

RESEARCH ARTICLE

Pharmacological Fingerprints of Contextual Uncertainty

Louise Marshall¹*, Christoph Mathys^{2,3}, Diane Ruge^{1,4}, Archy O. de Berker^{1,2}, Peter Dayan⁵, Klaas E. Stephan^{2,6,7}, Sven Bestmann¹

1 Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London, London, United Kingdom, **2** Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, London, United Kingdom, **3** Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London, United Kingdom, **4** Department of Psychology and Neurosciences, Leibniz Research Centre for Working Environment and Human Factors, Technical University Dortmund, Dortmund, Germany, **5** Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom, **6** Translational Neuromodeling Unit, Institute for Biomedical Engineering, University of Zurich & ETH Zurich, Zurich, Switzerland, **7** Max Planck Institute for Metabolism Research, Cologne, Germany

* These authors contributed equally to this work.

* louise.marshall@ucl.ac.uk



OPEN ACCESS

Citation: Marshall L, Mathys C, Ruge D, de Berker AO, Dayan P, Stephan KE, et al. (2016) Pharmacological Fingerprints of Contextual Uncertainty. *PLoS Biol* 14(11): e1002575. doi:10.1371/journal.pbio.1002575

Academic Editor: Matthew F. S. Rushworth, Oxford University, UNITED KINGDOM

Received: April 11, 2016

Accepted: October 5, 2016

Published: November 15, 2016

Copyright: © 2016 Marshall et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All data files are available from figshare.com (<http://dx.doi.org/10.6084/m9.figshare.3168682.v1>, <http://dx.doi.org/10.6084/m9.figshare.3796314.v1>, <http://dx.doi.org/10.6084/m9.figshare.3793410.v1>, <http://dx.doi.org/10.6084/m9.figshare.3796353.v1>, <http://dx.doi.org/10.6084/m9.figshare.3796356.v1>, <http://dx.doi.org/10.6084/m9.figshare.3796362.v2>, <http://dx.doi.org/10.6084/m9.figshare.3796401.v2>, <http://dx.doi.org/10.6084/m9.figshare.3796407.v1>, <http://dx.doi.org/10.6084/m9.figshare.3796413.v1>, <http://dx.doi.org/10.6084/m9.figshare.3796434.v1>, <http://dx.doi.org/10.6084/m9.figshare.3796443.v1>,

Abstract

Successful interaction with the environment requires flexible updating of our beliefs about the world. By estimating the likelihood of future events, it is possible to prepare appropriate actions in advance and execute fast, accurate motor responses. According to theoretical proposals, agents track the variability arising from changing environments by computing various forms of uncertainty. Several neuromodulators have been linked to uncertainty signalling, but comprehensive empirical characterisation of their relative contributions to perceptual belief updating, and to the selection of motor responses, is lacking. Here we assess the roles of noradrenaline, acetylcholine, and dopamine within a single, unified computational framework of uncertainty. Using pharmacological interventions in a sample of 128 healthy human volunteers and a hierarchical Bayesian learning model, we characterise the influences of noradrenergic, cholinergic, and dopaminergic receptor antagonism on individual computations of uncertainty during a probabilistic serial reaction time task. We propose that noradrenaline influences learning of uncertain events arising from unexpected changes in the environment. In contrast, acetylcholine balances attribution of uncertainty to chance fluctuations within an environmental context, defined by a stable set of probabilistic associations, or to gross environmental violations following a contextual switch. Dopamine supports the use of uncertainty representations to engender fast, adaptive responses.

Author Summary

Interacting with dynamic and ever-changing environments requires frequent updating of our beliefs about the world. By learning the relationships that link events in the current environmental context, it is possible to prepare and execute fast, accurate responses to

<http://dx.doi.org/10.6084/m9.figshare.3796446.v1>,
<http://dx.doi.org/10.6084/m9.figshare.3796449.v2>,
<http://dx.doi.org/10.6084/m9.figshare.3796452.v2>,
<http://dx.doi.org/10.6084/m9.figshare.3796407.v1>). Modelling scripts are available from <http://www.translationalneuromodeling.org/tapas/>.
 Matlab code can be downloaded directly from <http://www.tnu.ethz.ch/de/software/tapas/download.html>.

Funding: This work was supported by the Medical Research Council (LM and AODB), the European Research Council, grant number ActSelectContext 260424 (SB), the Gatsby Charitable Foundation (PD), the Clinical Research Priority Program “Molecular Imaging” (KES), the René and Susanne Braginsky Foundation (KES), the Tourette Syndrome Association (DR), a Dorothy Feiss Scientific Research Grant (DR), and core funding from the Wellcome Trust Centre for Neuroimaging (CM) and the University of Zurich (KES). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

Abbreviations: ACh, acetylcholine; BIS-11, Barratt Impulsiveness Scale; BOLD, blood-oxygen-level-dependent; BP, blood pressure; BPA, Bayesian parameter average; CFQ, Cognitive Failures Questionnaire; DA, dopamine; DOSPERT, Domain-Specific Risk Taking; HGF, Hierarchical Gaussian Filter; HR, heart rate; ITI, intertrial interval; LC, locus coeruleus; NA, noradrenaline; ns, nonsignificant; PE, prediction error; RM-ANOVA, repeated-measures analysis of variance; RT, reaction time; SD, standard deviation; SED, standard error of the difference; SEM, standard error of the mean; TM, transition matrix; VAS, visual analogue scale; YD, Yu and Dayan.

those events that are predictable. However, the world’s complex dynamics give rise to uncertainty about the relationships that exist between events and uncertainty about how these relationships might change over time. Several neuromodulators have been proposed to signal these different forms of uncertainty, but their relative contributions to updating beliefs and modulating responses have remained elusive. Here we combine a probabilistic reaction time task, pharmacological interventions, and a hierarchical Bayesian learning model to identify the roles of noradrenaline, acetylcholine, and dopamine in individual computations of uncertainty. We propose that noradrenaline modulates learning about the instability of the relationships that link environmental events. Acetylcholine balances the attribution of uncertainty to unexpected events occurring within an environmental context or to gross violations of our expectations following a context change. In contrast, dopamine sensitises our actions to our beliefs about uncertainty.

Introduction

Adaptive performance in dynamic environments depends on our ability to represent and manipulate internal estimates of the world’s statistical structure [1–4]. By tracking the environment’s underlying regularities, an agent can learn the causes of its sensory input and thus the likelihood that a particular event will occur. In turn, this permits anticipatory action preparation and the rapid execution of responses [5].

However, the environment’s richly complicated sources of noise and latent structure present us with various forms of uncertainty. For instance, a London commuter predicting her journey time to work faces three distinct forms. First, there is irreducible uncertainty, which captures the randomness inherent in any complex environment and is undiminished by learning; an unplanned station closure could cause journey delays and thus influence the accuracy of the commuter’s estimated arrival time on any given day. Second, after moving to a new part of town, the duration of the commuter’s chosen route to work may be unclear, producing uncertainty about how likely she is to arrive at work on time. Over repeated journeys, this estimation uncertainty falls as the commuter learns about the probabilistic relationships within the current environmental context. For example, she learns to predict the frequent delays on this new route due to congestion during the morning rush hour, although these delays may vary with local dips and surges in the number of passengers using the service. Volatility uncertainty arises from our beliefs about the stability of the environment and thus how quickly probabilistic relationships are changing between contexts. A major sporting event, such as the London Olympics, may bring a large influx of additional passengers for an unknown period of time and with unexpected effects on transport performance, making it harder to predict future journey times until these changes have been taken into account.

The Brain Computes Different Forms of Uncertainty

An assortment of theoretical, behavioural, and neurobiological evidence suggests that the brain computes uncertainty estimates relating to the environment’s sensory events, contextual associations, and their changes over time [6–17]. Uncertainty about the validity of one’s own perceptual beliefs about the world should have the effect of suppressing top-down prior expectations relative to new bottom-up sensory evidence, promoting learning about the current environmental context [18]. With their broad distribution and extensive connectivity, neuromodulatory networks are well placed to facilitate the widespread changes in gain required to support such a function [19,20]. In particular, acetylcholine (ACh) and noradrenaline (NA)

are known to enhance bottom-up, feedforward thalamocortical transmission of sensory information relative to top-down, intracortical, and feedback processing [4,21–30].

Proposed roles for NA and ACh in computations of uncertainty. Within a stable environmental context, uncertainty arises from ignorance about, and the unreliability of, probabilistic cues predicting upcoming sensory events. Learning the environment's underlying statistical regularities means that responses to explicitly and predictably cued events are typically faster than to those believed improbable [31–34]. This effect is modulated by pharmacological [35,36], surgical [37,38], and neurodegenerative [39] manipulations of ACh. For example, human cholinergic basal forebrain activity has been shown to reflect an individual's uncertainty about probabilistic cue-outcome relationships [15]. More specifically, ACh appears to increase the rate at which humans learn probabilistic relationships under estimation uncertainty, supporting the idea that ACh enhances learning accorded to stimuli with uncertain predictive consequences [40] by suppressing the use of outdated top-down cues and boosting bottom-up sensory processing [4].

While NA plays no consistent role in probabilistic cueing [41,42], it is thought to offer an interrupt signal when volatility uncertainty arises between contexts [41–47]. Learning to make accurate predictions from the strongly unexpected observations that follow a contextual switch necessitates heightened sensory vigilance and a disregard for outdated top-down expectations. NA, with its broad neural network capable of triggering multiple, simultaneous changes across the brain [48], is well placed to rapidly coordinate this process. Indeed, neurons in the locus coeruleus (LC), the primary source of cortical NA, show strong responses to unexpected environmental changes [49,50]. Pharmacologically up-regulating NA accelerates the detection of unexpected switches in the predictive properties of sensory stimuli [51], while noradrenergic deafferentation of rat medial frontal cortex impairs behavioural adaptation to contextual switches [52]. Moreover, blood-oxygen-level-dependent (BOLD) activity in human LC has been shown to dynamically track volatility uncertainty [16], and pupil dilation—which is influenced by (nor)adrenergic afferents [53]—correlates with unexpected changes in probabilistic context [54,55].

Motor Responses Are Sensitive to Uncertainty

Thus, uncertainty representations existing within and between environmental contexts are crucial for optimal predictions about the probability of future events. While good predictions facilitate anticipatory preparation of appropriate behavioural responses [5], they are not sufficient for adaptive performance in dynamic environments. An additional mechanism is required to modify action selection based on one's own beliefs about the latent changes in the environment and/or the occurrence of unexpected events. Indeed, when an unexpected event occurs, humans are capable of engaging resources to inhibit a prepared response and replace it with an alternative [56,57], albeit at the expense of a prolonged reaction time (RT) [58,59].

Proposed role for dopamine in response modulation. There is considerable evidence linking the neuromodulator dopamine (DA) to flexible behaviour [60–64]. Dopaminergic deficits due to Parkinson's disease are associated with specific flexibility impairments in both motor [59,65] and cognitive domains [60,66], with performance restored by dopaminergic medication [59,67]. In healthy individuals, pharmacological DA depletions impair adaptive reactions to unexpected events occurring within a broadly predictable context [58]. However, it remains unclear whether DA supports accurate response selection by facilitating perceptual belief updating [15] or by modulating the sensitivity of response selection to perceptual beliefs.

A Unified Framework of Uncertainty

In sum, while physiological, pharmacological, behavioural, and theoretical work has suggested separable neuromodulatory involvement in different uncertainty computations, attempts to characterise the relative contributions of NA, ACh, and DA within a single computational scheme are lacking (but see noticeable exceptions from Varazzani et al. [68] contrasting the roles of NA and DA in motivation and Brown et al. [69] assessing NA and ACh involvement in orienting responses to environmental novelty). Here we employ a Hierarchical Gaussian Filter (HGF) model [9,10] to track human learning and response modulation in dynamic, probabilistic environments under pharmacological NA, ACh, and DA interventions.

The HGF has been successfully applied in several recent studies of learning under volatility [15,17,70–74]. We developed a novel instantiation of the HGF with two components: a three-level perceptual model of an agent's mapping from environmental causes to sensory inputs and a response model that maps those inferred environmental causes to observed RT responses [75]. Thus, we sought to disentangle the effects of pharmacological interventions of the NA, ACh, and DA neuromodulatory systems on participant-specific perceptual belief updating from those on the sensitivity of motor responses to perceptual estimates.

Results

According to a double-blind, between-subject design, 128 healthy participants undertook a serial probabilistic RT task (Fig 1) after having received a noradrenergic α 1-receptor antagonist (prazosin; NA- group), a cholinergic M1-receptor antagonist (biperiden; ACh- group) or a dopaminergic D1/D2 receptor antagonist (haloperidol; DA- group), or a placebo. Data from 124 participants are reported. Four participants were excluded from analyses, three due to high missed response rates ($\geq 11\%$) and one because the behavioural model parameter estimation (using the Broyden-Fletcher-Goldfarb-Shanno algorithm) did not converge. The four drug groups were matched for gender (Kruskal-Wallis test: $H_3 = 0.53$, $p = 0.912$), age (one-way ANOVA: $F_{3,120} = 0.46$, $p = 0.714$), body weight ($F_{3,120} = 2.24$, $p = 0.087$), education level ($H_3 = 1.31$, $p = 0.727$), and all baseline psychometric measures taken (Table 1).

On each trial, participants were required to respond to the presentation of one of four visual stimuli by making a speeded button press before the end of a 1,200 ms intertrial interval (ITI) (Fig 1). Participants were trained on the stimulus-response mappings, which remained consistent within an experimental session but were counterbalanced across participants. The experimental sequence of 1,200 trials was generated by a hidden probabilistic rule that switched, without explicit indication to the participant, every 50 trials. At any given time, stimulus transitions were generated by one of eight different transition matrices (TMs). Trials were drawn from each TM three times. The order of TMs was pseudorandom, with no consecutive repeats. The overall probability of each stimulus was equal across the 1,200 trials.

This created transient contexts that participants could infer from stimulus observations, allowing them to reduce their uncertainty about events before they occurred [76]. Nonetheless, the probabilistic nature of these contexts also produced unexpected stimulus outcomes, i.e., a sensory prediction error (PE). For fast and accurate responses, participants had to track three forms of uncertainty: irreducible uncertainty arising from the inherent randomness of the probabilistic transitions between consecutive stimuli, estimation uncertainty arising from their imperfect knowledge of the probabilistic relationships governing stimulus transition contingencies within contexts, and volatility uncertainty maintained by the unsignalled contextual instability.

