

Noradrenergic and Cholinergic Modulation of Belief Updating

Marieke Jepma^{1,2*}, Stephen B. R. E. Brown^{1*}, Peter R. Murphy^{1,3}, Stephany C. Koelewijn⁴, Boukje de Vries⁴, Arn M. van den Maagdenberg⁴, and Sander Nieuwenhuis¹

Abstract

■ To make optimal predictions in a dynamic environment, the impact of new observations on existing beliefs—that is, the learning rate—should be guided by ongoing estimates of change and uncertainty. Theoretical work has proposed specific computational roles for various neuromodulatory systems in the control of learning rate, but empirical evidence is still sparse. The aim of the current research was to examine the role of the noradrenergic and cholinergic systems in learning rate regulation. First, we replicated our recent findings that the centroparietal P3 component of the EEG—an index of phasic catecholamine release in the cortex—predicts trial-to-trial variability in learning rate and mediates the effects of surprise

and belief uncertainty on learning rate (Study 1, $n = 17$). Second, we found that pharmacological suppression of either norepinephrine or acetylcholine activity produced baseline-dependent effects on learning rate following nonobvious changes in an outcome-generating process (Study 1). Third, we identified two genes, coding for $\alpha 2A$ receptor sensitivity (*ADRA2A*) and norepinephrine reuptake (*NET*), as promising targets for future research on the genetic basis of individual differences in learning rate (Study 2, $n = 137$). Our findings suggest a role for the noradrenergic and cholinergic systems in belief updating and underline the importance of studying interactions between different neuromodulatory systems. ■

INTRODUCTION

According to contemporary ideas in neuroscience, the brain's overarching function is to construct a model of its environment to optimally predict its sensory inputs and thus minimize the amount of surprise or "prediction error" (Friston, 2010; Doya, 2007; Rao & Ballard, 1999). In a dynamically changing environment, accurate prediction requires the continuous updating of this internal model in response to new observations. There is considerable evidence that optimal belief updating can be captured by Bayesian learning models, in which the impact of each new observation on existing beliefs (i.e., the "learning rate") depends on the unexpectedness and the precision of the observation (i.e., the extent to which it can be used to predict future outcomes; Nassar, Wilson, Heasly, & Gold, 2010; Friston, 2009; Behrens, Woolrich, Walton, & Rushworth, 2007). The unexpectedness and precision of observations are in turn determined by various types of uncertainty related to environmental stochasticity and volatility and the observer's ignorance about the current state of the environment. These various types of uncertainty can be estimated

based on the sequence of previous observations (Nassar et al., 2010; Behrens et al., 2007).

Neuroimaging studies have identified the neural correlates of various elements of this belief-updating process (Kolossa, Kopp, & Fingscheidt, 2015; McGuire, Nassar, Gold, & Kable, 2014; Iglesias et al., 2013; O'Reilly et al., 2013; Payzan-LeNestour, Dunne, Bossaerts, & O'Doherty, 2013; Chumbley et al., 2012; Behrens et al., 2007). Furthermore, theoretical and computational modeling work has proposed specific roles for different neuromodulatory systems. There is an extensive literature linking the dopaminergic system to the encoding of reward prediction errors (Daw & Doya, 2006; Schultz, Dayan, & Montague, 1997). More recently, it has been suggested that dopamine (DA) has a key role in the precision-weighting of prediction errors—not restricted to the reward domain—and thereby controls learning rate (Friston et al., 2012; Friston, 2009). Norepinephrine (NE) has been proposed to encode "unexpected uncertainty" arising from unanticipated changes in a task context (Yu & Dayan, 2005), which may specifically drive belief updating following sudden environmental change. In contrast, acetylcholine (ACh) has been proposed to encode "expected uncertainty," arising from stochasticity inherent in the environment (noise or risk) and/or from known ignorance about the environment (estimation uncertainty or ambiguity; Yu & Dayan, 2005). Whereas high levels of stochasticity in otherwise stable environments

¹Leiden University, ²University of Amsterdam, ³University Medical Center Hamburg-Eppendorf, ⁴Leiden University Medical Centre
*Co-first author.

