outcomes (noise), observed events should have less influence on beliefs, whereas if uncertainty is caused by a change in the process being estimated (volatility) the influence of events should increase. Previous work has demonstrated that humans respond appropriately to changes in volatility, but there is less evidence of a rational response to noise. Here, we test adaptation to variable levels of volatility and noise in human participants, using choice behaviour and pupillometry as a measure of the central arousal system. We find that participants adapt as expected to changes in volatility, but not to changes in noise. Using a Bayesian observer model, we demonstrate that participants are, in fact, adapting to estimated noise, but that their estimates are imprecise, leading them to misattribute it as volatility and thus to respond inappropriately.

It is much easier to respond appropriately to an event if we know what has caused it. For example, if heavy traffic means that our drive into work takes longer than normal, the best course of action the next time we have to make the journey depends on what caused the traffic to be heavier (Yu & Dayan, 2005 [↗](#)). If it was caused by a one-off or random event, such as a broken-down lorry, then we should continue using the same route as before, whereas if it was caused by some longer-term change, perhaps there are new road works nearby disrupting the traffic, we should consider a different route. Frequently, however, the causes of events are not obvious, we experience the heavy traffic but aren't sure why it has occurred. In these situations, the best we can do is make an educated guess, based on our experience, about what broad type of causal process has led to recent events. In the case of the drive into work, if the traffic has been heavier for a number of days in a row it is likely that some prolonged shift has occurred, and we should change routes, whereas if the traffic changes noisily from day to day, then we should probably stick with our usual route. In the learning literature, this problem is often framed as a competitive attribution of uncertainty to one of two types; expected uncertainty, which is caused by the variability of noisy associations and unexpected uncertainty, which is caused by longer lasting changes (sometimes called volatility) in an association (Behrens et al., 2007 [↗](#); Browning et al., 2015 [↗](#); Nassar et al., 2012 [↗](#); Pulcu & Browning, 2017 [↗](#); Yu & Dayan, 2005 [↗](#)). The behavioural importance of this attribution process can be seen in the driving example given above; an event caused by noise requires the opposite behavioural response (continuing to use the same route) than the same event caused by volatility (switching routes). Consequently, effective decision making often depends on the accurate attribution of uncertainty, with misattribution having a substantial detrimental effect on choice (Pulcu & Browning, 2019 [↗](#)).

The influence of events on subsequent choice can be estimated within a reinforcement learning framework as the learning rate parameter (Sutton & Barto, 2018 [↗](#)), with a higher learning rate indicating a greater influence of the event on behaviour. As described above, the normative response to changes in volatility and noise is to use a higher learning rate when volatility is high and/or noise is low (Pulcu & Browning, 2019 [↗](#); Yu & Dayan, 2005 [↗](#)). A large number of studies have found the predicted increase in learning rates in response to higher outcome volatility in human learners (Behrens et al., 2007 [↗](#), 2008 [↗](#); Browning et al., 2015 [↗](#); Gagne et al., 2020 [↗](#); Nassar et al., 2012 [↗](#); Pulcu & Browning, 2017 [↗](#)). In contrast, the evidence for adaptation of learning in response to changes in outcome noise is less complete. Previous studies have described the expected reduction of learning rates when outcome noise is high, but only when the level of noise is explicitly signalled in a task (Diederer & Schultz, 2015 [↗](#)), or when it is made unambiguous by virtue of being very much smaller than changes in outcome caused by volatility (Nassar et al., 2010 [↗](#), 2012 [↗](#)). As illustrated in the driving example above, we are often faced with situations in which there exists significant uncertainty about whether an event has been caused by volatility or noise. To date however, the degree to which human learners are able to discriminate between these types of uncertainty, when they are not explicitly labelled, has not been closely examined.

From a neurobiological perspective, activity of central modulatory neurotransmitter systems have been argued to represent distinct sources of uncertainty during learning, with central norepinepheric (NE)/locus coeruleus (LC) activity described as signalling changes in the associations (i.e. volatility) and central cholinergic activity representing noise (Yu & Dayan, 2005 [↗](#)). Electrophysiological measures of LC activity in non-human primates have been shown to correlate with pupil dilation (Joshi et al., 2016 [↗](#)) suggesting it may be possible to estimate activity in this system in humans using pupilometry. Taking this approach, indirect support for this role of the NE system has been provided by studies of human participants that report greater pupillary size in volatile relative to stable contexts (Browning et al., 2015 [↗](#); Nassar et al., 2012 [↗](#); Pulcu & Browning, 2017 [↗](#)). However, the pupil also responds to other learning signals, such as surprise (Browning et al., 2015 [↗](#); O'Reilly et al., 2013 [↗](#); Preuschoff et al., 2011 [↗](#)) and has been reported as being smaller when outcome noise is high (Nassar et al., 2012 [↗](#)). Neuroimaging evidence suggests an association between activity in other central neurotransmitter nuclei, including the cholinergic

basal forebrain, and pupil dilation (de Gee et al., 2017 [↗](#)). Overall, this suggests that the pupillary signal may reflect a more general belief updating process (O'Reilly et al., 2013 [↗](#)) rather than a specific volatility signal and thus that, like learning rates, pupil size should increase when noise is reduced as well as when volatility is increased.

In this paper we test whether human participants modify their learning in situations in which the attribution of uncertainty as volatility or noise is challenging (**Figure 1a-c** [↗](#)). We report the results of a study in which participants completed a learning task during which the noise and volatility of both win and loss outcomes were independently manipulated. Participant behaviour was characterised using learning rate parameters derived from reinforcement learning models of choice, while interpretation of the results was facilitated by a Bayesian Ideal Observer model that was developed to provide a benchmark comparator to participant behaviour (Behrens et al., 2007 [↗](#); Nassar et al., 2012 [↗](#); Piray & Daw, 2021 [↗](#); Pulcu et al., 2022 [↗](#)) and by the collection of pupillometry data as a physiological marker of central neurotransmitter function (de Gee et al., 2017 [↗](#); Joshi et al., 2016 [↗](#)). It was predicted that human participants would be able to adapt appropriately to the cause of the events they encountered—using a higher learning rate, and displaying increased pupil size, when volatility was high and when noise was low for both win and loss outcomes (**Figure 1d** [↗](#)).

Results

Participant Demographics

70 participants (see Supplementary Table 1 for demographic information) completed a learning task in which they had to choose one of two stimuli based on the separately estimated magnitudes of win and loss outcomes associated with the stimuli (**Figure 1** [↗](#)).

Experimental Manipulation of Volatility and Noise Influences Participant Choice Behaviour

As explained above, high levels of volatility and low levels of noise should increase the degree to which outcomes influence choice behaviour. A crude metric of this effect is provided by examining participant choice as a function of the previous outcome. In the task, a win outcome of >50p or a loss outcome of <50p associated with Shape A prompts participants to select Shape A in the subsequent trial, with the other outcomes (i.e. win <50P and loss <50p) prompting choice of Shape B. The influence of the outcomes on choice can therefore be roughly estimated as the relative proportion of trials in which Shape A was chosen when it was prompted by a previous outcome of a given magnitude, compared to when Shape B was prompted. Analysis of this choice metric (**Figure 2a-b** [↗](#)) found the expected effect of volatility, with participant choice being more influenced by previous outcomes when volatility was higher ($F(1,696)=99.8, p<0.001$). An effect of noise was observed, but in the opposite direction to expected, with outcomes influencing choice more when noise was increased ($F(1,696)=4.79, p=0.03$). No significant difference between the influence of win and loss outcomes was found ($F(1,696)=1, p=0.32$) and there was no interaction between volatility and noise ($F(1,693)=0.61, p=0.4$). Having found some evidence of an impact of the uncertainty manipulations on a crude measure of subject choice we next sought to characterise this effect using reinforcement learning models fitted to participant choices.




Figure 1.