In total, participants responded correctly on $90.3 \pm 0.79\%$ (mean \pm standard error of the mean [SEM]), $88.4 \pm 1.23\%$, $87.7 \pm 1.27\%$, and $89.2 \pm 0.91\%$ of the trials in the Placebo, NA-,

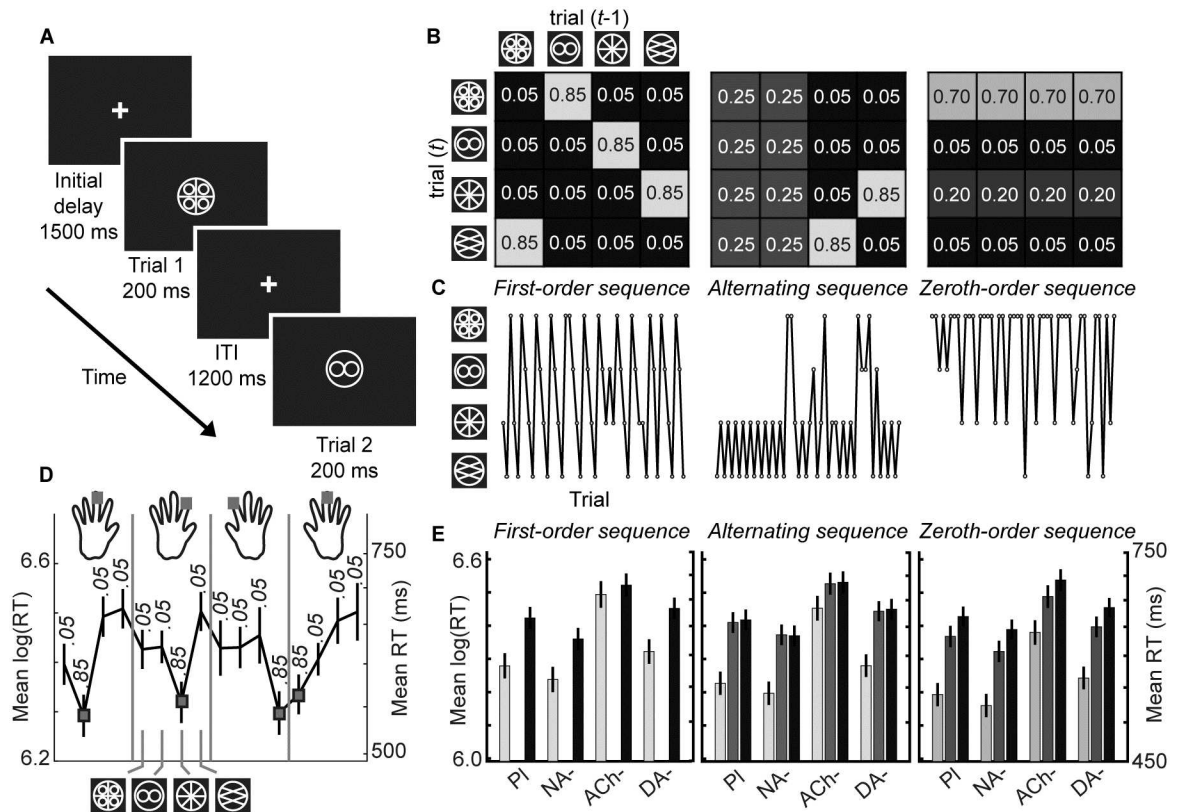


Fig 1. Task design. (A) Trial sequence. A trial began with the presentation of a central white fixation cross against a black background. After an initial delay of 1,500 ms at the start of each block, one of four visual stimuli was presented for 200 ms. Participants were required to make a speeded button-press response before the end of a 1,200 ms intertrial interval (ITI). (B) Stimulus transitions were generated by one of eight different transition matrices (TMs), which changed every 50 trials without explicit indication to the participant. These TMs comprised two different first-order sequences, two alternating sequences, and four zeroth-order sequences, each of which occurred three times in a pseudorandom order across 1,200 trials. The overall probability of each stimulus was equal across the 1,200 trials. For full details, see [S1 Fig](#). (C) Example trial sequences generated from the three example TMs in 1B. (D) By tracking the transition probabilities, participants could learn to predict high probability events and prepare to make the correct button press accordingly. Faster responses were observed for predictable stimuli compared to unexpected stimuli. Here we depict Placebo group log(RTs) (mean \pm standard error of the mean [SEM]) for each of the 16 possible combinations between consecutive stimuli for the first-order sequence shown in B. Grey boxes indicate stimulus combinations with a high transition probability. (E) Indeed, across all types of TM, responses were faster for stimuli with higher transition probabilities (mean \pm SEM). <http://dx.doi.org/10.6084/m9.figshare.3796314.v1>.

doi:10.1371/journal.pbio.1002575.g001

ACh-, and DA- groups, respectively. The percentages of correct responses did not differ between groups ($F_{3,123} = 1.12, p = 0.345$). Since a significant time x drug interaction on self-reported alertness was identified (see the Physiological and Subjective Control Measures results in [S1 Text](#)), the participant-specific difference in alertness between baseline and the time corresponding to peak drug concentration, $\Delta alertness$, was used as a covariate in our analyses to control for any interparticipant variability in subjective drug effect. Since psychopharmacological drugs can have dose-dependent effects [77–79], we also included body weight as an additional covariate in our analyses.

Model-Agnostic Analyses

We first conducted a series of conventional, model-agnostic analyses of behaviour, to assess whether participants learned about the underlying stimulus transition contingencies and whether learning was influenced by our pharmacological interventions ([Fig 2](#)). First, we

Table 1. Participant details for each experimental group.

	Placebo (n = 32)	NA- (n = 31)	ACh- (n = 29)	DA- (n = 32)	Between-Group Difference?
Gender (Number Male) [#]	13	15	14	14	ns, <i>p</i> = 0.912
Age (Years)	23.0 ± 4.6	23.1 ± 4.0	22.0 ± 3.6	22.8 ± 4.5	ns, <i>p</i> = 0.714
Weight (kg)	61.7 ± 1.5	69.0 ± 2.4	64.9 ± 2.5	64.1 ± 1.7	ns, <i>p</i> = 0.087
Education Level (1–5) [#]	2.7 ± 0.2	2.8 ± 0.1	2.6 ± 0.1	2.7 ± 0.1	ns, <i>p</i> = 0.727
Digit Span (Forwards + Backwards) [#]	13.1 ± 0.4	13.0 ± 0.5	12.7 ± 0.4	13.8 ± 0.5	ns, <i>p</i> = 0.252
Impulsivity: BIS-11	61.9 ± 1.5	65.4 ± 1.5	64.7 ± 1.6	63.3 ± 2.0	ns, <i>p</i> = 0.465
Risk Taking: DOSPERT (Total)	104.3 ± 3.0	113.5 ± 3.6	107.7 ± 3.7	105.5 ± 3.5	ns, <i>p</i> = 0.245
Distractibility: CFQ	38.3 ± 1.9	41.7 ± 1.6	43.9 ± 1.8	40.3 ± 1.9	ns, <i>p</i> = 0.185
Sleep Quantity on the Previous Night (Hours) [#]	7.3 ± 0.2	7.2 ± 0.2	6.7 ± 0.3	7.1 ± 0.2	ns, <i>p</i> = 0.513
Sleep Quality on the Previous Night (1–8) [#]	5.5 ± 0.3	5.7 ± 0.2	5.3 ± 0.3	5.4 ± 0.2	ns, <i>p</i> = 0.620
Fatigue during Task (0–100)	44.6 ± 3.9	44.8 ± 3.6	43.5 ± 2.6	41.2 ± 3.7	ns, <i>p</i> = 0.876
Active Drug (%) [#]	50	77	86	44	<i>p</i> = 0.001

Between-group comparisons revealed no significant differences (ns = nonsignificant) for gender, age, body weight, education level, baseline cognitive capacity (Digit Span), impulsivity (Barratt Impulsiveness Scale, BIS-11), risk taking (Domain-Specific Risk-Taking Scale, DOSPERT), distractibility (Cognitive Failures Questionnaire, CFQ), fatigue during the task, or sleep quality or quantity on the previous night. For continuous data, one-way ANOVAs were used to test for any between-group differences. For discrete data ([#]), Kruskal-Wallis tests were applied. Education level refers to the highest attained from the following: 1 = compulsory education (≤ 12 y); 2 = further education (13–14 y); 3 = undergraduate degree (15–17 y); 4 = one postgraduate degree (≥ 18 y); and 5 = multiple postgraduate degrees. Age data are mean \pm standard deviation (SD). Remaining data are mean \pm SEM. Active drug refers to the percentage of participants within each group who reported at the end of the experiment that they believed they had received an active drug. <http://dx.doi.org/10.6084/m9.figshare.3168682.v1>.

[#] Discrete data.

doi:10.1371/journal.pbio.1002575.t001

conducted a repeated-measures analysis of variance (RM-ANOVA) of the log(RTs) for correct responses on trials binned according to the five true conditional probabilities that existed in each of the TMs, grouped into High (0.85 and 0.70), Mid (0.25 and 0.20), and Low (0.05) transition probabilities, with drug as a between-subject factor. This revealed a significant decrease in log(RTs) with increasing transition probability (main effect of probability: $F_{1,27,150.10} = 28.32$, $p < 0.001$, effect size $\eta_p^2 = 0.19$), which was modulated by drug type (probability \times drug interaction: $F_{3,82,150.10} = 12.33$, $p < 0.001$, $\eta_p^2 = 0.24$).

Moreover, across the course of a contextual block (Fig 2B), participants became faster at responding to High and Mid probability stimuli and slower at responding to the Low probability stimuli (significant main effects of probability: $F_{1,28,151.41} = 27.80$, $p < 0.001$, $\eta_p^2 = 0.19$, and time: $F_{1,118} = 12.01$, $p = 0.001$, $\eta_p^2 = 0.09$; probability \times time interaction: $F_{2,236} = 6.55$, $p = 0.002$, $\eta_p^2 = 0.05$). The effect was modulated by drug type (probability \times time \times drug interaction: $F_{6,236} = 3.16$, $p = 0.005$, $\eta_p^2 = 0.07$) but again not systematically related to differences in Δ alertness or body weight between drug groups (all $p > 0.05$). Post hoc (Benjamini-Hochberg-corrected) pairwise comparisons indicated that the impact of drug was driven by the ACh- group, which showed significant log(RT) slowing compared to Placebo ($t_{57} = 2.98$, $p = 0.003$, effect size Cohen's $d = 0.79$). Together, these results indicate that participants learned about the true stimulus transition contingencies and that this learning was modulated by our pharmacological manipulations.

As is common in such behavioural tasks [80–83], participants showed evidence of post-error slowing on correct trials following those on which they made an error ($F_{1,118} = 10.92$, $p = 0.001$, $\eta_p^2 = 0.09$; Fig 2C). This effect was not modulated by drug group (trial-type \times drug interaction: $p = 0.933$) or by Δ alertness or body weight (both $p > 0.29$). Participants also demonstrated significant log(RT) slowing on correct post-infrequent trials (true transition

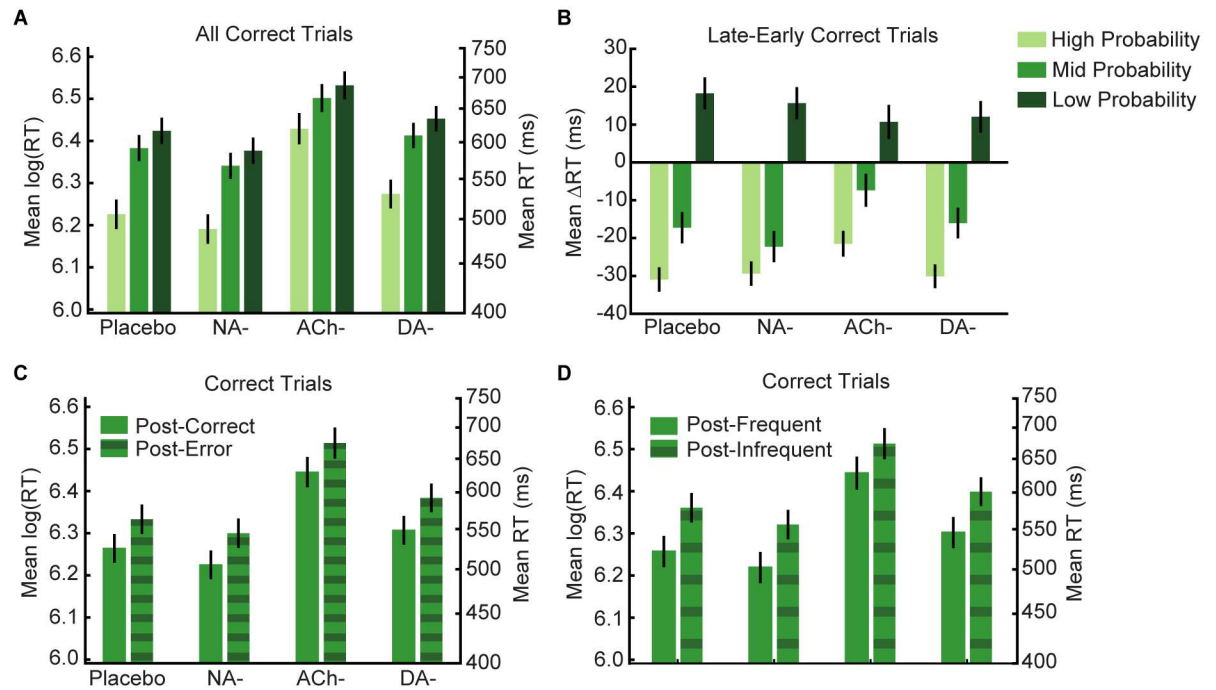


Fig 2. Model-agnostic changes in log(RT) indicate that participants learned to predict the stimulus transitions. (A) In all four groups, log(RT) increased as a stimulus' true transition probability decreased. (B) A median split on each 50-trial contextual block was used to compare mean log(RTs) on Early (1–25) and Late (25–50) trials at each probability level. Over the course of a context, participants became faster at responding to High and Mid probability stimuli and slower at responding to Low probability stimuli. (C) Across drug groups, participants showed evidence of post-error slowing on correct trials that followed an erroneous response compared to those following correct responses. (D) They also showed evidence of slowing on correct trials that followed an infrequent stimulus transition. Results are mean \pm SEM, corrected for the covariates Δ alertness and body weight. Results shown in A, B, and D were modulated by drug group. <http://dx.doi.org/10.6084/m9.figshare.3796353.v1>.

doi:10.1371/journal.pbio.1002575.g002

probability = 0.05) compared to all other correct trials ($F_{1,118} = 36.00, p < 0.001, \eta_p^2 = 0.23$; Fig 2D), which was modulated by drug group ($F_{3,118} = 4.58, p = 0.005, \eta_p^2 = 0.10$), but not by Δ alertness or body weight (both $p > 0.15$). This effect was driven by the ACh- group, with pairwise comparisons revealing significant slowing compared to Placebo ($t_{57} = 3.38, p = 0.001, d = 0.90$). Error rates significantly decreased with increasing transition probability (main effect of probability: $F_{1,57, 182.11} = 5.04, p = 0.013, \eta_p^2 = 0.04$). The effect was again modulated by drug type (probability x drug interaction: $F_{4,71, 182.11} = 4.72, p = 0.001, \eta_p^2 = 0.11$), but not by Δ alertness or body weight (both $p > 0.78$). There was no between-subject effect of drug group ($p = 0.768$). For additional model-agnostic analyses, please refer to S1 Text. Raw RT data can be found at <https://dx.doi.org/10.6084/m9.figshare.3793410.v2>.

Model-Based Analyses

The HGF model (Fig 3) allows us to map an individual's beliefs about stimulus transitions, transition contingencies, and volatility—and the respective irreducible, estimation, and volatility uncertainty about these beliefs—onto his/her observed RT responses. The HGF is hierarchical not only in that learning occurs simultaneously at multiple levels, but also in that belief updating at one level is constrained by beliefs at the level above. This provides a generic framework for implementing dynamic learning rates, which underlie learning in volatile environments [1,84]. We predicted that antagonising NA and ACh would impact on participants'

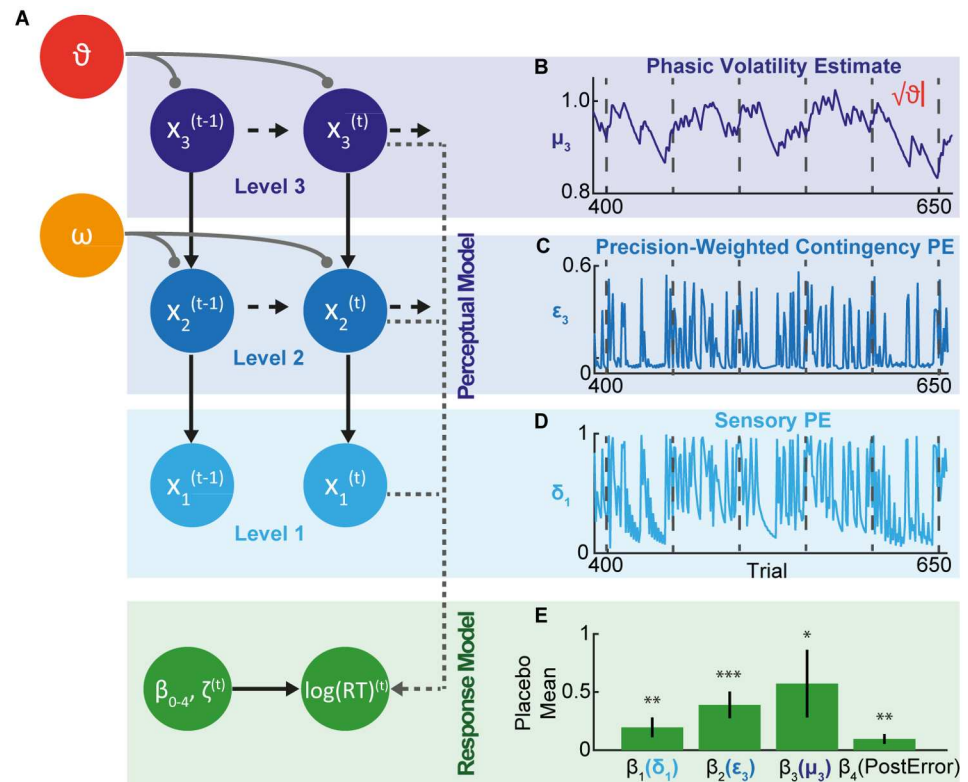


Fig 3. Hierarchical Gaussian Filter (HGF). (A) The perceptual model tracks an individual's learning of the task's structure across three levels. State x_1 represents trialwise stimulus transitions from one stimulus to the next, x_2 the probability of the transitions (i.e., transition contingencies), and x_3 the phasic volatility, where t is the current trial number. Participants hold and update beliefs about the true quantities at each level, with a mean μ and a variance σ . ϑ and ω are participant-specific parameters that couple the levels and determine the respective speed of belief updating about phasic volatility and transition contingencies. The response model describes the mapping from a participant's trialwise beliefs onto their observed $\log(RT)$ responses. (B) Example of the trialwise dynamics at level 3 from Placebo Participant 2. μ_3 reflects the participant's belief about the true phasic volatility (x_3). Vertical dashed lines indicate true context switches. μ_3 tends to increase following a context change and then decreases over the course of a context as the participant learns the new contextual rule and thus perceives the environment to be increasingly stable. $\sqrt{\vartheta}$ is a variance determining the step-size of μ_3 and therefore how quickly the participant updates their phasic volatility estimates. (C) As in B, but for precision-weighted contingency prediction error [PE] (ϵ_3) at level 2. This estimate results from weighting the contingency PE (δ_2) by a precision ratio that captures uncertainty about input from the level below relative to the level above. The higher the precision at level 2, the more meaningful a deviation from the predicted stimulus transition contingency. This in turn increases the impact on phasic volatility belief updating at level 3. For simplicity, we only depict the ϵ_3 trajectory for true transition changes. (D) As in B and C, but for sensory PE (δ_1) at level 1. This estimate arises from irreducible uncertainty about stimulus transitions. Trialwise values are equivalent to $1 - \hat{x}_1$, where \hat{x}_1 is equal to the participant's trialwise prediction of the occurring transition. Again, for simplicity we only show δ_1 values for true transitions here. (E) Mean β values for the Placebo group indicate that increases in sensory PE (β_1), precision-weighted contingency PE (β_2), and phasic volatility estimates (β_3) slowed participants' trialwise $\log(RT)$ s. There was also evidence of post-error slowing (β_4). Results are mean \pm SEM. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. <http://dx.doi.org/10.6084/m9.figshare.3796356.v1>.

doi:10.1371/journal.pbio.1002575.g003

computation of volatility uncertainty and estimation uncertainty respectively, while DA antagonism would impede motor adaptation to uncertain outcomes without perturbing the course of learning.