warrant a low learning rate to avoid excessive belief updating, a high degree of ignorance warrants a high learning rate to facilitate learning about the environment. Other computational accounts have proposed that ACh controls the learning rate in reinforcement-learning algorithms (Doya, 2002), and facilitates perceptual learning by boosting the bottom-up signaling of sensory stimuli (Moran et al., 2013). In addition, physiological studies have shown that ACh and NE both boost thalamocortical (afferent) relative to intracortical (intrinsic) synaptic transmission, which presumably enhances the impact of new incoming stimuli relative to preexisting internal representations and thus facilitates learning of new information (Hasselmo, 1995, 2006; Yu & Dayan, 2002; Hsieh, Cruikshank, & Metherate, 2000; Kimura, Fukuda, & Tsumoto, 1999).

Although the dopaminergic, noradrenergic, and cholinergic systems have all been associated with the control of learning rate, empirical evidence for their roles in belief updating is sparse (but see Marshall et al., 2016). The notion that DA controls learning rate is supported by findings that dopaminergic genotype predicts individual differences in learning rate (Set et al., 2014; Krugel, Biele, Mohr, Li, & Heekeren, 2009). A role for the noradrenergic system in belief updating following abrupt task changes is broadly consistent with evidence that pharmacological manipulations and lesions of this system in animals affect reversal learning and attentional set shifting (Seu, Lang, Rivera, & Jentsch, 2009; McGaughy, Ross, & Eichenbaum, 2008; Newman, Darling, & McGaughy, 2008; Lapid, Bondi, & Morilak, 2007; Tait et al., 2007; Lapid & Morilak, 2006; Devauges & Sara, 1990). More specific, although indirect, evidence for a role of the noradrenergic system in learning rate regulation has been provided by studies that used pupil dilation as an index of locus coeruleus (LC; the main noradrenergic nucleus in the brain) activity in humans (Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015; Silvetti, Seurinck, van Bochove, & Verguts, 2013; Nassar et al., 2012). Finally, pharmacological studies in humans have provided evidence that cholinergic stimulation and suppression respectively promote and suppress belief updating about contextual probabilities (Marshall et al., 2016; Vossel et al., 2014).

We recently reported evidence for a role of the catecholamine (NE and DA) systems in belief updating (Jepma et al., 2016), using a “predictive inference” task that requires participants to repeatedly predict the next location on a number line in a sequence (Nassar et al., 2010). In this task, the location outcome on each trial is drawn from a Gaussian distribution that is centered on a specific point of the number line but occasionally shifts to a new point on the number line; hence the outcome-generating process contains both noise and unpredictable “change points.” Trial-by-trial variation in the outcome-evoked centroparietal P3 component of the EEG predicted learning rate in this task and formally mediated the effects of model-based estimates of

outcome-evoked surprise and preexisting belief uncertainty on learning rate (Jepma et al., 2016). In light of evidence for an association between the centroparietal P3 and stimulus-evoked phasic catecholamine release in the cortex (Polich, 2007; Nieuwenhuis, Aston-Jones, & Cohen, 2005; Pineda, Foote, & Neville, 1989), these findings are consistent with a role of the catecholamine systems in learning rate regulation. Corroborating this idea, we found that a pharmacological manipulation of catecholamine activity—using the NE transporter blocker atomoxetine—affected learning rate following change points, in a way that depended on participants’ baseline (placebo session) learning rate (Jepma et al., 2016).