The Magnitude Learning Task

(a) Timeline of one trial from the learning task. On each trial participants were presented with two abstract shapes and were asked to choose one of them. The empty bars above and below the fixation cross represented the total available wins and losses for the trial, the full length of each bar was equivalent to £1. Participants chose a shape and then were shown the proportion of each outcome that was associated with their chosen shape as coloured regions of the bars (green for wins and red for losses). The empty portions of the bars indicated the win and loss magnitudes associated with the unchosen option, allowing participants to infer which shape would have been the better option on every trial. The task consisted of six blocks of sixty trials each. The volatility and noise of the two outcomes varied independently between blocks with different shapes used in each block. Panel **b** illustrates outcomes from the four block types. As can be seen blocks with high volatility and low noise (top left), and those with low volatility and high noise (bottom right), present participants with a similar range of magnitudes. Participants therefore have to distinguish whether variability in the outcomes is caused by volatility or noise from the temporal structure of the outcomes rather than the size of changes in magnitude (cf; [Diederer & Schultz, 2015](#); [Krishnamurthy et al., 2017](#); [Nassar et al., 2012](#)). Panel **c** shows two example blocks (one block in grey, the other in white) with both win (green) and loss outcomes (red) displayed. Panel **d** shows the expected adaptation of learning rates in response to the manipulation of volatility and noise; for both win and loss outcomes, learning rates should be increased when volatility is high and when noise is low.




Figure 2.

The impact of uncertainty manipulations on participant choice.

Panels **a** and **b** report a summary metric for the effect of win and loss outcomes on subsequent choice. The metric was calculated as the proportion of trials in which an outcome of magnitude 51-65 associated with Shape A was followed by choice of the shape prompted by the outcome (i.e. Shape A for win outcomes, Shape B for loss outcomes) relative to when the outcome magnitude was 49-35 (see methods and materials for more details). We focused on this outcome range as these range of magnitudes were covered by all volatility x noise conditions and it was dictated by the relatively smaller range coverage in the low volatility low noise condition (also see **Figure 1C** loss outcomes shown in red between trials 60-120). The higher this number, the greater the tendency for a participant to choose the shape prompted by an outcome. As can be seen, the outcome of previous trials had a greater influence on participant choice when volatility was high, with a small effect of noise, in the opposite direction to that predicted. Panels **c** and **d** report the win and loss learning rates estimated from the same data. Again, the expected effect of volatility is observed, this time with no consistent effect of noise. Bars represent the mean (±SEM) of the data, with individual data points superimposed.

Participants Adjust Normatively to Changes in Volatility but not Noise

We aimed to capture the computational process that underlies participant choice behaviour by fitting different reinforcement learning models to choice data separately for each block of the task and each participant. The best fitting RL model included separate learning rates for win and loss outcomes allowing estimation of the degree to which participants adjusted these learning rates in response to the block-wise changes in outcome volatility and noise (see supplementary Materials and Methods for model comparison and selection analyses).

Consistent with the analysis of choice data reported above, there was a significant main effect of volatility (**Figure 2c-d**; $F(1,696)=22.2$, $p<0.001$), with a higher learning rate used when volatility was high. There was no main effect of noise ($F(1,696)=0.63$, $p=0.43$) on learning rate or outcome valence ($F(1,696)=0.15$, $p=0.7$). An interaction between volatility and noise ($F(1,693)=7.74$, $p=0.006$) was also significant. A higher volatility led to a significantly raised learning rate when noise was low ($F(1,383)=27.1$, $p<0.01$), with a non-significant increase when noise was high ($F(1,311)=1.13$, $p=0.29$). Higher noise was associated with a non-significant reduction in learning rates when volatility was high ($F(1,347)=2.57$, $p=0.11$) but to a significant increase in learning rate when volatility was low ($F(1,347)=4.7$, $p=0.031$).

In summary, analysis of both crude choice data and learning rates indicates that participants adapted appropriately to changes in the volatility of learned associations but did not show a consistent response to changes in noise. In the next section we utilise a Bayesian Observer Model (BOM) to investigate potential causes for this relative insensitivity to noise.

Using a Bayesian Observer Model to Characterise Noise Insensitivity

Bayesian Observer Models (BOM) can be used as normative benchmarks against which human behaviour may be compared (Behrens et al., 2007; Nassar et al., 2012; Piray & Daw, 2021; Pulcu et al., 2022). BOMs are generally not fit to participant choice, rather these models invert a generative process assumed to underlie observed events and provide an estimate of the belief of an idealised agent exposed to the same outcomes as participants. We developed a BOM (Pulcu et al., 2022) based on the generative process underlying the outcome magnitudes of our task (**Figure 3a**). The BOM explicitly estimates the volatility and noise of the outcomes and uses these estimates to influence its belief about the likely magnitude of upcoming outcomes (see methods for more details). We first tested whether the BOM reproduced the normative learning rate adaptation to changes in volatility and noise described in the introduction, by exposing the model to the same outcomes as participants and using the model's belief about the likely magnitude of the win and loss outcome on each trial to generate choices. We then estimated the effective learning rate of the model by fitting the same RL model used to analyse participants' choices to the model's choices. These learning rates are presented in **Figure 3f** (**Figure 3e** reproduces the learning rates of participants, averaged across wins and losses, for comparison). As can be seen the BOM adapts as expected, using a higher learning rate both when volatility increases ($F(1,696)=422$, $p<0.001$) and when noise decreases ($F(1,696)=21.2$, $p<0.001$). No effect of outcome valence or interaction between volatility and noise (all $p>0.09$) was observed.

Having shown that an optimal learner adjusts its learning rate to changes in volatility and noise as expected, we next sought to understand the relative noise insensitivity of participants. In these analyses we “lesion” the BOM, to reduce its performance in some way, and then assess whether doing so recapitulates the pattern of learning rate adaptation observed for participants (**Fig 3e**). First, we tested the impact of completely removing the ability of the BOM to adjust to changes in




Figure 3.

The Behaviour of Bayesian Observer Models.

Bayesian Observer Models (BOM) invert generative descriptions of a process, indicating how an idealised observer may learn. We developed a BOM based on the generative model of the task we used (**a**). Details of the BOM are provided in the methods, briefly it assumes that observations (y_i) are generated from a Gaussian distribution with a mean (μ_i) and standard deviation (SD_i). Between observations, the mean changes with the rate of change controlled by the volatility parameter ($v\mu_i$). The standard deviation and volatility of this model estimate the noise and volatility described for the task. The last parameters control the change in volatility ($k\mu$) and standard deviation (vSD) between observations, allowing the model to account for different periods when these types of uncertainty are high and others when they are low. The BOM adjusts its learning rate in a normative fashion (**f**), increasing it when volatility is higher, or noise is lower. The BOM was lesioned in a number of different ways in an attempt to recapitulate the learning rate adaptation observed in participants (shown in panel **e**). Removing the ability of the BOM to adapt to changes in volatility (**b**) or noise (**c**) did not achieve this goal (**g,h**). However, degrading the BOMs representation of uncertainty (**d**) was able to recapitulate the behavioural pattern of participants (**i**). Bars represent the mean ($\pm$ SEM) of participant learning rates, with raw data points presented as circles behind each bar.

either volatility (**Figure 3b**) or noise (**Figure 3c**) by removing the top nodes of the model (i.e. *kmu* or *vs* respectively). Removing these nodes forces the BOM to estimate the mean volatility or noise across all task blocks rather than estimating local periods where they are higher or lower (see supplementary video). As illustrated in **Figure 3g-h**, neither of these lesions recapitulates the pattern of learning rates observed in participants, with the volatility lesioned model attributing increased volatility to noise and thus decreasing its learning rate during periods of higher volatility (main effect of volatility; $F(1,696)=11.9$, $p<0.001$) and the SD-lesioned model treating any form of uncertainty as volatility and thus increasing its learning rate in response to increased noise (main effect of noise; $F(1,696)=227$, $p<0.001$). This suggests that human participants are able to adapt to changes in outcome volatility and noise to some degree but are less sensitive to these changes than the intact BOM.

We next assessed whether a relative degrading of the model's representation of volatility and noise (**Figure 3d**) altered its behaviour in a manner similar to participants. This was achieved by independently coarsening the model's representation of volatility and noise, with the degree of coarsening selected to make the model's choices as similar as possible to those of a given participant. Details of this coarsening process are provided in the methods section, but in simple terms, at one extreme, the intact model's beliefs about current volatility and noise are represented as probability distributions over many possible values, with the number of values used gradually reduced during coarsening, until the coarsest model treats each form of uncertainty as being either "high" or "low". As can be seen from **Figure 3i**, this relative degrading of the model's representation of uncertainty more closely recapitulated the learning rates observed in participants, with a significant increase in learning rate in response to increased volatility ($F(1,696)=59$, $p<0.001$) and no effect of noise ($F(1,696)=2.3$, $p=0.13$). In the next sections we characterise how coarsening the BOM changes its behaviour and assess whether it provides an accurate account of participants' noise insensitivity.