Perceptual model. The HGF's perceptual model tracks a participant's learning of the task's structure: the trialwise stimulus transitions at level 1, the probability of the transitions

(i.e., transition contingencies) at level 2, and the volatility of transition contingencies at level 3 (Fig 3A). Trialwise trajectories of each participant's perceptual estimates at each level evolve according to the predictions made and outcomes experienced by that individual (Fig 3B–3D). At levels 2 and 3, these estimates are modelled by Gaussian distributions with a mean (μ) and variance (σ), the latter reflecting the uncertainty of the estimate. Precision (π) of the estimate is equal to inverse variance ($1/\sigma$). Irreducible uncertainty at level 1 gives rise to sensory PE, δ_1 . Estimation uncertainty at level 2 gives rise to contingency PE, δ_2 . We can weight PEs according to their precision (inverse uncertainty). At level 1, this gives us precision-weighted sensory PE, ϵ_2 , and at level 2 precision-weighted contingency PE, ϵ_3 . Volatility uncertainty arises from phasic volatility beliefs, μ_3 , at level 3. See [Materials and Methods](#) and [S1 Text](#) for full details.

Importantly, the HGF does not assume fixed learning across the population but rather contains participant-specific parameters that couple the hierarchical levels and allow for individual expression of approximate Bayes-optimal learning. ϑ determines the speed of learning about volatility, i.e., the rate at which estimates of phasic volatility (μ_3) are updated. As such, ϑ encapsulates metavolatility, i.e., the rate at which volatility changes, with higher values implying a belief in a more unstable world and leading to a more variable learning rate (as expressed in phasic volatility belief updating). By contrast, ω is a constant component of the volatility and captures how rapidly individuals generally update their beliefs about transition contingencies at level 2. Changes in ω therefore lead to a tonic alteration of the learning rate. Comparing ϑ and ω estimates for each of the drug groups to the Placebo group allowed us to assess the effects of NA, ACh, and DA antagonism on perceptual belief updating.

We note that while model-agnostic analyses can provide an indication of learning and possible drug effects, the model we endorse here permits us to probe (our best guess at) subjective expectations about the transitions that are driven by data-limited observations. Moreover, it separates the set of relatively complex and interacting factors that influence RTs in a computationally limpid way and provides us with insight into the individual effects of our pharmacological manipulations.

Overall, the model tracked the true stimulus transitions well (Fig 4). We note that the model is uninformed about the true stimulus transition probabilities but rather bases its estimates on the observed stimulus transitions only. The punctate change points contained in the true generative process are detected implicitly by the HGF as an increase in learning rate (α_1), which reflects the influence of increased uncertainty and formally corresponds to a reduced contribution of belief precision (denominator in Eq 3) to the weighting of PE.

Moreover, when we categorise trials according to participants' trialwise estimates of transition contingencies, as provided by model parameter $\hat{\mu}_1$ (five bins: 0.8–1, 0.6–0.8, 0.4–0.6, 0.2–0.4, and 0–0.2), we observe the same decrease in log(RT) with increasing transition probability established in our model-agnostic results (cf. Fig 5 with Fig 2A; significant effect of $\hat{\mu}_1$: $F_{1,67,185.64} = 17.56, p < 0.001, \eta_p^2 = 0.14$). As in the model-agnostic results, this was modulated by drug group (significant $\hat{\mu}_1 \times$ drug interaction: $F_{5,02,185.64} = 9.51, p < 0.001, \eta_p^2 = 0.20$), but not by Δ alertness or body weight (both $p > 0.11$).

Response model. The response model describes the mapping from a participant's perceptual trialwise beliefs, as provided by the perceptual model, onto their observed log(RT) responses (Fig 3A). The response model we report in the following (Eq 1) describes log(RT) as a linear function of a constant component of log(RT), sensory PE (δ_1), precision-weighted contingency PE (ϵ_3), phasic volatility (μ_3), post-error slowing, and Gaussian noise (ζ). (See [Materials and Methods](#) and [S1 Text](#) for random effects Bayesian model comparison of

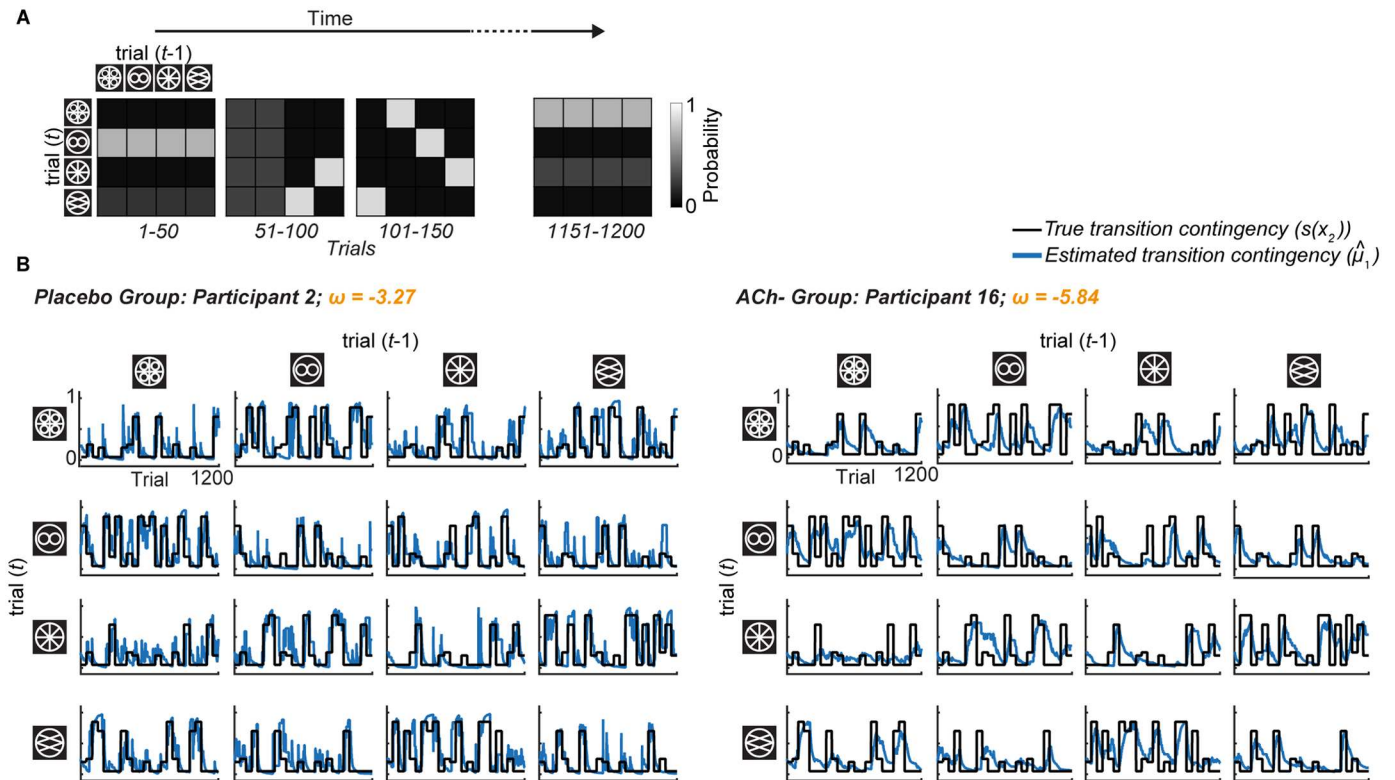


Fig 4. True and estimated stimulus transition contingencies for two example participants. (A) Transitions between pairs of stimuli, from trial $t-1$ to trial t , were defined by TMs. Every 50 trials, the TM switched to a different matrix. (B) Each panel corresponds to 1 of 16 possible transitions between stimuli across 1,200 trials. The black lines indicate the true transition contingencies. The blue lines reflect the participant's inferred estimates (i.e., the posterior expectation of these contingencies, $\hat{\mu}_i$), before seeing the stimulus outcome on each trial. The model tracked the true underlying contingencies and detected change points. In a representative participant from the Placebo group, the model tracked the true transition contingencies closely. An example participant from the ACh- group showed a greater discrepancy in the tracking of the true transition contingencies. This is reflected in the two participants' ω estimates: Placebo Participant 2 showed a higher transition contingency learning rate ($\omega = -3.27$) than ACh- Participant 16 ($\omega = -5.84$). <http://dx.doi.org/10.6084/m9.figshare.3796362.v2>.

doi:10.1371/journal.pbio.1002575.g004

alternative response models.)

$$\log(\text{RT})^{(t)} = \beta_0 + \beta_1(\delta_1^{(t)}) + \beta_2(\epsilon_3^{(t)}) + \beta_3(\mu_3^{(t)}) + \beta_4(\text{PostError}^{(t)}) + \zeta^{(t)} \quad (1)$$

All regression coefficients for the Placebo group were significantly greater than 0 (Fig 3E), meaning that sensory PE ($\beta_1(\delta_1)$: $t_{29} = 2.63, p = 0.010, d = 1.38$), precision-weighted contingency PE ($\beta_2(\epsilon_3)$: $t_{29} = 3.74, p < 0.001, d = 1.97$), and phasic volatility estimates ($\beta_3(\mu_3)$: $t_{29} = 2.05, p = 0.043, d = 1.08$) all had slowing influences on $\log(\text{RT})$ and that there was evidence of post-error slowing ($\beta_4(\text{PostError})$: $t_{29} = 3.09, p = 0.003, d = 1.62$). Each of the drug groups showed equivalent post-error slowing to the Placebo group (all $p > 0.48$; Fig 6H), mirroring our model-agnostic result. The lack of a difference in the noise parameter ζ between the Placebo group and any of the drug groups (all $p > 0.34$; Fig 6C) indicates that the model's ability to predict $\log(\text{RT})$ was unaltered under our drug manipulations. Comparing the regression coefficients for each drug group to the Placebo group allowed us to assess whether NA, ACh, and DA antagonism altered the mapping from perceptual beliefs onto motor responses.

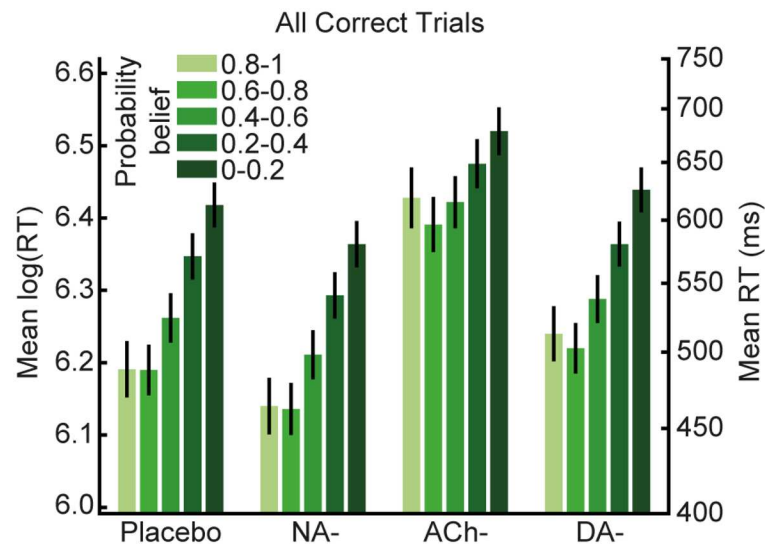


Fig 5. Model-based changes in log(RT) mirror our model-agnostic finding that participants learned to predict the stimulus transitions. In all four groups, faster responses were observed as participants' estimates of the true transition contingencies increased, mirroring our model-agnostic result (see Fig 2A). Results are mean \pm SEM, corrected for Δ alertness and body weight. <http://dx.doi.org/10.6084/m9.figshare.3796401.v2>.

doi:10.1371/journal.pbio.1002575.g005

The Influence of NA and ACh in Perceptual Uncertainty Computations

NA antagonism increased phasic volatility learning rate. The noradrenergic ($\alpha 1$ -receptor) antagonist prazosin increased the rate at which individuals updated their volatility estimates, as reflected by an increase in ϑ (linear model: $t_{59} = 2.11$, $p = 0.037$, effect size Cohen's $d = 0.56$; Fig 6A). A higher ϑ leads to greater fluctuations in participants' phasic volatility estimates, μ_3 , resulting in a more variable phasic learning rate. By contrast, there was no effect on ω ($p = 0.574$; Fig 6B), indicating that the tonic learning rate about the probabilistic contexts remained unchanged. We note that when we apply a Benjamini-Hochberg correction for three comparisons, the reported effect of NA- on ϑ does not survive. However, we do observe a significant increase in ϑ for the NA- group when we run permutation tests randomising over drug assignment (see S1 Text and S5 Table for details). Further, when we rerun the linear model analyses excluding body weight, which was equivalent across drug groups ($F_{3,120} = 2.24$, $p = 0.087$), as a covariate, the significant increase in ϑ for the NA- group survives the more stringent correction for three comparisons ($t_{60} = 2.32$, $p = 0.022$, $d = 0.60$), and all other reported significant effects are unchanged.

ACh antagonism slowed learning about stimulus transition contingencies. Muscarinic ACh receptor antagonism under biperiden had more widespread perceptual effects. While ϑ was again significantly increased compared to Placebo ($t_{57} = 2.86$, $p = 0.005$, $d = 0.79$), ω estimates in the ACh- group were significantly reduced ($t_{57} = -2.59$, $p = 0.011$, $d = -0.71$). The lower estimate of ω indicates that participants were slower to update their transition contingency estimates under biperiden and thus slower to adapt to the probabilistic contexts.

DA antagonism had no effect on learning about task structure. The D1/D2 DA receptor antagonist haloperidol did not influence the rate at which participants learned about the task's volatility or contextual transition contingencies compared to Placebo (ϑ and ω : both $p > 0.23$).

To summarise, both NA and ACh antagonism altered learning of uncertain events arising from unexpected contextual changes in the environment. Only ACh antagonism disrupted learning of transition contingencies within probabilistic contexts.