The present research had three aims. First, we examined whether we could replicate the relationships between P3 amplitude, surprise and uncertainty, and learning rate found in our previous study (Study 1). Second, we aimed to gain further insight into the specific role of the noradrenergic system in belief updating. To this end, we examined the effects of clonidine—a centrally acting α_2 agonist that predominantly acts to attenuate baseline NE activity by agonizing presynaptic α_2 receptors (Svensson, Bunney, & Aghajanian, 1975)—on performance of the predictive inference task (Study 1). Clonidine has been shown to reduce P3 amplitude in previous studies (Brown et al., 2016; Nieuwenhuis et al., 2005; but see Brown, van der Wee, van Noorden, Giltay, & Nieuwenhuis, 2015); hence, we reasoned that it may also reduce learning rate via its effects on the α_2 receptor. In addition, we examined whether interindividual variation in noradrenergic genotypes were associated with individual differences in learning rate (Study 2). Third, we aimed to examine the role of the cholinergic system in belief updating. Therefore, in Study 1, we compared the effects of clonidine to those of scopolamine, an anticholinergic agent that blocks the activity of the muscarinic ACh receptor. Based on the proposed computational roles of ACh and NE in encoding, respectively, expected and unexpected uncertainty (Yu & Dayan, 2005), one could predict that scopolamine influences the effects of noise and estimation uncertainty (both forms of expected uncertainty) on learning rate, whereas clonidine influences the effect of change points (unexpected uncertainty) on learning rate. However, given the evidence that ACh and NE both enhance the influence of incoming sensory stimuli relative to internal representations, due to their effects on intrinsic versus afferent synaptic transmission (Hasselmo, 1995), an alternative prediction is that pharmacological suppression of either system will result in an overall reduction of learning rate. Similar effects of clonidine and scopolamine would also be consistent with recent findings that these two drugs similarly affect measures of temporal attention, perceptual sensitivity, and ERPs, presumably due to bidirectional interactions between the basal forebrain and the LC (Brown et al., 2016; Brown, Tona, et al., 2015; Brown, van der Wee, et al., 2015).

STUDY 1

Methods

Participants

Eighteen healthy individuals (mean age = 21 years, range = 18–26 years; 15 women) took part in three experimental sessions, separated by 1 week, in return for €140. Exclusion criteria included history or presence of neurological or psychiatric disorders; use of prescribed medication, smoking, pregnancy, and a systolic blood pressure below 100 mm Hg; a diastolic blood pressure below 70 mm Hg; or a heart rate below 65 beats per minute in rest. Participants were instructed to abstain from using recreational drugs, caffeine, or alcohol from 15 hr before the start of each session. All participants provided written informed consent, and the study was approved by the medical ethics committee of the Leiden University Medical Center.

Participants received a single oral dose of clonidine, a single oral dose of scopolamine (1.2 mg), and a placebo in a randomized, double-blind, counterbalanced double-dummy crossover design. The first 11 participants received a clonidine dose of 175 µg. As the 11th participant showed an unexpected large drop in blood pressure of 35 mm Hg systolic, but without clinical consequences, 60 min after the ingestion of clonidine 175 µg (blind was broken by the supervising physician), we reduced the dose of clonidine to 150 µg for the final seven participants.

One participant was excluded from the analyses because of poor performance on the predictive inference task in all sessions (see below); hence, our analyses of this task were based on 17 participants.

Procedure

During each session, participants received a capsule of clonidine or placebo at 09:35 a.m. and a capsule of scopolamine or placebo at 10:35 a.m. The different kinetic profiles of clonidine and scopolamine necessitated administrations at different times before testing. This double-dummy design resulted in one clonidine session (i.e., clonidine verum plus scopolamine placebo), one scopolamine session (clonidine placebo plus scopolamine verum), and one placebo session (clonidine plus scopolamine placebos). Treatment order was counterbalanced across participants.

At the start of each session ($t = -20$ min), a peripheral intravenous cannula was placed and connected to an intravenous 0.9% NaCl (saline) drip to be able to increase blood pressure through volume expansion and to have an entryway to administer escape medication in the case of a severe drop in tension and/or heart rate. Furthermore, three cardio electrodes were applied to the participant's chest and connected to an electrocardiography monitor. At $t = 0$ min, participants ingested a microcrystalline cellulose-filled capsule with either clonidine or

placebo. At $t = 60$ min, participants ingested a microcrystalline cellulose-filled capsule with either scopolamine or placebo. At $t = 120$ min, participants performed the predictive inference task, which lasted approximately 30 min. This task was part of a test battery of four cognitive tasks, the other three tasks have been reported elsewhere (Brown et al., 2016; Brown, Tona, et al., 2015; Brown, van der Wee, et al., 2015).