The Degraded BOM Misattributes Noise as Volatility

The BOM was degraded by reducing the number of bins it used to represent volatility and/or noise, until its behaviour most closely matched that of participants. This process led to a greater coarsening of the noise than the volatility dimension (**Figure 4a**; $F(1,69)=49$, $p<0.001$), with no effect of outcome valence ($F(1,69)=0.73$, $p=0.4$), suggesting that the degraded model maintained a generally less precise representation of noise than volatility. In order to investigate the impact of this coarsening on the model's beliefs, we used the degraded BOM's estimates of volatility and noise to categorise task trials as either high or low volatility/noise (i.e. trials in which the model's estimates of these variables were higher/lower than the mean) and compared these to the same trial labels generated by the intact BOM. Consistent with the greater degradation of the noise dimension, coarsening the model caused it to miscategorise more trials which the intact BOM had labelled as having high than low noise (**Figure 4b**; $F(1,69)=30.7$, $p<0.01$) with no effect of volatility ($F(1,69)=1.9$, $p=0.17$) or outcome valence ($F(1,69)=0.004$, $p=0.95$). As illustrated in **Figure 4c**, when the degraded BOM miscategorised high noise trials, it tended to label them as having high, rather than low, volatility. Overall, these results indicate that coarsening the BOM caused it, relative to the intact BOM, to misattribute high noise trials as high volatility trials.

The Degraded BOM Rescues Optimal Behaviour

The process of fitting the degraded BOM to participant behaviour can be understood as searching for a configuration of the model (akin to a grid-based maximum likelihood estimation) in which participant choice conforms to the normative response to volatility and noise coded in the model's structure. In other words, participants' learning rates should increase when the degraded BOM's estimate of volatility is high and, critically, when it estimates that noise is low. We demonstrate this by reanalysing participant behaviour, using the trial labels of the degraded BOM to indicate periods of low/high volatility and noise in place of the task block labels used in the original




Figure 4.

Analysis of the behaviour of the degraded BOM.

The process of degrading the BOM involved reducing the number of bins used to represent the volatility and noise dimensions independently until the choice of the model matched that of participants. Panel **a** illustrates the number of bins selected by this process for the volatility and noise dimensions (averaged across win and loss outcomes). As can be seen the degraded BOM maintained a less precise representation of noise than volatility. In order to understand the behaviour of the degraded model, the model's estimated vmu_i and SD_i were used to label individual trials as high/low volatility and noise (NB greater than or less than the mean value of the estimates). These trial labels were compared with the same labels from the intact model, which were used as an ideal comparator (panels **b** and **c**). Panel **b** illustrates the proportion of trials in which the labels of the two models agreed, arranged by the ground truth labels of the full model and averaged across win and loss outcomes. The dotted line indicates the agreement expected by chance. The degraded model trial labels differed from those of the full model particularly for high noise trials, with no impact of trial volatility. Panel **c** provides more details on how the degraded model misattributes trials. In this figure, the labels assigned by the full model are arranged along the x axis. The colour of each square represents the proportion of trials with a specific full model label that received the indicated label of the degraded model (arranged along the y axis). The diagonal squares illustrate agreement between models as reported in panel **b**. As highlighted by the red outlines, trials which the full model labelled as having high noise were generally mislabelled by the degraded model as having high volatility. Reanalysis of participant choices using the trial labels provided by the full (panel **e**) and degraded (panel **f**) models indicate that participants adapt their learning rates in a normative fashion when the degraded model trial labels are used (panel **f**), but not when the full model labels are used (panel **e**). Panel **d** illustrates the same analysis using the original task block labels for comparison. Bars represent the mean (±SEM) of participant learning rates, with raw data points presented as circles behind each bar. See supplementary figure 1 for a comparison of the behaviour of the degraded BOM with an alternative fitted model.

analysis. In effect, this approach allows us to test whether participants make internally consistent errors during learning, i.e. reducing their learning rates for outcomes that they thought — instead of the actual task structure — were associated with high volatility and/or low noise. As can be seen (**Figure 4f**), participants significantly increased their learning rate when the degraded BOM estimated volatility to be high ($F(1,566)=86, p<0.001$) and noise to be low ($F(1,566)=81, p<0.001$). In control analyses, this normative response to uncertainty was not seen when the labels from the intact rather than the degraded BOM were used (**Figure 4e**), or when the BOM's representation of outcome mean was degraded, rather than its estimates of volatility and noise (supplementary materials).

Assuming Human Participants Use the Degraded BOM's Estimates of Volatility and Noise also Rescues Normative Pupillary Response

If the degraded BOM is a fair representation of how participants are performing the learning task, then we would expect it to be better able to explain physiological markers of uncertainty estimation than the simple task block structure or the intact BOM. Specifically, participants' pupils should be larger when the degraded BOM thinks that volatility is high and when it thinks noise is low. We first show (**Figures 5a-c**) that participants' pupils do not adapt normatively to the task block structure, with no main effect of block volatility ($F(1,1723)=0.002, p=0.9$) and an increase of pupil size in response to higher noise ($F(1,1723)=13.8, p<0.001$). In contrast, analysis using the trial labels derived from the degraded model (**Figure 5d-f**) recovered the expected increase in pupil size in response to both raised volatility ($F(1,2067)=105, p<0.001$) and reduced noise ($F(1,2067)=42.3, p<0.001$) suggesting that the model provides a reasonable measure of participants' estimates of these parameters. Finally, we tested whether the degraded BOM was able to explain more variance in the pupil data than the intact BOM. In order to do this, we first regressed participants' pupil data against the estimated volatility and noise of the intact BOM, as well as a range of other task related factors (**Figure 5g**; see methods for more details of analysis). Having removed the variance accounted for by these factors, we then regressed the residuals of this first level analyses against the degraded model's estimates of volatility and noise. This second level analysis (**Figure 5h-i**) indicated that the degraded model was able to account for variance associated with outcome noise that was not explained by the full model ($F(1,286)=4.1, p=0.04$), but did not explain additional variance associated with outcome volatility ($F(1,286)=0.1, p=0.75$). In summary, assuming that participants used the degraded BOM's estimates of outcome volatility and noise rescued the normative pattern of physiological adaptation during the task.

Discussion

Humans respond in a rational, if approximate, manner to the causal statistics of dynamic environments. We found that participants adapted as expected to changes in outcome volatility, but were relatively insensitive to changes in noise. Using a degraded Bayesian Observer Model (BOM) to characterise participants' behaviour suggested that they responded appropriately to a relatively coarse estimation of the level of noise, that led to its misattribution as volatility. Analysis of pupillometry data using the degraded model again suggested that participants were responding normatively to changes in estimated noise, but that these estimates diverged from the true noise of experienced outcomes. These results illustrate that human learners are able adapt to the statistical properties of their environment, but during this process they make internally consistent errors, utilising higher learning rates as a result of misattributing environmental noise as volatility which also leads to suboptimal choice.

Using a task in which volatility and noise varied independently between blocks, we found that human learners adapted as expected (Behrens et al., 2007; Browning et al., 2015; Nassar et al., 2012; Pulcu & Browning, 2019) to blockwise changes in the volatility of both win and loss outcomes, increasing the learning rate used when volatility was high vs. low. In contrast, the




Figure 5

Analysis of pupillometry data.