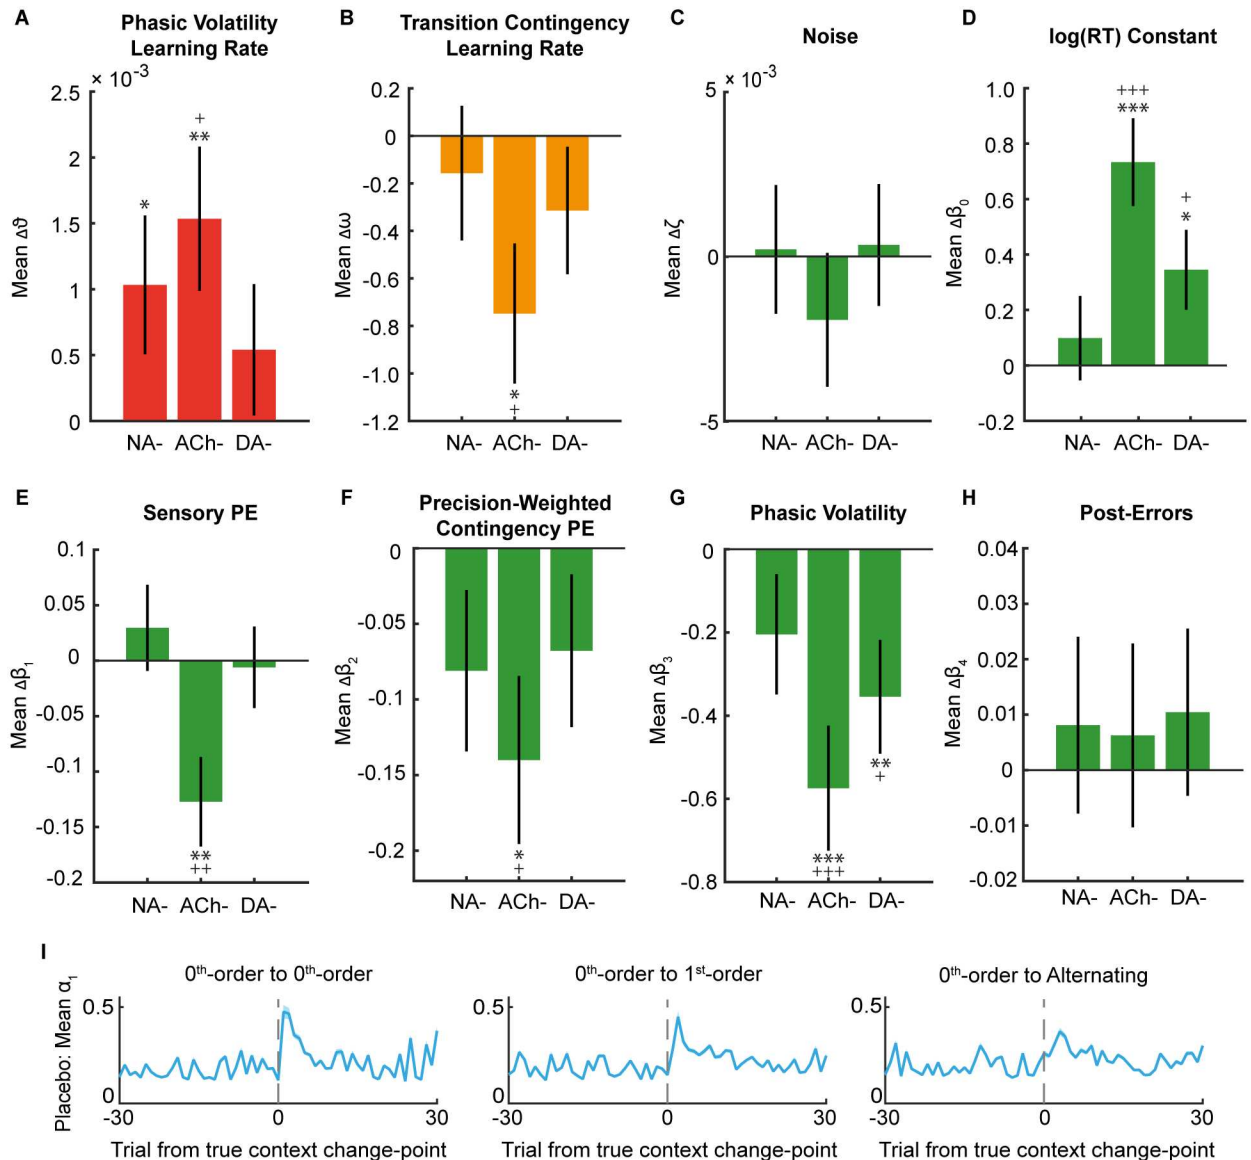


Fig 6. Model parameter results. (A–H) Compared to the Placebo group, NA and ACh antagonism modulated participants’ perceptual belief updating. NA- increased the rate at which participants updated their volatility estimates (increased ϑ). ACh- decreased the rate at which participants learned about stimulus contingencies (decreased ω) and increased participants’ ϑ estimates. DA- antagonism decreased the sensitivity of participants’ trialwise responses to their phasic volatility estimates (β_3). DA- and ACh- antagonism also caused some general response slowing (β_4). The three drug groups and the Placebo group showed equivalent post-error slowing (β_4) and Gaussian noise (ζ). Results are (mean Drug) – (mean Placebo), \pm the standard error of the difference (SED) between the means of the two samples and corrected for the covariates Δ alertness and body weight. For uncorrected comparisons: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. For corrections for three comparisons, ++s are used for Benjamini-Hochberg-corrected significance thresholds at + $p < 0.05$, ++ $p < 0.01$, +++ $p < 0.001$. See S2 Table for Placebo group means. (I) First-level learning rate (α_1) trajectories for truly occurring trials from the Placebo group. Increases in α_1 are observed following a true change in context. This α_1 increase is amplified for a more obvious switch from one easy-to-detect zeroth-order context to a different zeroth-order context. In contrast, a switch to an alternating context, which is trickier to detect, is accompanied by a modest, more gradual increase in α_1 . Results are mean \pm SEM. <http://dx.doi.org/10.6084/m9.figshare.3796407.v1>.

doi:10.1371/journal.pbio.1002575.g006

Neuromodulatory Effects on Response Modulation

NA antagonism had no influence on responses. The response model output revealed no significant effects of NA antagonism on participants' capacity to modulate their motor responses according to their perceptual estimates of uncertainty (Fig 6E–6G).

ACh antagonism reduced response sensitivity to perceptual beliefs. Compared to Placebo, ACh antagonism reduced the sensitivity of participants' motor responses to sensory PE (β_1 : $t_{57} = -3.25$, $p = 0.002$, $d = -0.90$), precision-weighted contingency PE (β_2 : $t_{57} = -2.56$, $p = 0.012$, $d = -0.71$), and phasic volatility estimates (β_3 : $t_{57} = -3.91$, $p < 0.001$, $d = -1.08$) (Fig 6E–6G).

DA antagonism reduced response sensitivity to phasic volatility. Compared to Placebo, DA antagonism led to a decrease in the influence of phasic volatility estimates on log(RT) (β_3 : $t_{60} = -2.65$, $p = 0.009$, $d = -0.67$; Fig 6G). This indicates that DA antagonism suppressed the sensitivity of motor responses to higher-level inference.

In addition to the effects reported above, the log(RT) constant output indicated that suppressing DA and ACh also led to some general log(RT) slowing (β_0 : $t_{60} = 2.48$, $p = 0.015$, $d = 0.62$; $t_{57} = 4.78$, $p < 0.001$, $d = 1.32$, respectively; Fig 6D). Subjective Δ alertness systematically modulated the effects observed on ϑ ($t_{118} = 2.62$, $p = 0.010$, $d = 0.02$), sensory PE (β_1 : $t_{118} = 2.52$, $p = 0.013$, $d = 0.02$) and precision-weighted contingency PE (β_2 : $t_{118} = -3.21$, $p = 0.002$, $d = -0.02$). None of the effects were systematically modulated by any between-group differences in body weight (all $p > 0.10$).

Control Analyses

In [S1 Text](#), we provide details of the physiological and subjective control measures. We also provide additional assessments of the HGF model and its ability to provide a good fit to the behavioural data by examining model parameter correlations, residuals and simulations ([S2–S6 Figs](#)). Finally, we report the results of additional permutation tests ([S5 Table](#)), which allow us to make distribution-free comparisons of the effects of our drug manipulations on the model parameters.

Discussion

By implementing a novel probabilistic serial RT task in conjunction with three pharmacological manipulations and placebo, we characterised the roles of three neuromodulatory systems during perceptual belief updating and response selection. Leveraging a hierarchical Bayesian learning model to decompose hierarchically related forms of uncertainty enabled us to pinpoint processes linked to NA, ACh, and DA. While manipulating NA and ACh modulated perceptual uncertainty computations, DA receptor antagonism reduced the sensitivity of the motor system to uncertainty estimates.

A key benefit of our pharmacological approach is that we were able to directly manipulate the function of three different neuromodulatory systems and compare the resulting psychopharmacological effects to a placebo condition. This is relevant given likely functional overlap between the different neuromodulatory systems, as observed here. Indeed, manipulation of a single neuromodulatory system, or use of a single drug, will be agnostic to such an overlap and may make any one effect appear more relevant and specific than it is. We are also able to extend interpretations of earlier neuroimaging studies [[15,16](#)], from which it is not possible to infer with certainty that activations in particular brain regions, with inhomogeneous cellular compositions, reflect the activity of specific neuromodulatory neurons.

Overlapping, but Dissociable, Noradrenergic and Cholinergic Influences on Perceptual Belief Updating

We found considerable overlap in the influence of NA and ACh antagonism on perceptual belief updating but also quantitative differences between drug conditions. While part-synergistic, part-antagonistic interactions between the two neuromodulators during uncertainty processing have been theorised previously [4], to our knowledge this is the first study to directly assess these putative computational roles, and to distinguish them from dopaminergic effects, under three pharmacological manipulations and within the same computational framework. We propose that ACh guides probabilistic learning within environmental contexts, while NA has a more circumscribed role in modulating the rate at which an agent learns about the volatility latent in the environment.

NA influences beliefs about unexpected environmental changes. Our results suggest that NA antagonism under prazosin altered the rate at which individuals updated their volatility beliefs, as indicated by an increase in the model parameter ϑ . We note, however, that this result should be interpreted conservatively since it did not survive correction for three comparisons. Nonetheless, a permutation test, without any distributional assumptions, mirrored the result of the uncorrected comparisons (see [S5 Table](#)). Moreover, the result did survive correction when body weight (which was equivalent across drug groups) was excluded as a covariate from the linear model analyses. An influence of NA- on ϑ fits with the theorised role for NA in computing uncertainty arising from changes in environmental context [4]. Numerous studies have offered evidence that the NA system is sensitive to highly unexpected events that arise from a hidden contextual change. Noradrenergic neurons in the rat and nonhuman primate LC are responsive to environmental novelty and unexpected changes in reward contingencies [48–50,85]. Additionally, changes in pupil diameter, attributed at least in part to noradrenergic LC activity [53,68,86], have been shown to correlate with unexpected outcomes [54,55].

More specifically, in the present study we observed faster volatility belief updating following NA antagonism. In our model, ϑ represents the volatility of the volatility, and thus, our results suggest that NA stabilises an agent's estimate of environmental volatility. This is compatible with the notion that the volatility estimate has a relatively low baseline level, to which it returns after being pushed away. In a volatile environment, this is not an adaptive feature. Rather, the volatility estimate should remain high to enable revision of one's beliefs. It is possible that NA prevents the volatility estimate from falling by reducing an agent's ϑ estimate.

The neurophysiological literature has distinguished two functional modes of LC noradrenergic release [44,87]. A phasic mode, characterised by a relatively low baseline firing rate and high phasic responsiveness to task relevant stimuli, has been linked to enhanced task engagement, and a tonic mode to increased distractibility, attention shifting, and exploratory behaviour [87–89] (but see Jepma et al. [90]). More recently, human LC BOLD activity was demonstrated to correlate with “unexpected uncertainty” induced by a switch in reward probabilities associated with familiar stimuli [16], although the negative sign of this correlation still seems to lack explanation. In both our task and that used by Payzan-LeNestour et al., contextual switches required participants to identify discrete changes in underlying transitions between familiar stimuli. To continue making accurate predictions in light of new transition probabilities, participants had to increase their attentional engagement to facilitate an augmented learning rate. It is likely that in both cases a phasic LC activity mode was recruited and that this would be recognised as a decrease in BOLD activity at a neural population level. Speculatively, it also suggests that our pharmacological NA manipulation may have enabled more phasic NA responsiveness to emerge under suppression of tonic NA firing. Future investigations of the impact of noradrenergic drugs on LC activity profiles are needed to validate this theory.

ACh balances the attribution of uncertainty within and between environmental contexts. Muscarinic ACh receptor antagonism by biperiden led to slower updating of beliefs about stimulus transition contingencies, and therefore slower adaptation to the probabilistic contexts, as reflected by a decrease in the model parameter ω . We argue that this slowed adaptation also had knock-on effects higher up in the inferential hierarchy. Namely, we propose participants attributed perceived violations of their expectations to gross contextual switches as opposed to chance fluctuations in stimulus outcomes, which would be expressed as an increase in ϑ . In light of previous work, which we discuss next, it seems reasonable to suggest that by setting the rate at which an agent learns probabilistic associations, ACh facilitates the appropriate attribution of violated expectations to chance fluctuations in an environment's statistical regularities or to gross switches in environmental context.

According to the structure of our Bayesian model, a reduction in ω maps onto a reduced precision weighting of perceptual belief updates at level 2; compare Eq 2 and Equation B in S1 Text. Our findings indicate that under biperiden less weight was given to sensory evidence, and updates of probability estimates became more reliant on current beliefs. This supports proposed roles for ACh in regulating the relative influences of stimulus-driven versus expectation-guided processing [52,91] and attentional deployment [37,40]. For instance, it has been shown that pharmacologically stimulating ACh augments bottom-up sensory signalling in human primary auditory cortex in response to auditory stimuli, possibly by enhancing the gain of superficial pyramidal cells, to bias inference towards sensory data [30].

In a recent study, Vossel et al. examined perceptual belief updating during a probabilistic attentional cueing paradigm. By applying a similar instantiation of the HGF to saccadic RTs, the authors demonstrated faster learning about contextual probabilities following administration of galantamine, an acetylcholinesterase inhibitor which increases the synaptic availability of ACh, as indicated by an increase in model parameter ω [17]. In the present study, we observe the opposite behavioural effect with the opposite pharmacological manipulation (ACh receptor antagonism), offering independent evidence that ACh signalling guides belief updating about probabilistic associations within environmental contexts.

Our results also indicate that ACh antagonism led individuals to update their volatility estimates more rapidly, reflected by an increase in the model parameter ϑ . This is consistent with the notion that ACh- participants' impaired ability to learn transition contingencies led them to infer that contexts changed at a faster rate. Notably, in their theoretical framework, Yu and Dayan predicted that ACh depletions should cause an agent to underestimate the amount of randomness in a given context. In turn, this causes chance events occurring within a context to seem more significant than they are, and thus, they are more likely to be incorrectly taken as indicative of a context change (see Fig 6D in Yu and Dayan 2005 [4]). Our experimental observations support this hypothesis and are compatible with data indicating that cholinergic antagonists increase distractibility [92], while agonists suppress it [93–95].

We note that, although the quantities used in this current work are not identical to those previously introduced by Yu and Dayan 2005 [4], the HGF does embody versions of the same forms of uncertainty. The highest level of uncertainty in Yu and Dayan's (YD's) framework was induced by abrupt, discrete changes in contingencies, which induced what YD call "unexpected uncertainty" (and ascribed to NA). By contrast, the highest level of uncertainty in the HGF is the overall instability of the world, i.e., the rate at which volatility changes. It is this that we found to be modulated by our NA antagonist. Conversely, YD's notion of "expected uncertainty" (ascribed to ACh) suggests that it arises from the known unreliability of predictive relationships within a familiar environmental context. Amongst other effects, the lower the expected uncertainty, the slower the learning—consistent with the effect of parameter ω in the HGF, which was found to decrease under cholinergic antagonism. Along with YD, we also

argue that this change in learning has further knock-on effects for what participants perceive to be a chance random event or a change of context (and hence unexpected uncertainty).

In sum, our findings offer empirical support for the theoretical proposal that ACh and NA interact to construct appropriate cortical representations of volatile contexts, which facilitates optimal inferences about the current environment [4]. By regulating high-level uncertainty representations, the two neuromodulators contribute to the updating of one's perceptual beliefs, both within and between environmental contexts, an idea that is broadly supported by recent neuroimaging [15,16] and pharmacological [17] evidence.

DA Sensitises Motor Responses to Environmental Volatility

As per its construction, our computational model allowed us not only to characterise perceptual belief updating under our three pharmacological manipulations but also to assess how each intervention influenced the deployment of motor responses in response to individual estimates of uncertainty. Pharmacologically manipulating DA and ACh altered the degree to which participants' perceptual beliefs modulated the preparation of their speeded responses to uncertain stimuli. In contrast, NA antagonism had no significant impact on the sensitivity of participants' motor responses to their current perceptual beliefs relative to placebo.

We had originally predicted that an individual's capacity to modulate response selection following a sensory PE would be dependent on DA. Indeed, it has previously been shown that pharmacological DA depletion impedes adaptive reactions to unexpected events occurring within predictable contexts [58]. However, in the present study, we found no evidence that DA receptor antagonism influenced participants' reactions to low-level sensory PE (δ_1). Rather, we observed that suppressing DA significantly reduced β_3 , which we interpret as a reduction in the sensitivity of participants' motor responses to their higher-level volatility estimates (μ_3).