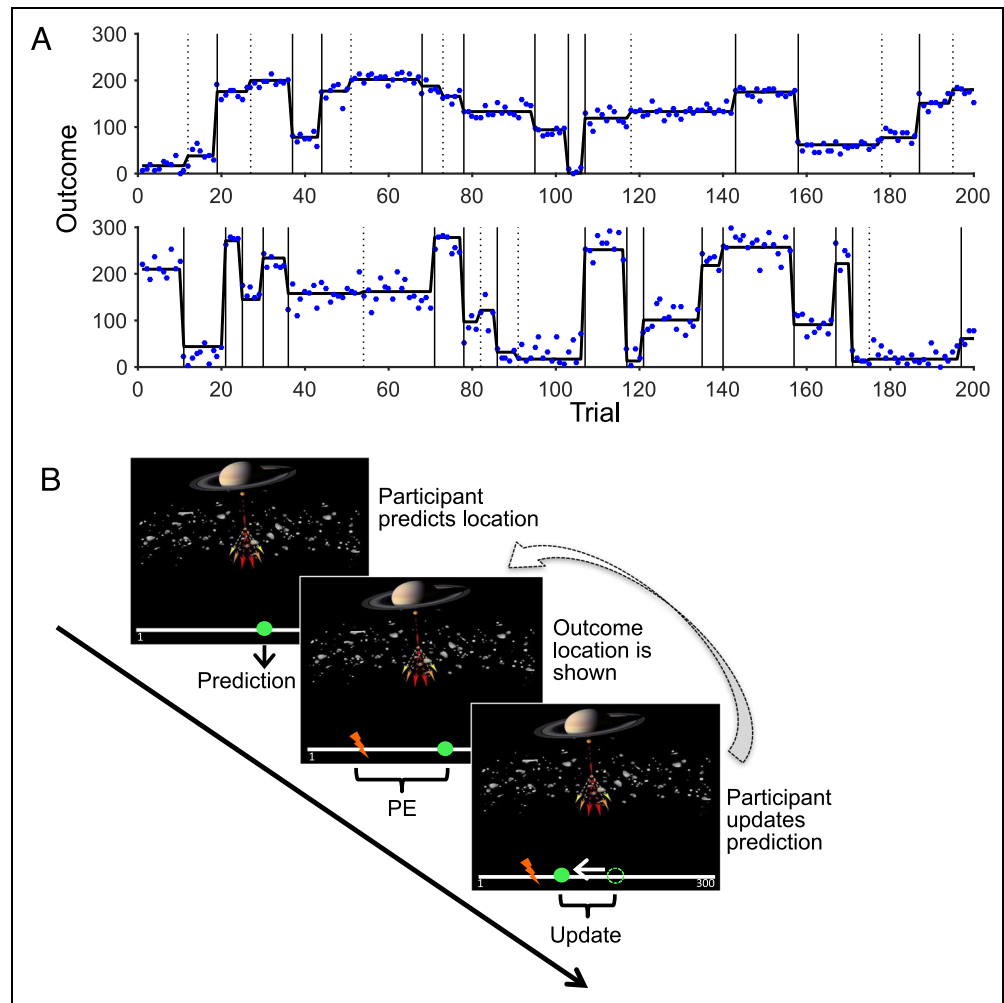
To measure the effects of clonidine and scopolamine on alertness, we administered a 40-trial simple RT (SRT) task upon participants' arrival in the lab, as well as right before and after performance of the predictive inference task. Participants had to respond as quickly as possible whenever a white circle appeared on the computer screen. SOA was jittered between 500 and 1250 msec, with a mean of 1000 msec. This task lasted less than 2 min. In addition, blood pressure and heart rate were measured four times per hour from $t = 0$ onward with an Omron M10-IT automatic sphygmomanometer. Participants' fitness was checked at $t = 240$ min, and participants were sent home via public transportation if their blood pressure and heart rate were close to the values measured at $t = -20$ min. If their blood pressure and heart rate had not returned to normal yet, they were kept for further observation.