Z-scored pupil area from 2 seconds before to 6 seconds after win (panel **a**) and loss (panel **b**) outcomes, split by task block. Lines illustrate average size, with shaded area illustrating SEM. Panel **c** Pupil size averaged across whole outcome period and both win and loss outcomes. Pupil size did not systematically vary by task block. Panels **d-f**, as above but using the trial labels derived from the degraded model. Pupil size was significantly larger for trials labelled as having high vs. low volatility and low vs. high noise. Panel **g** displays the mean (SEM) effect of volatility and noise as estimated by the full BOM derived from a regression analysis of pupil data. The residuals from this analysis were then regressed against the estimated volatility and noise from the degraded model. A time course of the regression weights from this analysis is shown in panel **h**, with the mean coefficients across the whole period shown in panel **i**. The degraded model's estimated noise accounted for a significant amount of variance not captured by the full model (pink line in **h** is below 0, the mean effect across the period is represented by dashed lines and arrows in panel **i**). See supplementary figure 2 for comparison of the degraded BOM with an alternative fitted model.

expected reduction of learning rates in response to increased outcome noise was not apparent, with participants employing a significantly higher learning rate in response to increased noise when volatility was low and a numerically lower learning rate when volatility was high. The absence of a normative response to blockwise changes in noise is at odds with previous work which has described a reduction in learning rates during periods of high noise (Diederen & Schultz, 2015 [DOI](#); Nassar et al., 2010 [DOI](#), 2012 [DOI](#)). However, in this previous work the level of noise was either explicitly presented to participants (as a bar on screen representing the standard deviation of the generative process in Diederen & Schultz) or was made unambiguous by being very different from changes caused by volatility (in Nassar et al, noise was generated using an SD of 5 or 10, while the average change due to volatility was 100). By design, in the current task high noise and volatility resulted in a similar range of magnitudes (**Figure 1b** [DOI](#)) forcing participants to use the temporal sequence of outcomes to discriminate between the different forms of uncertainty. Our behavioural results suggest that, in the absence of unambiguous differences between outcomes caused by volatility and those caused by noise, participants' ability to estimate and/or adapt to changes in noise is reduced. Interestingly, a recent study reported that participants do not adjust their choice or estimated confidence in response to variability in the orientation of arrays of visual gratings (Herce Castañón et al., 2019 [DOI](#)), suggesting that an insensitivity to outcome noise may be a general feature of human decision making, rather than a specific component of learning.

Noise fundamentally limits the reliability of information (MacKay, 2003 [DOI](#)) and ignoring it has a clear detrimental impact on inference (**Figure 3h** [DOI](#)), causing agents to be unnecessarily influenced by chance events (Pulcu & Browning, 2019 [DOI](#)). It would therefore be surprising if human learners were completely insensitive to this process, particularly given evidence that they can respond normatively when the level of noise is unambiguous (Diederen & Schultz, 2015 [DOI](#); Nassar et al., 2010 [DOI](#), 2012 [DOI](#)). We developed an ideal Bayesian Observer Model (BOM; Behrens et al., 2007 [DOI](#); Nassar et al., 2010 [DOI](#); Piray & Daw, 2021 [DOI](#); Pulcu et al., 2022 [DOI](#)) to investigate the degree to which participants were adapting to noise. The intact BOM displayed the expected behavioural response to changes in both volatility and noise (**Figure 3f** [DOI](#)) and, as a result, did not accurately capture the behaviour of participants (**Figure 3e** [DOI](#)). Completely removing the BOM's ability to adapt to noise (or volatility) did not recapitulate participant choice behaviour (**Figure 3g-h** [DOI](#)), whereas coarsening its representation of volatility and noise, produced a much closer match (**Figure 3i** [DOI](#)). This suggests that participants were relatively, rather than completely insensitive to noise and that they tended to misattribute high noise as volatility (**Figure 4** [DOI](#)). However, an important caveat to this interpretation is that the degree of coarsening was selected using participants' choices. The better behavioural match of the coarsened BOM to participant learning rates may therefore be simply because this model was fitted to the same choices used to calculate the learning rates, whereas the intact and fully lesioned models were not. We therefore sought to validate the coarsened BOM by assessing its ability to account for participants' pupillary data, and by comparing it with an alternative fitted BOM which coarsened the representation of the generative mean, rather than the estimated uncertainty (see supplementary materials). Participants' pupil size did not vary systematically between different block types, whereas they were significantly larger when the degraded BOM estimated volatility to be high and noise to be low (**Figure 5a-f** [DOI](#)). Similarly, the estimated noise of the degraded BOM accounted for additional variance in pupil size, over and above the intact BOM (**Figure 5g-i** [DOI](#)). In contrast, the alternative mean-degraded BOM did not recapitulate participants' learning rates (supplementary figure 3) and was not able to account for changes in participant pupil size (supplementary figure 4). The finding that participants' pupil size covaries in the expected direction with the degraded BOM's estimated levels of both volatility and noise provides some reassurance that the model is capturing the dynamics of participants' uncertainty estimates. More generally, the presence of both volatility and noise signals in this data, indicate that, as suggested previously (Nassar et al., 2012 [DOI](#); O'Reilly et al., 2013 [DOI](#)), the pupillometry signal reflects general belief updating rather than specifically volatility.

An outstanding question is why participants might be particularly insensitive to changes in outcome noise. It is tempting to try to answer this question by reference to the processes by which the BOM was coarsened (i.e. the insensitivity was caused by a reduction in the precision by which noise was represented in a multi-dimensional probability distribution). However, the BOM described here was developed as an algorithmic description of how the learning task may be solved. As far as we are aware, there is little evidence that it accurately describes the cognitive or neural implementation of uncertainty estimation. Alternative algorithmic approaches to the general problem of uncertainty estimation have been described (Kalman, 1960 [↗](#); Nassar et al., 2010 [↗](#); Piray & Daw, 2021 [↗](#); Pulcu & Browning, 2019 [↗](#)), including simpler approaches that avoid computationally expensive representations of multi-dimensional distributions (Kalman, 1960 [↗](#); Nassar et al., 2010 [↗](#)) and which therefore may be more likely implementational candidates. In other words, the current results indicate that human learners are relatively insensitive to changes in outcome noise, but do not specify the lower level mechanisms that determine this effect.

In conclusion, human learners adapt rationally, to estimates of the volatility and noise of experienced outcomes. However, these estimates are approximate leading to a relative insensitivity to outcome noise.

Methods

Experimental model and subject details

Participants

70 English-speaking participants aged between 18 and 65 were recruited from the general public using print and online advertisements. A previous study (Pulcu & Browning, 2017 [↗](#)) on behavioural response to changes in volatility reported an effect size of $d=0.7$. As the effect size of a noise manipulation was not clear, we recruited a sample size sufficient to detect an effect size of half this value ($d=0.35$) with 80% power. Participants were excluded from the study if they had any psychological or neurological disorders or were currently on psychotropic medication.

Method details

General procedure

Participants attended a single study visit during which they completed the learning task. The study was approved by the University of Oxford Central Research Ethics Committee (R49753/RE001). All participants provided written informed consent to take part in the study, in accordance with the Declaration of Helsinki.

Behavioural Paradigm

The reinforcement learning (RL) task consisted of six blocks, each comprising 60 trials. In each trial, participants were presented with two abstract shapes taken from the Agathodaimon font (i.e. shape A and shape B). Two different shapes were used in each block, with rest sessions between blocks. The shapes were presented randomly on either side of the screen. Participants were explicitly instructed that this randomised location did not influence the outcome magnitudes. Participants attempted to accumulate as much money as possible by learning the likely magnitude of the wins and losses associated with each shape and using this information to guide their choice. On each trial, participants chose one of two shapes, with their choice highlighted by a black frame (see **Fig 1a** [↗](#)). Following the choice, the win and loss amounts associated with the chosen shape were presented, in randomised order, for a jittered period (2-6 sec, mean: 4 sec) inside two empty

bars, above and below the fixation cross. The win amount was shown as a green area in the upper bar, and the loss amount represented as a red area in the lower bar. The total length of each bar represented £1 (i.e. of wins or losses) and thus the amount associated with the chosen shape was the proportion of the bar filled by the green/red areas (e.g. three quarters of the upper bar being green, would mean that the chosen option was associated with a win of 75p). Participants were informed that the unshaded area of each bar was the amount associated with the unchosen option. Thus, on each trial participants knew how much they had won/lost and how much they would have won/lost if they had chosen the other option. This feature simplified the task; rather than having to separately estimate the wins and losses associated with each shape, participants only had to estimate these values for one shape (with the other shape being £1 minus this value). For each trial participants received the difference between the win and loss amounts associated with their choice. A running total amount of money was displayed in the centre of the screen, under the bars and was updated at the beginning of the subsequent trial with the recent winnings.