It is important to note that some key differences distinguish our present experimental design from previous paradigms. In earlier work, participants were pre-trained to respond to stimuli presented within one predictable context, defined by one TM. Furthermore, switches from predictable to unpredictable contexts, consisting of random presentations of stimuli, were explicitly signalled [58]. Therefore, any probabilistic learning and higher-level perceptual uncertainty were removed. In this earlier setting, dopaminergic antagonism under haloperidol selectively impaired participants' reactions to unexpected events that elicited large sensory PEs.

In contrast, our present task created a more complex, and arguably more ecologically valid, scenario in which individuals had to infer the current context for themselves and adapt to any contextual changes. Here, uncertainty representations had to be acquired through direct sampling from a distribution of observations. To our knowledge, our study is the first attempt to interrogate the impact of DA, ACh, and NA on nonrewarded probabilistic learning within a single behavioural paradigm and a unified Bayesian framework. By estimating beliefs about various forms of uncertainty, we sought to identify neuromodulatory contributions specifically related to particular forms of uncertainty, as opposed to any confounding variables.

Our finding that haloperidol reduced the sensitivity of participants' responses to their phasic volatility estimates does sit well with an alternative line of work highlighting the importance of DA in cognitive switching [61,62]. For instance, Parkinson's disease patients with DA dysfunction have an impaired capacity to switch from naming digits to letters when both types of stimuli are presented simultaneously, even when the task shift is explicitly cued [60]. In summary, we propose that our DA manipulation suppressed response modulation by impeding cognitive switching following complex contextual rule changes.

Muscarinic ACh receptor antagonism under biperiden also led to decreased response modulation by parameters at all three hierarchical levels: sensory PE (δ_1), precision-weighted

contingency PE (ϵ_3), and phasic volatility estimates (μ_3) compared to Placebo. We propose that ACh receptor antagonism impeded participants' abilities to learn the statistical structure of the behavioural task, which in turn impaired their capacities to respond accordingly. Although both ACh and DA had effects on response modulation, in light of previous work, we suggest that DA's role is to modulate motor responses according to the widespread perceptual effects of ACh.

Limitations and Future Work

One of the main constraints of our study is that although prazosin, biperiden, and haloperidol are rather selective for NA, ACh, and DA receptors, respectively, there are complex interactions and dependencies between noradrenergic, cholinergic, and dopaminergic systems. Such interactions are a main reason why direct quantitative comparison between drug groups would not have provided direct comparisons between the action of different neuromodulators and therefore why our study was designed to detect changes relative to placebo instead. While our results highlight qualitative differences in how NA, ACh, and DA influence perceptual belief updating, future work will have to conduct direct quantitative comparisons of their roles.

Further, it is the receptors rather than the neuromodulators themselves that bring about psychophysiological effects, and there are dissociable roles of different receptor subtypes. For instance, the functions of nicotinic versus muscarinic cholinergic receptors in uncertainty signalling have yet to be directly compared. Distinctions have also been made between D1 and D2 dopaminergic receptor subtypes in regulating adaptive responses to unexpected stimuli [58]. Thus, future work could usefully be extended with genetic profiling and a range of selective agonists and antagonists for different receptor subtypes.

Finally, it is likely that all the neuromodulators operate over multiple timescales. For instance, separate, even competing, tonic and phasic effects have been a special target of investigation for NA. Teasing these timescales apart more fully is an ambition for the future, requiring a temporally richer design.

Nevertheless, our findings emphasise the necessity of studying the systems conjointly, as tasks associated with uncertainty will tend to involve them all.

Conclusion

In summary, our results offer novel and direct insight into the complex and intricate effects of NA, ACh, and DA during a probabilistic serial RT task. Employing a hierarchical Bayesian learning model that allowed us to assess various forms of uncertainty and PEs, we provide interventional evidence linking ACh and NA to uncertainty computations within and between behavioural contexts. In contrast, DA appears to be involved in sensitising motor responses to perceptual volatility estimates. While pharmacological manipulations do not selectively target particular neuromodulatory systems, our results offer a fresh perspective on the effects of noradrenergic, cholinergic, and dopaminergic neurotransmission on the computational mechanics of perceptual belief updating according to Bayesian principles. Future studies will verify the generality of our observed effects to different behavioural paradigms with and without learning, reward, prediction, and action. By characterising uncertainty computations and response modulation, our methodology can also be used to offer fresh insight into the numerous neurological and psychiatric disorders in which there is dysregulation of processes dependent on NA, ACh, and DA.

Materials and Methods

Participants

A total of 128 healthy participants (56 male, aged 18–38 y, 119 right-handed) with normal or corrected-to-normal vision undertook this study. Participants gave written informed consent in accordance with the Declaration of Helsinki. The experimental protocol was approved by the UCL Research Ethics Committee (Project ID: 3491/001). The following exclusion criteria applied: history of neurological or psychiatric disease, intake of medication (other than contraceptives), smoking, recreational drug use, and current participation in other pharmacological studies. Following a screening interview to rule out intolerances or contraindications, the study clinician (DR) assigned participants pseudorandomly (i.e., ensuring a balanced distribution of gender, age, and body weight) to receive a NA, ACh, or DA antagonist or a placebo. All other authors, including the experimenter (LM), were blind to the drug conditions.

General Procedure

We employed a double-blind, between-subject design. Each participant attended one experimental session during which they received a single oral dose of one of the following: 1 mg prazosin (α 1-receptor antagonist; NA- group), 6 mg biperiden (M1-receptor antagonist; ACh-group), 2.5 mg haloperidol (D1/D2-receptor antagonist; DA- group), or a placebo. We selected doses that were in line with previous studies showing clear behavioural and neurophysiological effects [58,96–98]. On arrival, participants completed computerised versions of the Digit Span test, Barratt Impulsiveness Scale (BIS-11) [99], Doman-Specific Risk-Taking (DOSPERT) Scale [100], and Cognitive Failures Questionnaire (CFQ) [101]. Participants also self-reported their baseline mood (alertness, calmness, and contentedness) with visual analogue scales (VASs) [102], and we measured their baseline heart rate (HR) and blood pressure (BP). To assess any subjective and/or physiological drug effects, the VAS, HR, and BP measurements were repeated before participants started the RT task and again once they completed it.

Two different drug administration times were used to match peak plasma concentration across drugs, based on previous pharmacokinetic data. To ensure that participants undertook the RT task when the drug was at its most active, haloperidol was administered 2 h in advance, while prazosin and biperiden were administered 1.5 h before the main experimental session [96–98]. A random 50% of participants from the Placebo group were administered a placebo tablet at the first time point, and the other 50% at the second time point. The study clinician administered the drug or placebo while the experimenter was away from the testing room. Participants were asked not to eat for at least 1 h before the first drug administration time.

Probabilistic Serial RT Task

Participants sat facing a computer screen positioned approximately 60 cm away. They were instructed to rest their left and right index and middle fingers on the four buttons of a custom-made button box placed in front of them and to maintain this position throughout the task. On each trial, participants were required to respond to the presentation of one of four visual stimuli by pressing an appropriate button as quickly as possible. Each stimulus was associated with one particular button. The stimulus-response mappings remained consistent within an experimental session but were counterbalanced across participants. Each participant acquired the stimulus-response mappings for their session during a training block in which they received visual error feedback after each trial. The training session comprised at least 100 trials and did not finish until the participant had reached a minimum performance criterion of 85% accuracy on the last 20 trials. Participants were then given 40 practice trials, in which the

stimuli were presented in a random order and without error feedback, to familiarise them with the timings of the main experiment. An additional refresher block, consisting of at least 26 trials with error feedback, was completed immediately before the main experiment. Again, participants had to achieve 85% accuracy in the last 20 trials to proceed. On average, participants reached this criterion in 28.1 ± 1.1 trials, indicating adequate learning and retention of the mappings. There was no difference in the number of refresher trials required between groups ($F_{3,120} = 1.17, p = 0.324$).

Each participant performed 1,200 trials of the probabilistic RT task. Fig 1A shows an example trial sequence. Anticipatory responses (< 80 ms RT) were recorded as incorrect. At any given time, the trial sequence was generated by one of eight TMs, which changed every 50 trials. In each case, there were 16 combinations that determined the probabilistic relationship between the stimuli presented on the current trial, t , and the previous trial, $t-1$. Three types of TM were utilised: first-order sequences, alternating sequences, and zeroth-order sequences (see Fig 1B and S1 Fig for further details). The order of TMs was pseudorandom, with no consecutive repeats. Importantly, the overall probability of each stimulus was equal across the 1,200 trials.

The pseudorandom order of TMs was used to generate one stimulus sequence that was used for all participants to ensure comparable learning processes and model parameter estimates. Rest periods occurred every 185 trials, orthogonal to TM switches. The importance of fast responses was stressed. Participants were told that by paying attention to any patterns in the order in which stimuli were presented and to any switches in these patterns, it might be possible to respond faster. No further information about the nature of the experiment was provided.

Combining our behavioural paradigm with three pharmacological manipulations allowed us to assess directly any separable roles for NA, ACh, and DA in belief updating under irreducible uncertainty, estimation uncertainty and volatility uncertainty and in sensitising the motor system to participants' individual perceptual beliefs. At the end of the experimental session, participants were debriefed, indicated whether they thought they had taken an active drug or placebo, and reported the quality and quantity of their sleep on the previous night [103].

Model-Agnostic Analyses

Trialwise RT was calculated as the time between stimulus onset and the subsequent button press. The RT data were log transformed [58]. To assess the interaction between stimulus transition probability and drug, trials were binned according to three probability levels corresponding to the presented stimuli's true conditional probabilities as existed in the TMs (High: 0.85 and 0.70; Mid: 0.25 and 0.20; Low: 0.05) [58,59]. Mean log(RTs) for each participant's correct responses were compared across these three probability levels and between drug groups. To obtain a model-agnostic indication of learning across the course of the probabilistic contexts, we performed a median split on each 50-trial contextual block and compared correct mean log(RTs) on Early (1–25) and Late (26–50) trials at each probability level and between drug groups. In reporting statistical differences, a significance threshold of $\alpha = 0.05$ was used. Where assumptions of sphericity were violated (Mauchly's test $p < 0.05$), the Greenhouse-Geisser correction was applied. For comparisons across the four groups, partial eta-squared (η_p^2) is reported as the effect size. For planned pairwise comparisons between each drug group and Placebo, we applied a Benjamini-Hochberg correction for three comparisons to the significance threshold. For pairwise comparisons, Cohen's d is reported as the effect size.

Model-Based Analyses

Model-based analyses permit quantification of the inferences participants made during the task and decomposition of the impact of different forms of uncertainty. Hierarchical Bayesian

models have proven powerful for explaining the adaptation of behaviour to probabilistic contexts in volatile environments [1,15,17,70,72–74,84,104,105]. Here we developed a novel instantiation of the HGF model [9] with a focus on TMs and two components: a three-level perceptual model and a response model (Fig 3A). We implemented the model using the HGF Toolbox (<http://www.translationalneuromodeling.org/tapas/>). Matlab (Mathworks, United States) code can be downloaded directly from <http://www.tnu.ethz.ch/de/software/tapas/download.html>. Note that, for the present study, we used the `tapas_logrt_linear_whatworld` family of scripts. The HGF is unconstrained by normative optimality assumptions of “ideal Bayesian” models that enshrine the true characterisation of the task (or alternatively Bayesian hierarchical uncertainty about this true characterisation) [76,106,107]; instead, it allows for interindividual variability in learning by fitting parameters that determine a participant-specific approximation to Bayes optimality.

Perceptual model. The perceptual model is a variant of the HGF as introduced in Mathys et al. [9,10]. It tracks a participant’s learning of the task’s structure in a similar way to previous studies using the HGF [15,17,70–74]. The HGF is a generative model of observed data that provides a mapping from hidden states of the world (x) to sensory inputs (u). Unlike previous applications of the HGF, in this application the data are observed transitions between stimuli that arise from a sequence of environmental states (x_1), where we use bold font to indicate a matrix. In our experiment, the jk^{th} element of x_1 is the transition from stimulus k to stimulus j , the probability of which participants must learn to perform the task well. There are 16 possible transitions induced by the trialwise presentation of one of four visual stimuli, meaning that x_1 is a four-by-four matrix. On each trial, an agent observes a sample from one column of the TM. Therefore, the current transition in the corresponding column of x_1 is 1, with all other elements in that column equal to 0.

The HGF comprises two parts: a generative (or forward) model that describes how data are generated, including the effect of such facets as the changing contexts, and a recognition model that performs (approximate) statistical inference on the observations of the actual data in order to determine the distribution over the values of random variables in the generative model appropriate to those particular observations. In contradistinction to models that assume that participants fashion the generative process to the task at hand, the HGF offers an inclusive scheme for explaining behaviour that generalises to a multitude of situations requiring hierarchical inference about the state of the world.

In the present application of the HGF, the generative model has two further levels above x_1 . Level 2 is a four-by-four matrix x_2 of real numbers governing the transition contingencies. These undergo random walks with increments that are independent of each other. At level 3, x_3 sets the variance of those random walks and therefore the rate of change (or volatility) of the elements of x_2 . Since we assume that all the elements experience the same volatility (see Mathys et al. 2011; 2014), x_3 is a scalar. Collectively, x_2 and x_3 capture stimulus transitions and their changes over time (albeit represented heuristically as a continuous random walk in logit space with a bijective mapping to the probability of specific discrete changes). More specifically, a sample of x_1 is generated by applying a logistic sigmoid transformation to the column of x_2 associated with the stimulus that was previously shown, to generate a probability distribution over the four possible next stimuli (see S1 Text for details). A sample is then drawn from that distribution.

The recognition part of the HGF takes observations of x_1 and infers approximate posterior distributions over the values of x_2 and x_3 . This amounts to a variant of predictive coding in which beliefs are dynamically updated across the levels via PEs that are weighted by their salience, or expected precision (equivalent to inverse variance, or uncertainty). Estimates of stimulus transition probabilities, i.e., the posterior distribution over x_2 , correspond to the posterior distribution over x_2 and are updated by PEs about stimulus occurrences. Estimates of

environmental volatility, i.e., the posterior distribution over x_3 , are updated in proportion to PEs about the transition contingencies. Thus, the effective learning rate is influenced by uncertainty about current beliefs and environmental stability.

More precisely, trial by trial, participants update their beliefs about the true quantities at each level, which at levels 2 and 3 are modelled by Gaussian distributions with a mean (μ) and variance (σ), the latter reflecting the uncertainty of the estimate. Precision ($\hat{\pi}$) of the prediction is equal to the inverse variance ($1/\hat{\sigma}$), where the hat denotes the participant's predicted estimate before seeing the stimulus outcome on each trial. At level 1, when the elements of \mathbf{x}_2 are each transformed by the logistic sigmoid to produce probabilities \hat{x}_1 there is irreducible uncertainty (the participant's estimate of which is written as $\hat{\mu}_1$), which gets its name since it is undiminished by learning [11]. Irreducible uncertainty arises from any probabilistic relationship and is closely related to entropy, with an inverted-U relationship to probability that peaks at $p = 0.5$. The quantity gives rise to sensory PE (δ_1) following the presentation of an unexpected stimulus that would require a participant to respond against their expectation.

Note that we do not enforce columnwise normalisation of \hat{x}_1 (i.e., the columns do not necessarily add up to one, as they would have to in order to represent a proper probability distribution over mutually exclusive events). We argue that ensuring that the probabilities sum to one would require a sort of certainty about the stimuli that participants do not necessarily have when performing the task; for instance, it would require precise a priori knowledge that each and every trial will present exactly one of four stimuli and that there is no possibility of novel stimuli occurring during the experiment. In practice, the statistics governing the sensory events that occurred in our task ensure that column sums of participants' $\hat{\mu}_1$ estimates never stray far from unity.

At level 2, $\hat{\sigma}_2$, which is informational in origin, represents estimation uncertainty about the true probabilistic relationships governing stimulus transitions, giving rise to a more abstract contingency PE (δ_2). At level 3, volatility uncertainty arises from the environment's volatility, i.e., how quickly the transition contingencies are changing.