Predictive Inference Task

During this task, participants observed locations on a horizontal number line—one location per trial—and were asked to predict each next location as accurately as possible. The number line ranged from 0 to 300 in units of 1, and the location on each trial was determined by the following process. On each trial, a number was randomly drawn from a Gaussian distribution, the mean of which changed at unsignaled moments—referred to as change points. The probability of a change point was .10 on each trial, except for the first three trials after the previous change point on which this probability was 0. When a change point occurred, a new mean for the number-generating distribution was randomly drawn from a uniform distribution ranging from 0 to 300 in units of 1. In each experimental session, participants completed two 200-trial blocks of this task. The *SD* of the number-generating distribution was 10 (low noise) in one block and 25 (high noise) in the other block (Figure 1A). We used six instantiations of this number-generating process—two for each noise level—such that participants experienced new sequences of outcomes in each session. The outcome sequences did not differ across participants.

Throughout the task, a horizontal number line, ranging from 1 to 300, was presented on the screen (Figure 1B). At the start of each trial, participants selected a specific location on the number line (their prediction) using a mouse, after which a small green oval was displayed underneath the selected location. One second later, an

Figure 1. The predictive inference task. (A) The number outcomes in a low-noise ($SD = 10$; top) and a high-noise ($SD = 25$; lower panel) block. The blue dots indicate the outcomes on each trial; these were drawn from a normal distribution of which the mean (horizontal black line) changed at unsigned moments (change points). The vertical lines indicate obvious (straight lines) and nonobvious (dotted lines) change points. (B) On each trial, participants predicted the next location on a number line, after which the outcome location was shown and participants could update their prediction for the next trial.



arrow accompanied by the number outcome on that trial was displayed in red above the corresponding location on the number line, and the difference between this location and the participant’s prediction was indicated by a gray bar. Half a second later, the next trial started, and participants could update their prediction. To ensure that learning rates were always in the 0–1 range, we constrained participants’ prediction space to the interval in between (and including) their previous prediction and the most recent outcome (Nassar et al., 2010, 2012).

We embedded the task in a cover story according to which the number line represented the earth and the outcomes reflected locations of missile attacks from outer space. We instructed participants that on each trial they could place a “laser shield” above a specific location on earth to prevent that location from being hit. To make the number-generating process as transparent as possible, we gave participants the following two additional instructions: (i) on their way to earth the missiles pass through an asteroid layer, causing random deflections of their direction and therefore variation in impact locations across trials, and (ii) the location at which the

missiles are aimed changes at unpredictable moments. These instructions provide intuitive information about the SD of the number-generating distribution (noise) and the occasional change points, respectively. Before starting the experimental blocks, participants completed two 30-trial practice blocks.

One participant fully updated her predictions to the most recent outcome on nearly all trials in all sessions, suggesting a misunderstanding of the task structure. We excluded this participant from further analyses.

Computation of Trial-specific Prediction Error and Learning Rate

We defined the prediction error on each trial t as the difference between the actual and predicted outcome location, that is, $\text{prediction error}(t) = \text{outcome}(t) - \text{prediction}(t)$. We defined the learning rate on each trial as the prediction update from that trial to the next, as a fraction of the most recent prediction error, that is, $\text{learning rate}(t) = [\text{Prediction}(t + 1) - \text{Prediction}(t)] / \text{prediction error}(t)$.

Normative Model

We also used an approximately Bayesian learning model that has been shown to capture key aspects of participants' performance on the predictive inference task (Nassar et al., 2010, 2012). The model updates beliefs about the current outcome-generating distribution according to a delta rule with a dynamic learning rate,

$$B_{t+1} = B_t + \alpha_t \times \delta_t \quad (1)$$

where B_t is the model's prediction about the generative distribution's mean, α_t is the learning rate, and δ_t is the prediction error (i.e., difference between observed and predicted outcome) on trial t .