The wins and the losses associated with each shape followed independent outcome schedules (**Figure 1b** [↗](#)), generated from a Gaussian distribution. In each block, the win and loss outcomes had either high or low volatility and high or low noise. When volatility was low, the mean of the Gaussian distribution remained constant, when volatility was high the mean changed from between 25-40 and 60-75 every 9-15 trials. When noise was low the standard deviation of the Gaussian was set to 5, whereas when noise was high the standard deviation was 35. As can be seen from **Figure 1b** [↗](#), these schedules resulted in similar ranges of outcome magnitudes for periods of high noise and high volatility. The first block for every participant had high volatility and low noise for both win and loss outcomes and was used to familiarise participants with the task. Choices from this block were not used in the analyses presented (although including them does not alter the reported pattern of results). The schedules in the remaining five blocks were presented in a randomised order with the constraint that, across both win and loss outcomes, each of the four combinations of volatility and noise level (**Figure 1B** [↗](#)) were presented either 2 or 3 times. Thus, while each participant completed at least two blocks with each of the four combinations of high/low volatility/noise, the specific pairings of win and loss volatility/noise levels, differed across participants. This approach was used in preference to a fully factorial design in order to keep the total task duration to a manageable level. At the end of the experiment, participants were paid one fifth of their total winnings, plus a £15 baseline rate for turning up to take part.

Pupillometry data was collected for 36 of the 70 participants. During collection of pupillary data, the task was presented on a VGA monitor connected to a laptop computer running Presentation software version 18.3 (Neurobehavioural Systems). An identical behavioural version of the task, presented using Psychtoolbox 3.0 on MATLAB (MathWorks Inc.), was used to collect behavioural data from the remaining 34 participants. In the pupillometry version, participants' heads were stabilised using a head-and-chin rest placed 70 cm from the screen on which the eye tracking system was mounted (Eyelink 1000 Plus; SR Research). The eye tracking device was configured to record the coordinates of both of the eyes and pupil area at a rate of 500 Hz. The task stimuli were drawn on either side of a fixation cross which marked the middle of the screen and were offset by 7° visual angle. The testing session lasted approximately 70 min per participant.

Analysis of Choice Data

Non-model based measure of the influence of outcomes

The manipulation of uncertainty in the reinforcement learning task is expected to alter the degree to which participants' choices are influenced by the outcomes they experience. A simple, if somewhat crude, measure of this influence can be calculated as the proportion of trials in a block in which participants select the choice prompted by the win or loss outcomes on the previous trial.

Generally, win outcomes of >50p and loss outcomes of <50p associated with a shape will prompt selection of the same shape on the next trial, whereas other outcomes will prompt selection of the alternative shape. The overall effect of win outcomes on choice can therefore be estimated as:

$$P_{(choice==A | previous\ win\ outcome\ for\ A>50p)} - P_{(choice==A | previous\ win\ outcome\ for\ A<50p)}$$

That is, the probability of choosing shape A, given that, on the previous trial, a win of >50p was associated with shape A – the probability of choosing Shape A, given that, on the previous trial a win of <50p was associated with Shape A. Similarly, the effect of loss outcomes is estimated as:

$$P_{(choice==A | previous\ loss\ outcome\ for\ A<50p)} - P_{(choice==A | previous\ loss\ outcome\ for\ A>50p)}$$

However, choice is also influenced by the magnitude of the outcome; a win of 90p will have a greater effect on subsequent choice than a win of 55p. Blocks with high levels of either volatility or noise have more extreme magnitudes than blocks with low levels of both ([Figure 1b](#)) which will bias any comparison of this metric between blocks. In order to limit the effect of this bias, we estimated the simple choice metric only for trials in which the previous outcome lay in the range of magnitudes common to all four blocks, 35-65.

Reinforcement Learning Model

While the choice metric described above provides a relatively transparent measure of the influence of task outcomes on choice, it does not account for differences in outcome magnitude making it liable to bias. We therefore fitted a simple reinforcement learning model to measure block-wise learning rates, which provide a more principled estimate of the degree to which choices are influenced by outcomes. The model combines a learning phase in which the magnitude of wins and losses associated with a shape are estimated (note that it is not necessary to learn the magnitudes associated with the other shape, as these are simply 1-those described below)

$$Q_{win_a(t+1)} = Q_{win_a(t)} + \alpha_{win}(win_{(t)} - Q_{win_a(t)})$$

$$Q_{loss_a(t+1)} = Q_{loss_a(t)} + \alpha_{loss}(loss_{(t)} - Q_{loss_a(t)})$$

In these equations, $Q_{win_a(t)}$ and $Q_{loss_a(t)}$ are the estimated win and loss magnitudes associated with Shape A on trial t , $win_{(t)}$ and $loss_{(t)}$ are the observed win and loss outcome magnitudes and α_{win} and α_{loss} are the win and loss learning rates. These values are then combined in a decision phase such that:

$$P_{choice_a(t)} = \frac{1}{1 + e^{-\beta(Q_{win_a(t)} - Q_{loss_a(t)})}}$$

Where $P_{choice_a(t)}$ is the probability that Shape A will be chosen on trial t and β is a single inverse decision temperature. This model was initiated with $Q_{win_a(0)} = Q_{loss_a(0)} = 0.5$ and the three free parameters ($win_{(t)}$, $loss_{(t)}$ and β) were estimated for each block and each participant by calculating the joint posterior probability given participant choice, marginalising each parameter and deriving the parameters' expected values (Behrens et al., 2007; Browning et al., 2015). See supplementary materials for model selection data.

Some analyses reported in the paper (i.e. where trials are labelled as high/low volatility and high/low noise by the Bayesian Observer Model rather than by task block) cannot be modelled using this block-wise approach (as different types of trial are interleaved throughout the task, rather than blocked). In these analyses a similar, single model was fit across all trials in the task. This model had 8 different learning rates (separate win and loss learning rates, for each combination of high/low volatility and high/low noise labelled trials) and a single inverse temperature parameter. Although this model is somewhat less flexible than the blockwise

modelling approach (i.e. it has 8, rather than 10 learning rates, and 1 rather than 5 inverse temperatures), it produces the same pattern of results when applied to participant choices split by task block (all estimated learning rates correlate at $r > 0.8$, **Figure 2c-d** show results from blockwise fitting, **Figure 3e** from the simpler model). This simpler model was fit using stan, with 5000 burn in and 5000 estimation trials, with posterior convergence visually checked and rhat values of less than 1.1 accepted.

Note that neither of these models describe how participants adjust to different levels of volatility and noise, they simply estimate the learning rates used in each block/type of trial, which are expected to vary in response to differences in levels of uncertainty (in contrast, the Bayesian Observer Model described below does estimate uncertainty and adjust to levels of uncertainty).

Bayesian Observer Model

A recursive, grid-based Bayesian Observer Model (BOM) was developed, similar to that described by Behrens and colleagues (Behrens et al., 2007; Pulcu et al., 2022). The BOM is based on a generative process (see **Figure 3**), and described fully in Pulcu et al. (Pulcu et al., 2022). Below we summarise the key aspects of the model.