Generally, at any level i of the hierarchy, the update of the belief on trial t (i.e., posterior mean $\mu_i^{(t)}$ of the state x_i) is proportional to the precision-weighted PE, $\epsilon_i^{(t)}$. This weighted PE is the product of the upward-propagating PE, $\delta_{i-1}^{(t)}$, and a precision ratio, $\psi_i^{(t)}$ (see S1 Text for details), capturing the uncertainty about input from the level below relative to the uncertainty about the state of the level being updated [15]. A general and didactically useful form of this precision-weighted PE (with subtle differences below level 3; see Mathys et al. 2014 [10]) is:

$$\Delta\mu_i^{(t)} \propto \epsilon_i^{(t)} = \psi_i^{(t)} \delta_{i-1}^{(t)}, \tag{2}$$

$$\psi_i^{(t)} = \frac{\hat{\pi}_{i-1}^{(t)}}{\pi_i^{(t)}}. \tag{3}$$

Thus, precision-weighted sensory PE (ϵ_2) is weighted by uncertainty at levels 1 and 2, and precision-weighted contingency PE (ϵ_3) by uncertainty at levels 2 and 3. The punctate change points contained in the true generative process are detected implicitly by the HGF via spikes in the precision weights. At levels 2 and 3, $\alpha_i^{(t)}$ is proportional to the precision ratio, $\psi_i^{(t)}$, defined in Eq 3. At level 1, the learning rate α_1 is simply defined as the update divided by the PE:

$$\alpha_1^{(t)} \propto \frac{\mu_2^{(t)} - \hat{\mu}_1^{(t)}}{\delta_1^{(t)}}. \tag{4}$$

Two participant-specific perceptual parameters, ϑ and ω , couple the hierarchical levels and allow for individual expression of approximate Bayes-optimal learning. ϑ determines the speed

of learning about the volatility of the environment, i.e., the rate at which estimates of trialwise phasic volatility (μ_3) are updated. Higher ϑ values imply a belief in a more unstable world and lead to a more variable learning rate. By contrast, ω is a constant component of the volatility and captures a tonic learning rate about stimulus transition contingencies (see Fig 3A and S1 Text for details).

Response model. The response model describes the mapping from a participant's trialwise beliefs, as provided by the perceptual model, onto his/her observed responses, $\log(\text{RTs})$. We reasoned that there are several variables that could influence trialwise $\log(\text{RT})$. Therefore, we constructed and compared three response models using random effects Bayesian model comparison [108,109] and associated techniques for assessing differences in model frequencies across groups as implemented in the VBA toolbox [110]. Each model proposed that $\log(\text{RT})$ on any given trial is a linear function of a constant component of $\log(\text{RT})$ and several other factors. Since there is evidence, both from earlier work [80–83] and the present study, that participants' RTs increase on a trial following an incorrect response, we included post-error slowing in each response model. While the perceptual model assumed that participants updated their beliefs according to the stimulus presented on each trial, the response model incorporated correct trials only.

The extra factors in the different models came from quantities at each level of the HGF that might influence $\log(\text{RT})$. The first response model contained the following parameters: δ_1 (sensory PE), because of evidence that DA sensitises motor responses to low-level PE [58,59]; ϵ_3 (precision-weighted contingency PE), which has been shown to correlate with activity in the cholinergic basal forebrain [15]; and μ_3 (estimated phasic volatility), which is relevant to cognitive switching tasks for which there is a proposed role for DA. For each parameter, the quantity relates to the stimulus transition that actually occurred on each trial. ζ is Gaussian noise.

Response Model 1:

$$\log(\text{RT})^{(t)} = \beta_0 + \beta_1(\delta_1^{(t)}) + \beta_2(\epsilon_3^{(t)}) + \beta_3(\mu_3^{(t)}) + \beta_4(\text{PostError}^{(t)}) + \zeta^{(t)} \quad (5)$$

Alternative research has indicated that activity in the dopaminergic midbrain correlates with the precision-weighted form of sensory PE, ϵ_2 [15]. To disambiguate whether motor responses are modulated according to raw sensory PE or the confidence one has in their sensory predictions, our second model contained ϵ_2 instead of δ_1 .

Response Model 2:

$$\log(\text{RT})^{(t)} = \beta_0 + \beta_1(\epsilon_2^{(t)}) + \beta_2(\epsilon_3^{(t)}) + \beta_3(\mu_3^{(t)}) + \beta_4(\text{PostError}^{(t)}) + \zeta^{(t)} \quad (6)$$

Since δ_1 and ϵ_2 are highly correlated, we also constructed a third response model containing both parameters so as to ascertain whether one had a higher degree of explanatory power in terms of determining $\log(\text{RT})$.

Response Model 3:

$$\log(\text{RT})^{(t)} = \beta_0 + \beta_1(\delta_1^{(t)}) + \beta_2(\epsilon_2^{(t)}) + \beta_3(\epsilon_3^{(t)}) + \beta_4(\mu_3^{(t)}) + \beta_5(\text{PostError}^{(t)}) + \zeta^{(t)} \quad (7)$$

Random effects Bayesian model comparison established that Response Model 1 (containing parameters δ_1 , ϵ_3 , and μ_3) was superior in all four pharmacological groups by a considerable margin: for the Placebo, NA-, ACh-, and DA- groups respectively, the posterior model probabilities were 0.911, 0.828, 0.636, and 0.829; protected exceedance probabilities (i.e., the probability that Response Model 1 is more likely than any other model in the comparison set) were 1.000, 1.000, 0.964, and 1.000 (Fig 7). Moreover, no significant difference in model frequencies between the Placebo group and any of the drug groups was identified (NA- versus Placebo: $p = 0.958$, ACh- versus Placebo: $p = 0.560$, DA- versus Placebo: $p = 0.955$). To further verify

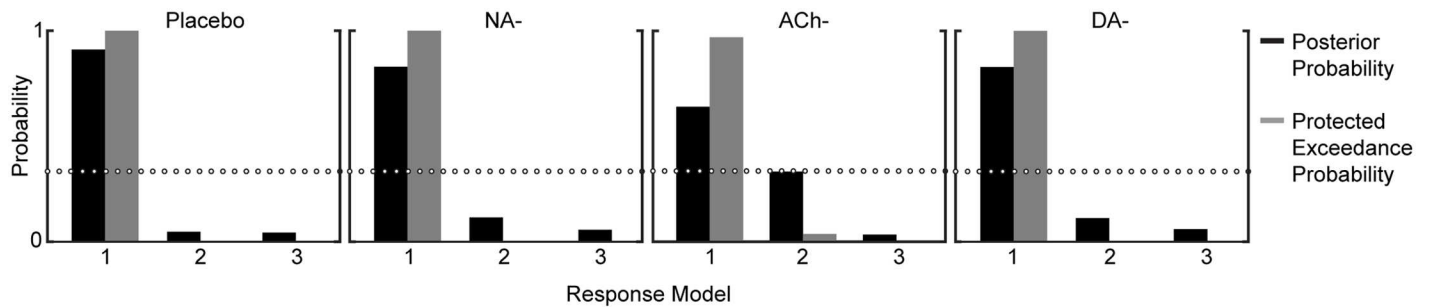


Fig 7. Random effects Bayesian model comparison confirmed that Response Model 1 was superior in all four groups. Posterior probabilities quantify the likelihood of each model given the data. Protected exceedance probabilities quantify how likely it is that any given model is more frequent than all other models in the comparison set while also protecting against the possibility that the observed variability in (log-) model evidences could be due to chance. The dotted line indicates the threshold for chance-level posterior probabilities (0.33). <http://dx.doi.org/10.6084/m9.figshare.3796413.v1>.

doi:10.1371/journal.pbio.1002575.g007

that Response Model 1 offered the best means by which to explain trialwise log(RT), we also compared a more exhaustive set of models containing different combinations of parameters from the HGF (see S3 and S4 Tables for details). Again, Response Model 1 was found to be superior and was therefore used for all subsequent analyses.

Model fitting. For each participant, individual maximum a posteriori estimates for perceptual and response model parameters were jointly obtained using the Broyden-Fletcher-Goldfarb-Shanno algorithm as implemented in the HGF Toolbox. Where priors were required, we defined them by inverting the perceptual model in isolation, given the known stimulus sequence (using the function “`tapas_bayes_optimal_whatworld_config`”), under suitably uninformative priors. We then used the resulting posterior estimates to define our priors for the subsequent inversion of the full model given the behavioural data (see S1 Table). In other words, the prior means in our empirical data analysis corresponded to those parameter values for which our stimulus sequence would generate minimal surprise (in an observer with the aforementioned uninformative priors).

Our key question pertained to how different neuromodulators influence learning and response modulation. Therefore, to make planned comparisons between the three active drug groups (NA-, ACh-, and DA-) and the Placebo group, we fit linear models separately for each participant-specific parameter (ϑ , ω , and each β , respectively), with additional covariates (body weight and Δ alertness). For pairwise comparisons, Cohen’s d was used as a measure of effect size. In Fig 6, we present the results of the uncorrected, planned pairwise comparisons taken from each linear model and those that survive a Benjamini-Hochberg correction for three pairwise comparisons, i.e., between each drug group and Placebo.

Control Analyses

We validated our model parameter results using permutation tests. Moreover, to further illustrate the implications of different model parameters and to demonstrate that the HGF provided a good fit to the behavioural data, we assessed the correlations between model parameters, analysed residuals, and used the HGF to generate simulated log(RT) data. Please see S1 Text for details.

Supporting Information

S1 Fig. Probabilistic structure of the experimental task. At any given time, the trial sequence was generated by one of eight TMs, which changed every 50 trials. In each case, there were 16 combinations that determined the probabilistic relationship between the stimuli presented

on the current trial, t , and the previous trial, $t-1$. TMs 1 and 2 generated first-order stimulus sequences in which there was a high probability of the sequences 1-2-3-4 and 4-3-2-1 occurring, respectively. TMs 3 and 4 resulted in a high probability of alternating between two stimuli. TMs 5–8 were zeroth-order and led to one stimulus occurring with a high probability, one stimulus with a mid-range probability, and two stimuli with a low probability. Over the course of the experiment, each of the TMs occurred three times in a pseudorandom order, with no consecutive repeats. The overall probability of each stimulus was equal across 1,200 trials. (TIF)

S2 Fig. Model parameter correlations. Unless stated otherwise, Bayesian parameter averages (BPAs) and posterior means for the model parameters were only moderately correlated across groups (all absolute $r < 0.660$ for BPAs and $r < 0.716$ for posterior means). Higher correlations existed between BPAs for ω (transition contingency learning rate) and μ_{3_0} (the initial phasic volatility estimate) ($r = -0.948, -0.764, -0.771$, and -0.983 for Placebo, NA-, ACh-, and DA-, respectively). Higher correlations also occurred between the BPAs ($r = -0.877, -0.736, -0.630$, and -0.880) and mean posteriors ($r = -0.938, -0.947, -0.873$, and -0.893) for β_0 (log(RT) constant) and $\beta_3(\mu_3)$ (the sensitivity of log(RTs) to phasic volatility estimates). σ_{3_0} is the initial value of σ_3 (the uncertainty about the phasic volatility estimate). <http://dx.doi.org/10.6084/m9.figshare.3796434.v1>. (TIF)

S3 Fig. Residuals between observed log(RTs) and log(RTs) predicted by the HGF. The distribution of residuals suggests that, across drug groups, the model captured any patterns in the data well. Data are mean \pm SEM. <http://dx.doi.org/10.6084/m9.figshare.3796443.v1>. (TIF)

S4 Fig. Autocorrelation between residuals across trials indicate that the model did not systematically under- or overestimate log(RTs) at true change points. Data are mean \pm SEM. <http://dx.doi.org/10.6084/m9.figshare.3796446.v1>. (TIF)

S5 Fig. Mean log(RTs) for empirical and simulated data. Empirical data (filled bars) indicated that log(RT) increased as a stimulus' true transition probability decreased. Simulated data (unfilled bars) generated using the posteriors for each participant in the Placebo group as model parameters faithfully reflect the empirical Placebo data. By shifting the parameters significantly altered by our different drug manipulations by the difference between the Placebo group mean for those parameters and the relevant drug-group mean, we simulated log(RT) data comparable to the empirical data observed in each drug-group. Note that there are no post-error slowing effects in the simulated data. Data are mean \pm SEM. <http://dx.doi.org/10.6084/m9.figshare.3796449.v2> (TIF)

S6 Fig. Responses to high probability events become faster with learning but are slower after change points. (A) Mean baseline-corrected RTs for the first 25 (High-probability) trials in each context, where the baseline is the mean RT of the last three High-probability trials in the previous context. RTs increase following a true change point but fall as participants learn the new contextual rule. (B) As in A, but for simulated RTs. The model neatly captures the increase in RTs following true change points, the reduction in RT that occurs with learning across the course of the new context, and the suppressed effect of both in the ACh- group. <http://dx.doi.org/10.6084/m9.figshare.3796452.v2>. (TIF)

S1 Table. A summary of HGF parameters and priors. All priors are specified in the space in which they are estimated. For an account of how this relates to the native space of that parameter, the reader is referred to the original description of the model [9].

(DOCX)

S2 Table. Placebo group averages for the perceptual and response model parameters. All results are mean \pm SEM, corrected for the covariates Δ alertness and body weight. β_0 reflects a constant component of $\log(\text{RT})$. β_{1-4} reflect the influence of sensory PE (δ_1), precision-weighted contingency PE (ϵ_3), phasic volatility estimates (μ_3), and post-error trials on $\log(\text{RT})$. All β values were significantly greater than zero (all $p < 0.05$), indicating that these parameters slowed $\log(\text{RT})$. <http://dx.doi.org/10.6084/m9.figshare.3796407.v1>.

(DOCX)

S3 Table. Results of familywise Bayesian model comparison. To further verify that Response Model 1 offered the best means by which to explain trialwise $\log(\text{RT})$, we compared a more exhaustive set of linear response models containing different combinations of parameters from the HGF on the Placebo group. We first ran a familywise model comparison on models containing every combination of the parameters δ_1 , ϵ_2 , ϵ_3 , and μ_3 (Family 1) versus models containing every combination of $\hat{\sigma}_1$, $\hat{\sigma}_2$, and $\hat{\sigma}_3$ (Family 2). Note that the quantities corresponded to the transition that actually occurred on each trial. All models included post-error slowing. Family 1 was found to be superior (posterior probability: 0.700; exceedance probability: 0.999).

(DOCX)

S4 Table. Results of exhaustive random effects Bayesian model comparison for models containing each combination of parameters δ_1 , ϵ_2 , ϵ_3 , and μ_3 . We ran random effects Bayesian model comparison on all the models in Family 1. Response Model 1 was found to be superior (posterior probability: 0.270; protected exceedance probability: 0.844).

(DOCX)

S5 Table. Results of the permutation tests, randomising over drug assignment, for the HGF model parameters. Aside from the effect of ACh- on β_2 , all significant effects observed in the uncorrected multiple comparisons (cf. Fig 6) are mirrored in the results of the permutation tests. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; ns = nonsignificant.

(DOCX)

S1 Text. Supporting information on physiological and subjective control measures, additional model-agnostic analyses, the HGF model, permutation tests, model parameter correlations, residuals, and simulations.

(DOCX)

Author Contributions

Conceptualization: LM SB.

Formal analysis: LM CM AOdB KES PD SB.

Funding acquisition: LM SB.

Investigation: LM DR.

Methodology: LM CM DR SB.

Project administration: LM.

Software: LM CM.

Supervision: SB.

Visualization: LM CM SB.

Writing – original draft: LM CM AOdB PD KES SB.

Writing – review & editing: LM CM DR AOdB PD KES SB.