The model uses the sequence of observed outcomes to compute two latent variables—change point probability (Ω) and relative uncertainty (τ)—on each trial. Together, these latent variables in turn determine trial-specific learning rate, according to

$$\alpha_t = \Omega_t + (1 - \Omega_t)\tau_t \quad (2)$$

where Ω_t reflects the posterior probability that a change point has occurred since the previous trial, which increases transiently following surprising outcomes. The model computes Ω_t following each new outcome (X_t) as a function of the likelihood of that outcome if a change point had occurred and the likelihood of that outcome if a change point had not occurred:

$$\Omega_t = \frac{U(X_t|0300)H}{U(X_t|0,300)H + \mathcal{N}(X_t|B_t, \sigma^2)(1-H)} \quad (3)$$

where U is the uniform distribution from which X_t is generated if a change point occurred, \mathcal{N} is the predictive normal distribution if a change point did not occur, B_t is the model's prediction on trial t , σ^2 is the total variance of the predictive distribution; and H is the hazard rate. The hazard rate is the proportion of trials on which a change point occurred (i.e., the prior probability of change points), which was .09 in our experiment.

The total uncertainty about the next outcome (σ^2 , i.e., the total variance of the predictive distribution) is the sum of the variance of the generative distribution (σ_N^2 , i.e., noise) and the uncertainty about the mean of the generative distribution (σ_μ^2). Although outcome uncertainty due to noise is constant within each task block, uncertainty attributable to imprecise knowledge of the mean of the generative distribution decreases with each observed outcome in a stable regime. Just as the gain in a Kalman filter, appropriate learning in our task depends on the proportion of total outcome uncertainty that is due to an imprecise estimate of the generative mean, and we define relative uncertainty (τ) as this proportion. On each trial, relative uncertainty is computed according to the variance on the predictive distribution over generative means (a weighted mixture of change point and

nonchange point conditional distributions) according to the following equation:

$$\tau_{t+1} = \frac{\Omega_t \sigma_N^2 + (1 - \Omega_t)\tau_t \sigma_N^2 + \Omega_t(1 - \Omega_t)(\delta_t(1 - \tau_t))^2}{\Omega_t \sigma_N^2 + (1 - \Omega_t)\tau_t \sigma_N^2 + \Omega_t(1 - \Omega_t)(\delta_t(1 - \tau_t))^2 + \sigma_N^2} \quad (4)$$

where the numerator reflects the variance on the predictive distribution over possible generative means (σ_μ^2) and the denominator is the total outcome variance, which also includes the noise variance (McGuire et al., 2014; Nassar et al., 2012). Relative uncertainty is computed in anticipation of each upcoming outcome and therefore reflects outcome-independent adjustments in learning (Nassar et al., 2010).

We applied the model to each participant's observed sequence of outcomes while fixing hazard rate (H) to the actual proportion of change point trials (.09) to obtain per-trial estimates of change point probability and relative uncertainty. Hazard rate can also be treated as a free parameter that is estimated by fitting the model to each participant's prediction data, thereby capturing interindividual variability in learning rate due to different prior expectations about the frequency of change points. To examine potential drug effects on the hazard rate parameter, we fitted the model to each participant's predictions in each session by minimizing the total squared difference between the participant's and the model's predictions, using a constrained search algorithm (*fmincon* in MATLAB).

In addition, we also used a recent extension of the normative model (Nassar et al., 2016) to obtain subjective estimates of change point probability and relative uncertainty. This involves running the normative model across the prediction errors experienced by participants, rather than the observed outcomes. To this end, model variables were computed recursively by first determining the uncertainty about the current mean of the outcome-generating distribution according to the relative uncertainty, change point probability and prediction error from the previous trial:

$$\sigma_\mu^2 = \Omega_t \sigma_N^2 + (1 - \Omega_t)\tau_t \sigma_N^2 + \Omega_t(1 - \Omega_t)\delta_t(1 - \tau_t) \quad (5)$$

Relative uncertainty was computed by expressing uncertainty about the generative mean as a fraction of total uncertainty about the next outcome location,

$$\tau_{t+1} = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_N^2} \quad (6)$$

This relative uncertainty estimate, along with the variance on the outcome distribution (noise; σ_N^2) was used to compute the change point probability associated with each new prediction error:

$$\Omega_{t+1} = \frac{H/300}{H/300 + \mathcal{N}\left(\delta_{t+1} \middle| 0, \frac{\sigma_N^2}{1 - \tau_{t+1}}\right)(1 - H)} \quad (7)$$

where H is the actual proportion of change point trials (.09) and δ_{t+1} is the new prediction error. Subjective estimates of change point probability and relative uncertainty were computed by evaluating Equations 5 and 6 according to the trial-by-trial prediction errors made by each individual participant.