The BOM assumes that the observed outcomes at a given time point t , y_t , are generated from a Gaussian distribution with an unknown mean, μ_t , and standard deviation, e^{SD_t} , with the later producing noise in the observed outcomes (**Figure 1b-c**).

$$y_t \sim N(\mu_t, e^{SD_t})$$

As illustrated in **Figure 1b-c**, the mean of this distribution may change between time points, leading to volatility in the task environment, with this change described by a second level Gaussian distribution, centered on the current mean and with a standard deviation of e^{vmut} . The mean of the generative Gaussian distribution in the following trial is drawn from:

$$P(\mu_{t+1}) \sim N(\mu_t, e^{vmut})$$

Both the noise (SD_t) and volatility ($vmut_t$) parameters can also change between time points with their change governed by Gaussian distributions centered on their current value with standard deviations of e^{vSD} and e^{kmu} respectively. These higher-level parameters allow the model to account for periods in which noise and volatility are high and other periods in which they are low (for example, as caused by the uncertainty changes between task blocks).

$$P(vmut_{t+1}) \sim N(vmut_t, e^{kmu})$$

$$P(SD_{t+1}) \sim N(SD_t, e^{vSD})$$

The BOM estimates the joint posterior probability of the five causal parameters, given the choice outcome it has observed. The joint probability distribution at time point t is defined as:

$$P(joint_t) = P(mu, vmut, kmu, SD, vSD | y_{t-1}, y_{t-2}, \dots, y_1)$$

This joint probability distribution can be thought of as the BOM's belief about the values of each parameter in the generative model. A Markovian assumption (i.e. that nodes of the model are sufficient to describe the generative process) simplifies this process and illustrates the recursive update performed by the BOM:

$$P(joint_t) = P(mu_t, vmut_t, kmu_t, SD_t, vSD_t | joint_{t-1}, y_{t-1})$$

We initialized the joint posterior, before observation of any task outcomes as a uniform distribution. The BOM performs the update, first using Bayes' rule to incorporate the effect of the most recently observed outcome, and then accounts for the drifting parameters by using the conditional probability of the new value of the drifting parameter, given the initial value and drift rate (See; Pulcu et al., 2022 [for a detailed account of this updating process](#)):

$$p(\text{joint}_t | \text{joint}_{t-1}, y_{t-1}) = \iiint p(\text{joint}_{t-1} | y_{t-1}) p(SD_t | SD_{t-1}, vSD) p(vmu_t | vmu_{t-1}, kmu) \dots p(mu_t | mu_{t-1}, vmu_t), dSD_{t-1}, dvmu_{t-1}, dmu_{t-1}$$

The value of each node is derived at every time point by marginalizing over all but the relevant dimension of the joint probability distribution and calculating the expected value of that dimension.

During the task, the shapes presented to participants change between each task block, which means that, at the start of each block, participants have to relearn the mean associated with each shape. This was dealt with in the BOM by flattening the μ dimension of the joint probability distribution at the start of each trial (i.e. replacing the values of the mean dimension, with the average of the joint distribution across this dimension). The effect of this is to reset the model's belief about the actual magnitude associated with the two new shapes, while maintaining its belief about the overall volatility and noise of the outcomes.

The BOM was provided with the win and loss outcomes (as values between 0 and 1) for each trial, across all trials in the task (excluding the first practice block, although including this did not alter the pattern of results). It treated the two outcomes as independent (i.e. the win outcome did not influence estimates for the loss outcome and vice versa) and transformed the outcomes to the infinite real line using the logistic transform before estimating the posterior probability (Pulcu et al., 2022 [for a detailed account of this updating process](#)).

Lesioning the Bayesian Observer Model

A number of different lesions were applied to the BOM. First, its ability to estimate changes in either volatility or noise was removed. This was achieved simply by removing the kmu or vSD nodes from the BOM (reducing the dimensionality of the joint distribution by one in each case). The effect of this is to force the BOM to estimate the mean volatility and noise (respectively) across the whole task, rather than to modify its estimates of these parameters between trials.

The second approach induced a graded, rather than absolute, lesion. This was achieved by reducing the precision with which the BOM represented the volatility-related nodes (vmu and kmu) and/or the noise related nodes (SD and vSD). More specifically, the BOM's estimates of the values of each of the five nodes are encoded on a five dimensional grid, with each dimension on the grid representing the possible range of values of a particular node, from low to high, using a fixed number of points. The probability ascribed by the model to a specific point on this dimension is the relative probability that the value of the node lies within the bin of values that is closer to the point, than to adjacent points. For example, say the value of volatility (vmu) ranged from 0 to 10 and was represented by 10 bins. In this case volatility would be represented by a probability mass function over the 10 bins ($<0.5, 0.5-1.5, 1.5-2.5, \dots, > 9.5$). Lesioning occurred by independently varying the number of bins used in the volatility-related and/or noise-related dimensions, from a maximum of 20, to a minimum of 2 (i.e. with only 2 bins volatility/noise would be represented as simply "high" or "low"). The degree of lesioning selected for each individual participant was determined as the number of bins for the volatility and noise dimensions that, after passing the model estimates through a softmax action selector with a single inverse temperature parameter (i.e. as described for the RL model), maximized the likelihood that the model would make the same


choices as the participant, across all task blocks. This process of lesioning therefore progressively coarsens the BOM's representation of the two types of uncertainty and selects the degree of coarsening that results in choices as similar as possible to participants (see supplementary materials for an alternative model that coarsens the representation of the mean values).

Pupillometry Data Preprocessing

Pupillometry data were collected using the Eyelink II system (SRresearch) from both eyes, sampled at 500Hz. Preprocessing involved the following steps: Eye blinks were identified using the built in filter of the Eyelink system and were removed from the data. A linear interpolation was implemented for all missing data points (including blinks). The resulting trace was subjected to a low pass Butterworth filter (cut-off of 3.75 Hz), z transformed across the session (Browning et al., 2015 [↗](#); Nassar et al., 2012 [↗](#)), and then averaged across the two eyes. The pupil response to the win and the loss outcomes were extracted separately from each trial, using a time window based on the presentation of the outcomes. This included a 2-s pre-outcome period, and a 6-s period following outcome presentation. Individual trials were excluded from the pupillometry analysis if more than 50% of the data from the outcome period had been interpolated (mean =6.7% of trials) (Browning et al., 2015 [↗](#)). The first 5 trials from each block were not used in the analysis as initial pupil adaption can occur in response to luminance changes in this period (Browning et al., 2015 [↗](#); Nassar et al., 2012 [↗](#)). The preprocessing resulted in two sets of timeseries per participant, one set containing pupil size data for each included trial when the win outcomes were displayed and the other when the loss outcomes were displayed. These pupil area data were binned into one second bins across the outcome period for analysis (NB **Figure 5a-f** [↗](#)). This analysis was supplemented by an individual regression approach (**Figure 5g-i** [↗](#)) in which individual participants' pupil area timeseries was first regressed against estimated trialwise volatility and noise from the intact BOM (**Figure 5g** [↗](#)), as well as a number of control variables (constant term, amount won/lost on trial (i.e. magnitude of outcome), valence of outcome (win or loss), order in which outcomes were presented (win first/loss first), trial number (1:360), whether shape chosen switched on next trial or not (1:0)). The residuals from this regression were then regressed against estimated trialwise volatility and noise from the degraded BOM (**Figure 5h,i** [↗](#)). These regression analyses resulted in timeseries of beta-weights that were analysed in the same manner as raw pupil size data.

Quantification and statistical analysis

Behavioural data were analysed using linear mixed effect models (*fitlme* function of Matlab (2022a)) with participant ID included as a random factor and volatility, noise and valence added as fixed factors. Two-way interactions between fixed effects were also tested (main effects are reported from models without interaction terms). Addition of random slopes for any of the fixed factors decreased LME model fit statistics and so were not included (Matuschek et al., 2017 [↗](#)). Analysis of timeseries pupillometry data included the additional fixed effect factor of time across the outcome period. Learning rates were transformed to the infinite real line using a logistic transform before analyses (untransformed data are displayed in figures for ease of interpretation). The normality of the distribution of the residuals of the LME analyses were checked both visually and with a one-sample Kolmogorov-Smirnov test. Changes in the classification of trials between the full and degraded BOM (**Figure 4b** [↗](#)) were analysed using a repeated measures ANOVA with within subject factors of volatility, noise and valence. Raw data are superimposed on all summary figures.



Code and Data Availability

Study data and analysis scripts, including code for the various models used are available at: <https://osf.io/j7md3/>.

Acknowledgements

We would like to thank James Gunnell for help in collecting the data. This study was funded by a MRC Clinician Scientist Fellowship awarded to MB (MR/N008103/1). MB was supported by the Oxford Health NIHR Biomedical Research Centre. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health.

Additional information

Author Contributions

MB and EP conceived the study, EP collected the data, MB and EP wrote the paper.

Disclosures

MB has received travel expenses from Lundbeck for attending conferences and consultancy from Jansen, CHDR and Novartis. EP declares no potential conflict of interest.