References

1. Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS. Learning the value of information in an uncertain world. *Nat Neurosci*. 2007; 10: 1214–1221. doi: [10.1038/nn1954](https://doi.org/10.1038/nn1954) PMID: [17676057](https://pubmed.ncbi.nlm.nih.gov/17676057/)
2. Conant RC, Ashby WR. Every good regulator of a system must be a model of that system. *Int J Syst Sci*. 1970; 1: 89–97. doi: [10.1080/00207727008920220](https://doi.org/10.1080/00207727008920220)
3. Körding KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature*. 2004; 427: 244–247. doi: [10.1038/nature02169](https://doi.org/10.1038/nature02169) PMID: [14724638](https://pubmed.ncbi.nlm.nih.gov/14724638/)
4. Yu AJ, Dayan P. Uncertainty, Neuromodulation, and Attention. *Neuron*. 2005; 46: 681–692. doi: [10.1016/j.neuron.2005.04.026](https://doi.org/10.1016/j.neuron.2005.04.026) PMID: [15944135](https://pubmed.ncbi.nlm.nih.gov/15944135/)
5. Bestmann S, Harrison LM, Blankenburg F, Mars RB, Haggard P, Friston KJ, et al. Influence of Uncertainty and Surprise on Human Corticospinal Excitability during Preparation for Action. *Curr Biol*. 2008; 18: 775–780. doi: [10.1016/j.cub.2008.04.051](https://doi.org/10.1016/j.cub.2008.04.051) PMID: [18485711](https://pubmed.ncbi.nlm.nih.gov/18485711/)
6. Averbach BB, Latham PE, Pouget A. Neural correlations, population coding and computation. *Nat Rev Neurosci*. 2006; 7: 358–366. doi: [10.1038/nrn1888](https://doi.org/10.1038/nrn1888) PMID: [16760916](https://pubmed.ncbi.nlm.nih.gov/16760916/)
7. Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. *Nat Neurosci*. 2006; 9: 1432–1438. doi: [10.1038/nn1790](https://doi.org/10.1038/nn1790) PMID: [17057707](https://pubmed.ncbi.nlm.nih.gov/17057707/)
8. Fiser J, Berkes P, Orbán G, Lengyel M. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn Sci*. 2010; 14: 119–130. doi: [10.1016/j.tics.2010.01.003](https://doi.org/10.1016/j.tics.2010.01.003) PMID: [20153683](https://pubmed.ncbi.nlm.nih.gov/20153683/)
9. Mathys C, Daunizeau J, Friston KJ, Stephan KE. A Bayesian foundation for individual learning under uncertainty. *Front Hum Neurosci*. 2011; 5: 39. doi: [10.3389/fnhum.2011.00039](https://doi.org/10.3389/fnhum.2011.00039) PMID: [21629826](https://pubmed.ncbi.nlm.nih.gov/21629826/)
10. Mathys CD, Lomakina EI, Daunizeau J, Iglesias S, Brodersen KH, Friston KJ, et al. Uncertainty in perception and the Hierarchical Gaussian Filter. *Front Hum Neurosci*. 2014; 8: 1–24.
11. Payzan-LeNestour E, Bossaerts P. Risk, Unexpected Uncertainty, and Estimation Uncertainty: Bayesian Learning in Unstable Settings. *PLoS Comput Biol*. 2011; 7: e1001048. doi: [10.1371/journal.pcbi.1001048](https://doi.org/10.1371/journal.pcbi.1001048) PMID: [21283774](https://pubmed.ncbi.nlm.nih.gov/21283774/)
12. Bach DR, Dolan RJ. Knowing how much you don't know: a neural organization of uncertainty estimates. *Nat Rev Neurosci*. 2012; 13: 572–586. doi: [10.1038/nrn3289](https://doi.org/10.1038/nrn3289) PMID: [22781958](https://pubmed.ncbi.nlm.nih.gov/22781958/)
13. Bland AR, Schaefer A. Different Varieties of Uncertainty in Human Decision-Making. *Front Neurosci*. 2012; 6: 1–11. doi: [10.3389/fnins.2012.00085](https://doi.org/10.3389/fnins.2012.00085) PMID: [22701401](https://pubmed.ncbi.nlm.nih.gov/22701401/)
14. Friston KJ, Shiner T, FitzGerald T, Galea JM, Adams R, Brown H, et al. Dopamine, affordance and active inference. *PLoS Comput Biol*. 2012; 8: e1002327. doi: [10.1371/journal.pcbi.1002327](https://doi.org/10.1371/journal.pcbi.1002327) PMID: [22241972](https://pubmed.ncbi.nlm.nih.gov/22241972/)
15. Iglesias S, Mathys C, Brodersen KH, Kasper L, Piccirelli M, den Ouden HEM, et al. Hierarchical Prediction Errors in Midbrain and Basal Forebrain during Sensory Learning. *Neuron*. 2013; 80: 519–530. doi: [10.1016/j.neuron.2013.09.009](https://doi.org/10.1016/j.neuron.2013.09.009) PMID: [24139048](https://pubmed.ncbi.nlm.nih.gov/24139048/)
16. Payzan-LeNestour E, Dunne S, Bossaerts P, O'Doherty JP. The Neural Representation of Unexpected Uncertainty during Value-Based Decision Making. *Neuron*. 2013; 79: 191–201. doi: [10.1016/j.neuron.2013.04.037](https://doi.org/10.1016/j.neuron.2013.04.037) PMID: [23849203](https://pubmed.ncbi.nlm.nih.gov/23849203/)
17. Vossel S, Bauer M, Mathys C, Adams RA, Dolan RJ, Stephan KE, et al. Cholinergic Stimulation Enhances Bayesian Belief Updating in the Deployment of Spatial Attention. *J Neurosci*. 2014; 34: 15735–15742. doi: [10.1523/JNEUROSCI.0091-14.2014](https://doi.org/10.1523/JNEUROSCI.0091-14.2014) PMID: [25411501](https://pubmed.ncbi.nlm.nih.gov/25411501/)
18. Yu A, Dayan P. Expected and unexpected uncertainty: ACh and NE in the neocortex. *Adv Neural Inf Process Syst*. 2003; 173–180.
19. Warren CM, Eldar E, van den Brink RL, Tona K-D, van der Wee NJ, Giltay EJ, et al. Catecholamine-Mediated Increases in Gain Enhance the Precision of Cortical Representations. *J Neurosci*. 2016; 36: 5699–5708. doi: [10.1523/JNEUROSCI.3475-15.2016](https://doi.org/10.1523/JNEUROSCI.3475-15.2016) PMID: [27225761](https://pubmed.ncbi.nlm.nih.gov/27225761/)

20. Berridge CW, Waterhouse BD. The locus coeruleus–noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Res Rev.* 2003; 42: 33–84. doi: [10.1016/S0165-0173\(03\)00143-7](https://doi.org/10.1016/S0165-0173(03)00143-7) PMID: [12668290](https://pubmed.ncbi.nlm.nih.gov/12668290/)
21. Hasselmo ME, Wyble BP, Wallenstein GV. Encoding and retrieval of episodic memories: role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus.* 1996; 6: 693–708. PMID: [9034856](https://pubmed.ncbi.nlm.nih.gov/9034856/)
22. Gil Z, Connors BW, Amitai Y. Differential Regulation of Neocortical Synapses by Neuromodulators and Activity. *Neuron.* 1997; 19: 679–686. doi: [10.1016/S0896-6273\(00\)80380-3](https://doi.org/10.1016/S0896-6273(00)80380-3) PMID: [9331357](https://pubmed.ncbi.nlm.nih.gov/9331357/)
23. Kimura F, Fukuda M, Tsumoto T. Acetylcholine suppresses the spread of excitation in the visual cortex revealed by optical recording: possible differential effect depending on the source of input. *Eur J Neurosci.* 1999; 11: 3597–3609. PMID: [10564367](https://pubmed.ncbi.nlm.nih.gov/10564367/)
24. Kobayashi M, Imamura K, Sugai T, Onoda N, Yamamoto M, Komai S, et al. Selective suppression of horizontal propagation in rat visual cortex by norepinephrine. *Eur J Neurosci.* 2000; 12: 264–272. PMID: [10651881](https://pubmed.ncbi.nlm.nih.gov/10651881/)
25. Yu AJ, Dayan P. Acetylcholine in cortical inference. *Neural Netw.* 2002; 15: 719–730. doi: [10.1016/S0893-6080\(02\)00058-8](https://doi.org/10.1016/S0893-6080(02)00058-8) PMID: [12371522](https://pubmed.ncbi.nlm.nih.gov/12371522/)
26. Hasselmo ME, McLaughly J. High acetylcholine levels set circuit dynamics for attention and encoding and low acetylcholine levels set dynamics for consolidation. In: Laurent Descarries KK and Mircea Steriade, editor. *Progress in Brain Research.* Elsevier; 2004. pp. 207–231. <http://www.sciencedirect.com/science/article/pii/S0079612303450152>
27. Sarter M, Hasselmo ME, Bruno JP, Givens B. Unraveling the attentional functions of cortical cholinergic inputs: interactions between signal-driven and cognitive modulation of signal detection. *Brain Res Rev.* 2005; 48: 98–111. doi: [10.1016/j.brainresrev.2004.08.006](https://doi.org/10.1016/j.brainresrev.2004.08.006) PMID: [15708630](https://pubmed.ncbi.nlm.nih.gov/15708630/)
28. Dayan P, Yu A. Norepinephrine and neural interrupts. *Adv Neural Inf Process Syst.* 2006; 18: 243.
29. Deco G, Thiele A. Cholinergic control of cortical network interactions enables feedback-mediated attentional modulation. *Eur J Neurosci.* 2011; 34: 146–157. doi: [10.1111/j.1460-9568.2011.07749.x](https://doi.org/10.1111/j.1460-9568.2011.07749.x) PMID: [21692884](https://pubmed.ncbi.nlm.nih.gov/21692884/)
30. Moran RJ, Campo P, Symmonds M, Stephan KE, Dolan RJ, Friston KJ. Free energy, precision and learning: the role of cholinergic neuromodulation. *J Neurosci.* 2013; 33: 8227–8236. doi: [10.1523/JNEUROSCI.4255-12.2013](https://doi.org/10.1523/JNEUROSCI.4255-12.2013) PMID: [23658161](https://pubmed.ncbi.nlm.nih.gov/23658161/)
31. Bowman EM, Brown VJ, Kertzman C, Schwarz U, Robinson DL. Covert orienting of attention in macaques. I. Effects of behavioral context. *J Neurophysiol.* 1993; 70: 431–443. PMID: [8360720](https://pubmed.ncbi.nlm.nih.gov/8360720/)
32. Downing CJ. Expectancy and visual-spatial attention: Effects on perceptual quality. *J Exp Psychol Hum Percept Perform.* 1988; 14: 188–202. doi: [10.1037/0096-1523.14.2.188](https://doi.org/10.1037/0096-1523.14.2.188) PMID: [2967876](https://pubmed.ncbi.nlm.nih.gov/2967876/)
33. Posner MI. Orienting of attention. *Q J Exp Psychol.* 1980; 32: 3–25. doi: [10.1080/00335558008248231](https://doi.org/10.1080/00335558008248231) PMID: [7367577](https://pubmed.ncbi.nlm.nih.gov/7367577/)
34. Vossel S, Mathys C, Daunizeau J, Bauer M, Driver J, Friston KJ, et al. Spatial Attention, Precision, and Bayesian Inference: A Study of Saccadic Response Speed. *Cereb Cortex.* 2014; 24: 1436–1450. doi: [10.1093/cercor/bhs418](https://doi.org/10.1093/cercor/bhs418) PMID: [23322402](https://pubmed.ncbi.nlm.nih.gov/23322402/)
35. Phillips JM, McAlonan K, Robb WGK, Brown VJ. Cholinergic neurotransmission influences covert orientation of visuospatial attention in the rat. *Psychopharmacology (Berl).* 2000; 150: 112.
36. Witte EA, Davidson MC, Marrocco RT. Effects of altering brain cholinergic activity on covert orienting of attention: comparison of monkey and human performance. *Psychopharmacology (Berl).* 1997; 132: 324.
37. Chiba AA, Bushnell PJ, Oshiro WM, Gallagher M. Selective removal of cholinergic neurons in the basal forebrain alters cued target detection. *Neuroreport.* 1999; 10: 3119–3123. PMID: [10549833](https://pubmed.ncbi.nlm.nih.gov/10549833/)
38. Voytko ML, Olton DS, Richardson RT, Gorman LK, Tobin JR, Price DL. Basal forebrain lesions in monkeys disrupt attention but not learning and memory [published erratum appears in *J Neurosci* 1995 Mar;15(3): following table of contents]. *J Neurosci.* 1994; 14: 167–186. PMID: [8283232](https://pubmed.ncbi.nlm.nih.gov/8283232/)
39. Parasuraman R, Greenwood PM, Haxby JV, Grady CL. Visuospatial attention in dementia of the Alzheimer type. *Brain.* 1992; 115: 711–733. PMID: [1628198](https://pubmed.ncbi.nlm.nih.gov/1628198/)
40. Bucci DJ, Holland PC, Gallagher M. Removal of Cholinergic Input to Rat Posterior Parietal Cortex Disrupts Incremental Processing of Conditioned Stimuli. *J Neurosci.* 1998; 18: 8038–8046. PMID: [9742170](https://pubmed.ncbi.nlm.nih.gov/9742170/)
41. Clark CR, Geffen GM, Geffen LB. Catecholamines and the covert orientation of attention in humans. *Neuropsychologia.* 1989; 27: 131–139. PMID: [2538773](https://pubmed.ncbi.nlm.nih.gov/2538773/)
42. Witte EA, Marrocco RT. Alteration of brain noradrenergic activity in rhesus monkeys affects the alerting component of covert orienting. *Psychopharmacology (Berl).* 1997; 132: 315–323.

43. Arnsten AFT, Contant TA. Alpha-2 adrenergic agonists decrease distractibility in aged monkeys performing the delayed response task. *Psychopharmacology (Berl)*. 1992; 108: 159–169.
44. Bouret S, Sara SJ. Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends Neurosci*. 2005; 28: 574–582. doi: [10.1016/j.tins.2005.09.002](https://doi.org/10.1016/j.tins.2005.09.002) PMID: [16165227](https://pubmed.ncbi.nlm.nih.gov/16165227/)
45. Coull JT, Sahakian BJ, Middleton HC, Young AH, Park SB, McShane RH, et al. Differential effects of clonidine, haloperidol, diazepam and tryptophan depletion on focused attention and attentional search. *Psychopharmacology (Berl)*. 1995; 121: 222–230.
46. Dayan P, Yu AJ. Phasic norepinephrine: a neural interrupt signal for unexpected events. *Netw Comput Neural Syst*. 2006; 17: 335–350.
47. Smith AP, Wilson SJ, Glue P, Nutt DJ. The effects and after effects of the α_2 -adrenoceptor antagonist idazoxan on mood, memory and attention in normal volunteers. *J Psychopharmacol (Oxf)*. 1992; 6: 376–381.
48. Bouret S, Sara SJ. Reward expectation, orientation of attention and locus coeruleus-medial frontal cortex interplay during learning. *Eur J Neurosci*. 2004; 20: 791–802. doi: [10.1111/j.1460-9568.2004.03526.x](https://doi.org/10.1111/j.1460-9568.2004.03526.x) PMID: [15255989](https://pubmed.ncbi.nlm.nih.gov/15255989/)
49. Aston-Jones G, Rajkowski J, Kubiak P. Conditioned responses of monkey locus coeruleus neurons anticipate acquisition of discriminative behavior in a vigilance task. *Neuroscience*. 1997; 80: 697–715. PMID: [9276487](https://pubmed.ncbi.nlm.nih.gov/9276487/)
50. Sara SJ, Segal M. Plasticity of sensory responses of locus coeruleus neurons in the behaving rat: implications for cognition. *Prog Brain Res*. 1991; 88: 571–585. PMID: [1813935](https://pubmed.ncbi.nlm.nih.gov/1813935/)
51. Devauges V, Sara SJ. Activation of the noradrenergic system facilitates an attentional shift in the rat. *Behav Brain Res*. 1990; 39: 19–28. PMID: [2167690](https://pubmed.ncbi.nlm.nih.gov/2167690/)
52. McGaughy J, Ross RS, Eichenbaum H. Noradrenergic, but not cholinergic, deafferentation of prefrontal cortex impairs attentional set-shifting. *Neuroscience*. 2008; 153: 63–71. doi: [10.1016/j.neuroscience.2008.01.064](https://doi.org/10.1016/j.neuroscience.2008.01.064) PMID: [18355972](https://pubmed.ncbi.nlm.nih.gov/18355972/)
53. Joshi S, Li Y, Kalwani RM, Gold JI. Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron*. 2016; 89: 221–234. doi: [10.1016/j.neuron.2015.11.028](https://doi.org/10.1016/j.neuron.2015.11.028) PMID: [26711118](https://pubmed.ncbi.nlm.nih.gov/26711118/)
54. Nassar MR, Rumsey KM, Wilson RC, Parikh K, Heasley B, Gold JI. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat Neurosci*. 2012; 15: 1040–1046. doi: [10.1038/nn.3130](https://doi.org/10.1038/nn.3130) PMID: [22660479](https://pubmed.ncbi.nlm.nih.gov/22660479/)
55. Preuschoff K, 't Hart BM, Einhäuser W. Pupil Dilation Signals Surprise: Evidence for Noradrenaline's Role in Decision Making. *Front Neurosci*. 2011; 5: 1–12. doi: [10.3389/fnins.2011.00115](https://doi.org/10.3389/fnins.2011.00115) PMID: [21994487](https://pubmed.ncbi.nlm.nih.gov/21994487/)
56. Hikosaka O, Isoda M. Switching from automatic to controlled behavior: cortico-basal ganglia mechanisms. *Trends Cogn Sci*. 2010; 14: 154–161. doi: [10.1016/j.tics.2010.01.006](https://doi.org/10.1016/j.tics.2010.01.006) PMID: [20181509](https://pubmed.ncbi.nlm.nih.gov/20181509/)
57. Isoda M, Hikosaka O. Cortico-basal ganglia mechanisms for overcoming innate, habitual and motivational behaviors. *Eur J Neurosci*. 2011; 33: 2058–2069. doi: [10.1111/j.1460-9568.2011.07698.x](https://doi.org/10.1111/j.1460-9568.2011.07698.x) PMID: [21645101](https://pubmed.ncbi.nlm.nih.gov/21645101/)
58. Bestmann S, Ruge D, Rothwell J, Galea JM. The Role of Dopamine in Motor Flexibility. *J Cogn Neurosci*. 2014; 27: 365–376. doi: [10.1162/jocn_a_00706](https://doi.org/10.1162/jocn_a_00706) PMID: [25170792](https://pubmed.ncbi.nlm.nih.gov/25170792/)
59. Galea JM, Bestmann S, Beigi M, Jahanshahi M, Rothwell JC. Action reprogramming in Parkinson's disease: response to prediction error is modulated by levels of dopamine. *J Neurosci*. 2012; 32: 542–550. doi: [10.1523/JNEUROSCI.3621-11.2012](https://doi.org/10.1523/JNEUROSCI.3621-11.2012) PMID: [22238089](https://pubmed.ncbi.nlm.nih.gov/22238089/)
60. Cools R, Barker RA, Sahakian BJ, Robbins TW. Mechanisms of cognitive set flexibility in Parkinson's disease. *Brain*. 2001; 124: 2503–2512. PMID: [11701603](https://pubmed.ncbi.nlm.nih.gov/11701603/)
61. Cools R, Frank MJ, Gibbs SE, Miyakawa A, Jagust W, D'Esposito M. Striatal dopamine predicts outcome-specific reversal learning and its sensitivity to dopaminergic drug administration. *J Neurosci*. 2009; 29: 1538–1543. doi: [10.1523/JNEUROSCI.4467-08.2009](https://doi.org/10.1523/JNEUROSCI.4467-08.2009) PMID: [19193900](https://pubmed.ncbi.nlm.nih.gov/19193900/)
62. van Holstein M, Aarts E, van der Schaaf ME, Geurts DE, Verkes RJ, Franke B, et al. Human cognitive flexibility depends on dopamine D2 receptor signaling. *Psychopharmacology (Berl)*. 2011; 218: 567–578.
63. Stelzel C, Basten U, Montag C, Reuter M, Fiebach CJ. Frontostriatal involvement in task switching depends on genetic differences in d2 receptor density. *J Neurosci*. 2010; 30: 14205–14212. doi: [10.1523/JNEUROSCI.1062-10.2010](https://doi.org/10.1523/JNEUROSCI.1062-10.2010) PMID: [20962241](https://pubmed.ncbi.nlm.nih.gov/20962241/)
64. Stelzel C, Fiebach CJ, Cools R, Tafazoli S, D'Esposito M. Dissociable fronto-striatal effects of dopamine D2 receptor stimulation on cognitive versus motor flexibility. *Cortex*. 2013; 49: 2799–2811. doi: [10.1016/j.cortex.2013.04.002](https://doi.org/10.1016/j.cortex.2013.04.002) PMID: [23660437](https://pubmed.ncbi.nlm.nih.gov/23660437/)