EEG Recording and Analyses

We recorded EEG from 64 Ag/AgCl scalp electrodes and from the left and right mastoids. We measured the horizontal and vertical EOG using bipolar recordings from electrodes placed approximately 1 cm lateral of the outer canthi of the two eyes and from electrodes placed approximately 1 cm above and below the participant's right eye. The EEG signal was preamplified at the electrode to improve the signal-to-noise ratio and amplified with a gain of $16\times$ by a BioSemi ActiveTwo system. The data were digitized at 24-bit resolution with a sampling rate of 512 Hz using a low-pass fifth-order sinc filter with a half-power cutoff of 102.4 Hz. Each active electrode was measured online with respect to a common mode sense active electrode producing a monopolar (nondifferential) channel.

EEG data were processed using a combination of BrainVision Analyzer 2 (Brain Products) and MATLAB (The Mathworks), the latter via custom scripting and subroutines from the EEGLAB toolbox (Delorme & Makeig, 2004). Continuous data were first rereferenced to the average of the left and right mastoid channels and high-pass filtered to 0.1 Hz (12 dB/octave). Ocular artifacts were removed using a regression-based algorithm (Gratton, Coles, & Donchin, 1983), after which the data were low-pass filtered up to 30 Hz (12 dB/octave). Noisy channels were then identified by visual inspection of signal variance and interpolated via spherical spline interpolation. Data epochs were extracted from 250 msec before to 1000 msec after outcome onset on each trial and baseline-corrected to the 250-msec interval preceding outcome onset. All epochs were then inspected for violations of amplitude (any sample from any scalp channel with an absolute voltage $> 120 \mu\text{V}$) and gradient (any scalp channel where absolute slope of a fitted line to the data was $> 65 \mu\text{V}/\text{sec}$) artifact criteria. In cases where no more than one channel was identified as artifactual, this channel was interpolated and the associated epoch was retained for subsequent analysis; otherwise, that epoch was discarded. A mean of $9.2 \pm 12.8\%$ of epochs for the placebo sessions, $11.1 \pm 12.0\%$ of epochs for the clonidine sessions, and $14.9 \pm 16.9\%$ of epochs for the scopolamine sessions were discarded. For all analyses, our measurement of the outcome-locked P3 component was based on the mean signal across a cluster of five centroparietal electrodes that was centered on the region of maximum component amplitude in the grand-averaged topography (corresponding to the location of Pz according the standard 10/20 measurement system). For single-

trial analysis of the P3, waveforms were low-pass filtered to 6 Hz to enhance signal-to-noise ratio and P3 amplitude was measured as the mean voltage between 340 and 480 msec postoutcome.

To identify significant effects of drug treatment on the trial-averaged centroparietal ERP and correct for multiple comparisons, we computed nonparametric permutation tests based on temporal clustering (Maris & Oostenveld, 2007). The following steps were followed for each pairwise comparison across the different treatment levels (placebo vs. clonidine, placebo vs. scopolamine, clonidine vs. scopolamine). First, a paired t test was performed to identify individual time points at which the effect of treatment was significant ($p < .05$) without correction for multiple comparisons, and such time points were combined into clusters based on their temporal adjacency. Next, the t scores of all time points comprising a cluster were summed, which yielded a cluster-level statistic (i.e., a cluster t value) for each identified cluster. Subsequently, a null distribution of cluster-level statistics was computed via the following procedure: For each of 1000 iterations, the treatment labels were randomly reassigned within participants, the above-described temporal clustering procedure was executed on the permuted data, and the maximum absolute cluster-level statistic derived from this procedure was stored. Finally, the absolute cluster-level statistics of all empirical clusters were compared with the distribution of values obtained from the permutation procedure. All time points comprising a cluster with a statistic larger than 95% of the permutation distribution were considered significant, corresponding to a cluster-corrected alpha level of .05.