Additional files

Supplementary Materials

References

- Behrens T. E. J., Hunt L. T., Woolrich M. W., Rushworth M. F. (2008) **Associative learning of social value** *Nature* **456**:245–249 <https://doi.org/10.1038/nature07538>
- Behrens T. E. J., Woolrich M. W., Walton M. E., Rushworth M. F. S. (2007) **Learning the value of information in an uncertain world** *Nature Neuroscience* **10**:1214–1221 <https://doi.org/10.1038/nn1954>
- Browning M., Behrens T. E., Jocham G., O'Reilly J. X., Bishop S. J. (2015) **Anxious individuals have difficulty learning the causal statistics of aversive environments** *Nature Neuroscience* **18**:590–596 <https://doi.org/10.1038/nn.3961>
- de Gee J. W., Colizoli O., Kloosterman N. A., Knapen T., Nieuwenhuis S., Donner T. H. (2017) **Dynamic modulation of decision biases by brainstem arousal systems** *ELife* **6** <https://doi.org/10.7554/eLife.23232>
- Diederen K. M. J., Schultz W. (2015) **Scaling prediction errors to reward variability benefits error-driven learning in humans** *Journal of Neurophysiology* **114**:1628–1640 <https://doi.org/10.1152/jn.00483.2015>
- Gagne C., Zika O., Dayan P., Bishop S. J. (2020) **Impaired adaptation of learning to contingency volatility in internalizing psychopathology** *ELife* **9** <https://doi.org/10.7554/eLife.61387>
- Herce Castañón S., Moran R., Ding J., Egner T., Bang D., Summerfield C. (2019) **Human noise blindness drives suboptimal cognitive inference** *Nature Communications* **10** <https://doi.org/10.1038/s41467-019-09330-7>
- Joshi S., Li Y., Kalwani R. M., Gold J. I. (2016) **Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex** *Neuron* **89**:221–234 <https://doi.org/10.1016/j.neuron.2015.11.028>
- Kalman R. E. (1960) **A New Approach to Linear Filtering and Prediction Problem** *Transactions of the ASME* **82**:35–45
- Krishnamurthy K., Nassar M. R., Sarode S., Gold J. I. (2017) **Arousal-related adjustments of perceptual biases optimize perception in dynamic environments** *Nature Human Behaviour* **1** <https://doi.org/10.1038/s41562-017-0107>
- Mackay D. J. C. (2003) **Information Theory, Inference, and Learning Algorithms** Cambridge University Press
- Matuschek H., Kliegl R., Vasishth S., Baayen H., Bates D. (2017) **Balancing Type I error and power in linear mixed models** *Journal of Memory and Language* **94**:305–315 <https://doi.org/10.1016/j.jml.2017.01.001>
- Nassar M. R., Rumsey K. M., Wilson R. C., Parikh K., Heasley B., Gold J. I. (2012) **Rational regulation of learning dynamics by pupil-linked arousal systems** *Nat Neurosci* **15**:1040–1046 <https://doi.org/10.1038/nn.3130>

- Nassar M. R., Wilson R. C., Heasly B., Gold J. I. (2010) **An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment** *Journal of Neuroscience* **30**:12366–12378 <https://doi.org/10.1523/JNEUROSCI.0822-10.2010>
- O'Reilly J. X., Schüfflgen U., Cuell S. F., Behrens T. E. J., Mars R. B., Rushworth M. F. S. (2013) **Dissociable effects of surprise and model update in parietal and anterior cingulate cortex** *Proceedings of the National Academy of Sciences of the United States of America* **110**:E3660–3669 <https://doi.org/10.1073/pnas.1305373110>
- Piray P., Daw N. D. (2021) **A model for learning based on the joint estimation of stochasticity and volatility** *Nature Communications* **12** <https://doi.org/10.1038/s41467-021-26731-9>
- Preuschoff K., Hart B. M., Einhäuser W. (2011) **Pupil Dilation Signals Surprise: Evidence for Noradrenaline's Role in Decision Making** *Frontiers in Neuroscience* **5** <https://doi.org/10.3389/fnins.2011.00115>
- Pulcu E., Browning M. (2017) **Affective bias as a rational response to the statistics of rewards and punishments** *eLife* **6** <https://doi.org/10.7554/eLife.27879>
- Pulcu E., Browning M. (2019) **The Misestimation of Uncertainty in Affective Disorders** *Trends in Cognitive Sciences* <https://doi.org/10.1016/j.tics.2019.07.007>
- Pulcu E., Saunders K. E. A., Harmer C. J., Harrison P. J., Goodwin G. M., Geddes J. R., Browning M. (2022) **Using a generative model of affect to characterize affective variability and its response to treatment in bipolar disorder** *Proceedings of the National Academy of Sciences of the United States of America* **119** <https://doi.org/10.1073/pnas.2202983119>
- Sutton R. S., Barto A. G. (2018) **Reinforcement Learning: An Introduction** MIT Press
- Yu A. J., Dayan P. (2005) **Uncertainty, neuromodulation, and attention** *Neuron* **46**:681–692 <https://doi.org/10.1016/j.neuron.2005.04.026>

Author information

Erdem Pulcu

Department of Psychiatry, University of Oxford, Oxford, UK
ORCID iD: [0000-0002-2170-0677](https://orcid.org/0000-0002-2170-0677)

For correspondence: erdem.pulcu@psych.ox.ac.uk

Michael Browning

Department of Psychiatry, University of Oxford, Oxford, UK, Oxford Health NHS Trust, Oxford, UK

Editors

Reviewing Editor

Andreea Diaconescu

University of Toronto, Toronto, Canada

Senior Editor

Michael Frank

Brown University, Providence, United States of America

Reviewer #1 (Public review):

Summary:

The authors present an interesting study using RL and Bayesian modelling to examine differences in learning rate adaptation in conditions of high and low volatility and noise respectively. Through "lesioning" an optimal Bayesian model, they reveal that apparently a suboptimal adaptation of learning rates results from incorrectly detecting volatility in the environment when it is not in fact present.

Strengths:

The experimental task used is cleverly designed and does a good job of manipulating both volatility and noise. The modelling approach takes an interesting and creative approach to understanding the source of apparently suboptimal adaptation of learning rates to noise, through carefully "lesioning" an optimal Bayesian model to determine which components are responsible for this behaviour.

Weaknesses:

The study has a few substantial weaknesses; the data and modelling both appear robust and informative, and it tackles an interesting question. The model space could potentially have been expanded, particularly with regard to the inclusion of alternative strategies such as those that estimate latent states and adapt learning accordingly.

<https://doi.org/10.7554/eLife.103734.1.sa2>

Reviewer #2 (Public review):

Summary:

In this study, the authors aimed to investigate how humans learn and adapt their behavior in dynamic environments characterized by two distinct types of uncertainty: volatility (systematic changes in outcomes) and noise (random variability in outcomes). Specifically, they sought to understand how participants adjust their learning rates in response to changes in these forms of uncertainty.

To achieve this, the authors employed a two-step approach:

(1) Reinforcement Learning (RL) Model: They first used an RL model to fit participants' behavior, revealing that the learning rate was context-dependent. In other words, it varied based on the levels of volatility and noise. However, the RL model showed that participants misattributed noise as volatility, leading to higher learning rates in noisy conditions, where the optimal strategy would be to be less sensitive to random fluctuations.

(2) Bayesian Observer Model (BOM): To better account for this context dependency, they introduced a Bayesian Observer Model (BOM), which models how an ideal Bayesian learner would update their beliefs about environmental uncertainty. They found that a degraded version of the BOM, where the agent had a coarser representation of noise compared to volatility, best fit the participants' behavior. This suggested that participants were not fully distinguishing between noise and volatility, instead treating noise as volatility and adjusting their learning rates accordingly.

The authors also aimed to use pupillometry data (measuring pupil dilation) as a physiological marker to arbitrate between models and understand how participants' internal representations of uncertainty influenced both their behavior and physiological responses. Their objective was to explore whether the BOM could explain not just behavioral choices but also these physiological responses, thereby providing stronger evidence for the model's validity.

Overall, the study sought to reconcile approximate rationality in human learning by showing that participants still follow a Bayesian-like learning process, but with simplified internal models that lead to suboptimal decisions in noisy environments.