65. Cools AR, van den Bercken JH, Horstink MW, Van Spaendonck KP, Berger HJ. Cognitive and motor shifting aptitude disorder in Parkinson's disease. *J Neurol Neurosurg Psychiatry*. 1984; 47: 443–453. PMID: [6736974](#)
66. Beatty WW, Monson N. Problem solving in Parkinson's disease: Comparison of performance on the Wisconsin and California Card Sorting Tests. *J Geriatr Psychiatry Neurol*. 1990; 3: 163–171. PMID: [2282133](#)
67. Cools R, Barker RA, Sahakian BJ, Robbins TW. Enhanced or Impaired Cognitive Function in Parkinson's Disease as a Function of Dopaminergic Medication and Task Demands. *Cereb Cortex*. 2001; 11: 1136–1143. doi: [10.1093/cercor/11.12.1136](#) PMID: [11709484](#)
68. Varazzani C, San-Galli A, Gilardeau S, Bouret S. Noradrenaline and Dopamine Neurons in the Reward/Effort Trade-Off: A Direct Electrophysiological Comparison in Behaving Monkeys. *J Neurosci*. 2015; 35: 7866–7877. doi: [10.1523/JNEUROSCI.0454-15.2015](#) PMID: [25995472](#)
69. Brown SBRE, van der Wee NJA, van Noorden MS, Giltay EJ, Nieuwenhuis S. Noradrenergic and cholinergic modulation of late ERP responses to deviant stimuli. *Psychophysiology*. 2015; 52: 1620–1631. doi: [10.1111/psyp.12544](#) PMID: [26352794](#)
70. Diaconescu AO, Mathys C, Weber LAE, Daunizeau J, Kasper L, Lomakina EI, et al. Inferring on the Intentions of Others by Hierarchical Bayesian Learning. *PLoS Comput Biol*. 2014; 10: e1003810. doi: [10.1371/journal.pcbi.1003810](#) PMID: [25187943](#)
71. Hauser TU, Iannaccone R, Ball J, Mathys C, Brandeis D, Walitza S, et al. Role of the medial prefrontal cortex in impaired decision making in juvenile attention-deficit/hyperactivity disorder. *JAMA Psychiatry*. 2014; 71: 1165–1173. doi: [10.1001/jamapsychiatry.2014.1093](#) PMID: [25142296](#)
72. Vossel S, Mathys C, Daunizeau J, Bauer M, Driver J, Friston KJ, et al. Spatial Attention, Precision, and Bayesian Inference: A Study of Saccadic Response Speed. *Cereb Cortex*. 2014; 24: 1436–1450. doi: [10.1093/cercor/bhs418](#) PMID: [23322402](#)
73. Vossel S, Mathys C, Stephan KE, Friston KJ. Cortical Coupling Reflects Bayesian Belief Updating in the Deployment of Spatial Attention. *J Neurosci*. 2015; 35: 11532–11542. doi: [10.1523/JNEUROSCI.1382-15.2015](#) PMID: [26290231](#)
74. de Berker AO, Rutledge RB, Mathys C, Marshall L, Cross GF, Dolan RJ, et al. Computations of uncertainty mediate acute stress responses in humans. *Nat Commun*. 2016; 7: 10996. doi: [10.1038/ncomms10996](#) PMID: [27020312](#)
75. Daunizeau J, den Ouden HEM, Pessiglione M, Kiebel SJ, Stephan KE, Friston KJ. Observing the Observer (I): Meta-Bayesian Models of Learning and Decision-Making. *PLoS ONE*. 2010; 5: e15554. doi: [10.1371/journal.pone.0015554](#) PMID: [21179480](#)
76. Harrison LM, Duggins A, Friston KJ. Encoding uncertainty in the hippocampus. *Neural Netw*. 2006; 19: 535–546. doi: [10.1016/j.neunet.2005.11.002](#) PMID: [16527453](#)
77. Chowdhury R, Guitart-Masip M, Bunzeck N, Dolan RJ, Düzel E. Dopamine modulates episodic memory persistence in old age. *J Neurosci*. 2012; 32: 14193–14204. doi: [10.1523/JNEUROSCI.1278-12.2012](#) PMID: [23055489](#)
78. Knecht S, Breitenstein C, Bushuven S, Wailke S, Kamping S, Flöel A, et al. Levodopa: faster and better word learning in normal humans. *Ann Neurol*. 2004; 56: 20–26. doi: [10.1002/ana.20125](#) PMID: [15236398](#)
79. Newhouse PA, Potter A, Singh A. Effects of nicotinic stimulation on cognitive performance. *Curr Opin Pharmacol*. 2004; 4: 36–46. doi: [10.1016/j.coph.2003.11.001](#) PMID: [15018837](#)
80. Botvinick MM, Braver TS, Barch DM, Carter CS, Cohen JD. Conflict monitoring and cognitive control. *Psychol Rev*. 2001; 108: 624. PMID: [11488380](#)
81. Cavanagh JF, Sanguinetti JL, Allen JJ, Sherman SJ, Frank MJ. The subthalamic nucleus contributes to post-error slowing. *J Cogn Neurosci*. 2014; 26: 2637–2644. doi: [10.1162/jocn_a_00659](#) PMID: [24800632](#)
82. Gehring WJ, Fencsik DE. Functions of the medial frontal cortex in the processing of conflict and errors. *J Neurosci*. 2001; 21: 9430–9437. PMID: [11717376](#)
83. Rabbitt PM. Errors and error correction in choice-response tasks. *J Exp Psychol*. 1966; 71: 264. PMID: [5948188](#)
84. den Ouden HE, Daunizeau J, Roiser J, Friston KJ, Stephan KE. Striatal prediction error modulates cortical coupling. *J Neurosci*. 2010; 30: 3210–3219. doi: [10.1523/JNEUROSCI.4458-09.2010](#) PMID: [20203180](#)
85. Vankov A, Hervé-Minvielle A, Sara SJ. Response to novelty and its rapid habituation in locus coeruleus neurons of the freely exploring rat. *Eur J Neurosci*. 1995; 7: 1180–1187. PMID: [7582091](#)

86. Murphy PR, O'Connell RG, O'Sullivan M, Robertson IH, Balsters JH. Pupil diameter covaries with BOLD activity in human locus coeruleus. *Hum Brain Mapp.* 2014; 35: 4140–4154. doi: [10.1002/hbm.22466](https://doi.org/10.1002/hbm.22466) PMID: [24510607](https://pubmed.ncbi.nlm.nih.gov/24510607/)
87. Aston-Jones G, Cohen JD. Adaptive gain and the role of the locus coeruleus–norepinephrine system in optimal performance. *J Comp Neurol.* 2005; 493: 99–110. doi: [10.1002/cne.20723](https://doi.org/10.1002/cne.20723) PMID: [16254995](https://pubmed.ncbi.nlm.nih.gov/16254995/)
88. Aston-Jones G, Rajkowski J, Kubiak P, Alexinsky T. Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *J Neurosci.* 1994; 14: 4467–4480. PMID: [8027789](https://pubmed.ncbi.nlm.nih.gov/8027789/)
89. Usher M, Cohen JD, Servan-Schreiber D, Rajkowski J, Aston-Jones G. The role of locus coeruleus in the regulation of cognitive performance. *Science.* 1999; 283: 549–554. PMID: [9915705](https://pubmed.ncbi.nlm.nih.gov/9915705/)
90. Jepma M, Te Beek ET, Wagenmakers E-J, Van Gerven JMA, Nieuwenhuis S. The role of the noradrenergic system in the exploration–exploitation trade-off: a psychopharmacological study. *Front Hum Neurosci.* 2010; 4: 170. doi: [10.3389/fnhum.2010.00170](https://doi.org/10.3389/fnhum.2010.00170) PMID: [21206527](https://pubmed.ncbi.nlm.nih.gov/21206527/)
91. Bentley P, Driver J, Dolan RJ. Cholinergic modulation of cognition: Insights from human pharmacological functional neuroimaging. *Prog Neurobiol.* 2011; 94: 360–388. doi: [10.1016/j.pneurobio.2011.06.002](https://doi.org/10.1016/j.pneurobio.2011.06.002) PMID: [21708219](https://pubmed.ncbi.nlm.nih.gov/21708219/)
92. Jones DNC, Higgins GA. Effect of scopolamine on visual attention in rats. *Psychopharmacology (Berl).* 1995; 120: 142–149.
93. O'Neill J, Siembieda DW, Crawford KC, Halgren E, Fisher A, Fitten LJ. Reduction in distractibility with AF102B and THA in the macaque. *Pharmacol Biochem Behav.* 2003; 76: 301–306. PMID: [14592682](https://pubmed.ncbi.nlm.nih.gov/14592682/)
94. Prendergast MA, Jackson WJ, Terry AV Jr, Decker MW, Arneric SP, Buccafusco JJ. Central nicotinic receptor agonists ABT-418, ABT-089, and (–)-nicotine reduce distractibility in adult monkeys. *Psychopharmacology (Berl).* 1998; 136: 50–58.
95. Terry AV, Risbrough VB, Buccafusco JJ, Menzaghi F. Effects of (±)-4-[[2-(1-methyl-2-pyrrolidinyl)ethyl]thio]phenol hydrochloride (SIB-1553A), a selective ligand for nicotinic acetylcholine receptors, in tests of visual attention and distractibility in rats and monkeys. *J Pharmacol Exp Ther.* 2002; 301: 284–292. PMID: [11907185](https://pubmed.ncbi.nlm.nih.gov/11907185/)
96. Korchounov A, Ziemann U. Neuromodulatory Neurotransmitters Influence LTP-Like Plasticity in Human Cortex: A Pharmacology-TMS Study. *Neuropsychopharmacology.* 2011; 36: 1894–1902. doi: [10.1038/npp.2011.75](https://doi.org/10.1038/npp.2011.75) PMID: [21544070](https://pubmed.ncbi.nlm.nih.gov/21544070/)
97. Meintschel F, Ziemann U. Modification of practice-dependent plasticity in human motor cortex by neuromodulators. *Cereb Cortex.* 2006; 16: 1106–1115. doi: [10.1093/cercor/bhj052](https://doi.org/10.1093/cercor/bhj052) PMID: [16221926](https://pubmed.ncbi.nlm.nih.gov/16221926/)
98. Ziemann U, Tergau F, Bruns D, Baudewig J, Paulus W. Changes in human motor cortex excitability induced by dopaminergic and anti-dopaminergic drugs. *Electroencephalogr Clin Neurophysiol Mot Control.* 1997; 105: 430–437.
99. Patton JH, Stanford MS, Barratt ES. Factor structure of the Barratt impulsiveness scale. *J Clin Psychol.* 1995; 768–74. PMID: [8778124](https://pubmed.ncbi.nlm.nih.gov/8778124/)
100. Blais A-R, Weber EU. A domain-specific risk-taking (DOSPERT) scale for adult populations. *Judgm Decis Mak.* 2006; 1: 33–47.
101. Broadbent DE, Cooper PF, FitzGerald P, Parkes KR. The cognitive failures questionnaire (CFQ) and its correlates. *Br J Clin Psychol.* 1982; 21: 1–16. PMID: [7126941](https://pubmed.ncbi.nlm.nih.gov/7126941/)
102. Bond A, Lader M. The use of analogue scales in rating subjective feelings. *Br J Med Psychol.* 1974; 47: 211–218.
103. Ellis BW, Johns MW, Lancaster R, Raptopoulos P, Angelopoulos N, Priest RG. The St. Mary's Hospital sleep questionnaire: a study of reliability. *Sleep.* 1980; 4: 93–97.
104. Nassar MR, Wilson RC, Heasley B, Gold JL. An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J Neurosci.* 2010; 30: 12366–12378. doi: [10.1523/JNEUROSCI.0822-10.2010](https://doi.org/10.1523/JNEUROSCI.0822-10.2010) PMID: [20844132](https://pubmed.ncbi.nlm.nih.gov/20844132/)
105. Wilson RC, Nassar MR, Gold JL. A Mixture of Delta-Rules Approximation to Bayesian Inference in Change-Point Problems. *PLoS Comput Biol.* 2013; 9: e1003150. doi: [10.1371/journal.pcbi.1003150](https://doi.org/10.1371/journal.pcbi.1003150) PMID: [23935472](https://pubmed.ncbi.nlm.nih.gov/23935472/)
106. Mars RB, Debener S, Gladwin TE, Harrison LM, Haggard P, Rothwell JC, et al. Trial-by-trial fluctuations in the event-related electroencephalogram reflect dynamic changes in the degree of surprise. *J Neurosci.* 2008; 28: 12539–12545. doi: [10.1523/JNEUROSCI.2925-08.2008](https://doi.org/10.1523/JNEUROSCI.2925-08.2008) PMID: [19020046](https://pubmed.ncbi.nlm.nih.gov/19020046/)
107. Strange BA, Duggins A, Penny W, Dolan RJ, Friston KJ. Information theory, novelty and hippocampal responses: unpredicted or unpredictable? *Neural Netw.* 2005; 18: 225–230. doi: [10.1016/j.neunet.2004.12.004](https://doi.org/10.1016/j.neunet.2004.12.004) PMID: [15896570](https://pubmed.ncbi.nlm.nih.gov/15896570/)

108. Rigoux L, Stephan KE, Friston KJ, Daunizeau J. Bayesian model selection for group studies—revisited. *Neuroimage*. 2014; 84: 971–985. doi: [10.1016/j.neuroimage.2013.08.065](https://doi.org/10.1016/j.neuroimage.2013.08.065) PMID: [24018303](https://pubmed.ncbi.nlm.nih.gov/24018303/)
109. Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ. Bayesian model selection for group studies. *Neuroimage*. 2009; 46: 1004–1017. doi: [10.1016/j.neuroimage.2009.03.025](https://doi.org/10.1016/j.neuroimage.2009.03.025) PMID: [19306932](https://pubmed.ncbi.nlm.nih.gov/19306932/)
110. Daunizeau J, Adam V, Rigoux L. VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioural data. *PLoS Comput Biol*. 2014; 10: e1003441. doi: [10.1371/journal.pcbi.1003441](https://doi.org/10.1371/journal.pcbi.1003441) PMID: [24465198](https://pubmed.ncbi.nlm.nih.gov/24465198/)