Strengths:

The generative model presented in the study is both innovative and insightful. The authors first employ a Reinforcement Learning (RL) model to fit participants' behavior, revealing that the learning rate is context-dependent—specifically, it varies based on the levels of volatility and noise in the task. They then introduce a Bayesian Observer Model (BOM) to account for this context dependency, ultimately finding that a degraded BOM—in which the agent has a coarser representation of noise compared to volatility—provides the best fit for the participants' behavior. This suggests that participants do not fully distinguish between noise and volatility, leading to the misattribution of noise as volatility. Consequently, participants adopt higher learning rates even in noisy contexts, where an optimal strategy would involve being less sensitive to new information (i.e., using lower learning rates). This finding highlights a rational but approximate learning process, as described in the paper.

Weaknesses:

While the RL and Bayesian models both successfully predict behavior, it remains unclear how to fully reconcile the two approaches. The RL model captures behavior in terms of a fixed or context-dependent learning rate, while the BOM provides a more nuanced account with dynamic updates based on volatility and noise. Both models can predict actions when fit appropriately, but the pupillometry data offers a promising avenue to arbitrate between the models. However, the current study does not provide a direct comparison between the RL framework and the Bayesian model in terms of how well they explain the pupillometry data. It would be valuable to see whether the RL model can also account for physiological markers of learning, such as pupil responses, or if the BOM offers a unique advantage in this regard. A comparison of the two models using pupillometry data could strengthen the argument for the BOM's superiority, as currently, the possibility that RL models could explain the physiological data remains unexplored.

The model comparison between the Bayesian Observer Model and the self-defined degraded internal model could be further enhanced. Since different assumptions about the internal model's structure lead to varying levels of model complexity, using a formal criterion such as Bayesian Information Criterion (BIC) or Akaike Information Criterion (AIC) would allow for a more rigorous comparison of model fit. Including such comparisons would ensure that the degraded BOM is not simply favored due to its flexibility or higher complexity, but rather because it genuinely captures the participants' behavioral and physiological data better than alternative models. This would also help address concerns about overfitting and provide a clearer justification for using the degraded BOM over other potential models.

<https://doi.org/10.7554/eLife.103734.1.sa1>

Author response:**Reviewer #1 (Public review):***Summary:*

The authors present an interesting study using RL and Bayesian modelling to examine differences in learning rate adaptation in conditions of high and low volatility and noise respectively. Through "lesioning" an optimal Bayesian model, they reveal that apparently a suboptimal adaptation of learning rates results from incorrectly detecting volatility in the environment when it is not in fact present.

Strengths:

The experimental task used is cleverly designed and does a good job of manipulating both volatility and noise. The modelling approach takes an interesting and creative approach to understanding the source of apparently suboptimal adaptation of learning rates to noise, through carefully "lesioning" an optimal Bayesian model to determine which components are responsible for this behaviour.

We thank the reviewer for this assessment.

Weaknesses:

The study has a few substantial weaknesses; the data and modelling both appear robust and informative, and it tackles an interesting question. The model space could potentially have been expanded, particularly with regard to the inclusion of alternative strategies such as those that estimate latent states and adapt learning accordingly.

We thank the reviewer for this suggestion. We agree that it would be interesting to assess the ability of alternative models to reproduce the sub-optimal choices of participants in this study. The Bayesian Observer Model described in the paper is a form of Hierarchical Gaussian Filter, so we will assess the performance of a different class of models that are able to track uncertainty-- RL based models that are able to capture changes of uncertainty (the Kalman filter, and the model described by Cochran and Cisler, *Plos Comp Biol* 2019). We will assess the ability of the models to recapitulate the core behaviour of participants (in terms of learning rate adaption) and, if possible, assess their ability to account for the pupillometry response.

Reviewer #2 (Public review):*Summary:*

In this study, the authors aimed to investigate how humans learn and adapt their behavior in dynamic environments characterized by two distinct types of uncertainty: volatility (systematic changes in outcomes) and noise (random variability in outcomes). Specifically, they sought to understand how participants adjust their learning rates in response to changes in these forms of uncertainty.

To achieve this, the authors employed a two-step approach:

(1) Reinforcement Learning (RL) Model: They first used an RL model to fit participants' behavior, revealing that the learning rate was context-dependent. In other words, it varied based on the levels of volatility and noise. However, the RL model showed that participants misattributed noise as volatility, leading to higher learning rates in noisy

conditions, where the optimal strategy would be to be less sensitive to random fluctuations.

(2) Bayesian Observer Model (BOM): To better account for this context dependency, they introduced a Bayesian Observer Model (BOM), which models how an ideal Bayesian learner would update their beliefs about environmental uncertainty. They found that a degraded version of the BOM, where the agent had a coarser representation of noise compared to volatility, best fit the participants' behavior. This suggested that participants were not fully distinguishing between noise and volatility, instead treating noise as volatility and adjusting their learning rates accordingly.

The authors also aimed to use pupillometry data (measuring pupil dilation) as a physiological marker to arbitrate between models and understand how participants' internal representations of uncertainty influenced both their behavior and physiological responses. Their objective was to explore whether the BOM could explain not just behavioral choices but also these physiological responses, thereby providing stronger evidence for the model's validity.

Overall, the study sought to reconcile approximate rationality in human learning by showing that participants still follow a Bayesian-like learning process, but with simplified internal models that lead to suboptimal decisions in noisy environments.

Strengths:

The generative model presented in the study is both innovative and insightful. The authors first employ a Reinforcement Learning (RL) model to fit participants' behavior, revealing that the learning rate is context-dependent-specifically, it varies based on the levels of volatility and noise in the task. They then introduce a Bayesian Observer Model (BOM) to account for this context dependency, ultimately finding that a degraded BOM - in which the agent has a coarser representation of noise compared to volatility - provides the best fit for the participants' behavior. This suggests that participants do not fully distinguish between noise and volatility, leading to the misattribution of noise as volatility. Consequently, participants adopt higher learning rates even in noisy contexts, where an optimal strategy would involve being less sensitive to new information (i.e., using lower learning rates). This finding highlights a rational but approximate learning process, as described in the paper.

We thank the reviewer for their assessment of the paper.

Weaknesses:

While the RL and Bayesian models both successfully predict behavior, it remains unclear how to fully reconcile the two approaches. The RL model captures behavior in terms of a fixed or context-dependent learning rate, while the BOM provides a more nuanced account with dynamic updates based on volatility and noise. Both models can predict actions when fit appropriately, but the pupillometry data offers a promising avenue to arbitrate between the models. However, the current study does not provide a direct comparison between the RL framework and the Bayesian model in terms of how well they explain the pupillometry data. It would be valuable to see whether the RL model can also account for physiological markers of learning, such as pupil responses, or if the BOM offers a unique advantage in this regard. A comparison of the two models using pupillometry data could strengthen the argument for the BOM's superiority, as currently, the possibility that RL models could explain the physiological data remains unexplored.

We thank the reviewer for this suggestion. In the current version of the paper, we use an extremely simple reinforcement learning model to simply measure the learning rate in each

task block (as this is the key behavioural metric we are interested in). As the reviewer highlights, this simple model doesn't estimate uncertainty or adapt to it. Given this, we don't think we can directly compare this model to the Bayesian Observer Model—for example, in the current analysis of the pupillometry data we classify individual trials based on the BOM's estimate of uncertainty and show that participants adapt their learning rate as expected to the reclassified trials, this analysis would not be possible with our current RL model. However, there are more complex RL based models that do estimate uncertainty (as discussed above in response to Reviewer #1) and so may more directly be compared to the BOM. We will attempt to apply these models to our task data and describe their ability to account for participant behaviour and physiological response as suggested by the Reviewer.

The model comparison between the Bayesian Observer Model and the self-defined degraded internal model could be further enhanced. Since different assumptions about the internal model's structure lead to varying levels of model complexity, using a formal criterion such as Bayesian Information Criterion (BIC) or Akaike Information Criterion (AIC) would allow for a more rigorous comparison of model fit. Including such comparisons would ensure that the degraded BOM is not simply favored due to its flexibility or higher complexity, but rather because it genuinely captures the participants' behavioral and physiological data better than alternative models. This would also help address concerns about overfitting and provide a clearer justification for using the degraded BOM over other potential models.

Thank you, we will add this.

<https://doi.org/10.7554/eLife.103734.1.sa0>