Statistical Analyses

We conducted multilevel regression and mediation analyses on single-trial measures of learning rate and P3 amplitude, using the Multilevel Mediation toolbox (wagerlab.colorado.edu/tools; Atlas, Bolger, Lindquist, & Wager, 2010; Wager, van Ast, et al., 2009; Wager, Waugh, et al., 2009). As the occasional change points in the outcome-generating process produced exceptionally large prediction errors, the data contained some strong "outlier" trials. To deal with this and to account for potential nonlinearity of the relationships of interest, we replaced the prediction error, learning rate, and P3 variables with their ranks in all regression and mediation analyses. We used bootstrapping (100,000 bootstrap samples) for significance testing, which does not require the assumption of normality for valid inference.

Regression analyses. We tested for effects of (absolute) prediction error, SD of the generative distribution, treatment (two dummy variables coding for clonidine and scopolamine), and their interactions on learning rate and P3 amplitude. Trials with prediction errors of 0 were excluded from the analysis on learning rate (6.0% of all

trials), as participants could not update their prediction on those trials (see task description above). In the analysis on P3 amplitude, we included a binary regressor that indicated whether or not the prediction error was exactly zero.

Mediation analyses. We further examined the relationships between trial-to-trial variation in prediction error, P3 amplitude, and learning rate using multilevel mediation. Mediation analyses test whether the relationship between an independent variable (X) and a dependent variable (Y) can be explained by a third variable (M ; the mediator).

In our first mediation model, we used prediction error as the X variable, learning rate as the Y variable, and P3 amplitude as the M variable. Thus, this model tested whether (i) there was an effect of prediction error on P3 amplitude (path a); (ii) P3 amplitude was predictive of learning rate, when controlling for prediction error (path b); and (iii) the relationship between prediction error and learning rate was formally mediated by P3 amplitude, that is, whether the relationship between prediction error and learning rate (path c) decreased when controlling for P3 amplitude ($c - c'$, equivalent to $a * b$). In additional mediation models, we used the computational variables change point probability and relative uncertainty (obtained by running the normative model either over the observed outcomes or over the experienced prediction errors; see Normative Model section above), in separate analyses, as the X variable.

We used the ranks of the X , M , and Y variables in all mediation models. Trials with prediction errors of 0 were excluded from the mediation analyses. We included treatment and the SD of the generative distribution as co-

variates in all mediation models and tested the significance of all effects using a bootstrap procedure (100,000 bootstrap samples).

Results

Physiological and Alertness Data

As expected, clonidine lowered systolic (mean tension 95 mm Hg) and diastolic (61 mm Hg) blood pressure relative to placebo (mean tension 110/76 mm Hg), also during performance of the predictive inference task ($t = 120\text{--}150$ min), $F(2, 34) = 34.6, p < .0005$ and $F(2, 34) = 19.8, p < .0005$ for systolic and diastolic pressure, respectively (Figure 2A). There were no differences between the placebo and scopolamine sessions in systolic or diastolic blood pressure. Scopolamine (65/min), as expected, lowered heart frequency relative to placebo (71/min) and clonidine (71/min), also during performance of the predictive inference task, $F(2, 34) = 5.1, p = .01$ (Figure 2B). The differences in blood pressure and heart rate were significant between clonidine and scopolamine (all $ps < .02$).

Results from the SRT task, administered before drug administration (at arrival of the participant), right before, and right after performance of the predictive inference task, suggest that clonidine increased SRT (336 msec) relative to placebo (279 msec) and scopolamine (317 msec), $F(2, 34) = 10.1, p < .0005$. Furthermore, mean SRT increased as the test session progressed, $F(2, 34) = 19.8, p < .0005$. As depicted in Figure 2C, clonidine increased SRT more strongly as the test session progressed than scopolamine and placebo, $F(4, 68) = 3.7, p = .009$. Pairwise comparisons for pretest and posttest indicated

Figure 2. Cardiovascular and alertness results. (A) Blood pressure data for the three treatments. The shaded gray area indicates significant pairwise comparisons between clonidine and placebo ($p < .05$, uncorrected). (B) Heart frequency for the three treatments. The shaded gray area indicates significant pairwise comparisons between scopolamine and placebo ($p < .05$). (C) Results from a SRT task, administered at the start of the test session (predrug administration) and right before (pretest) and after (posttest) participants performed the predictive inference task.

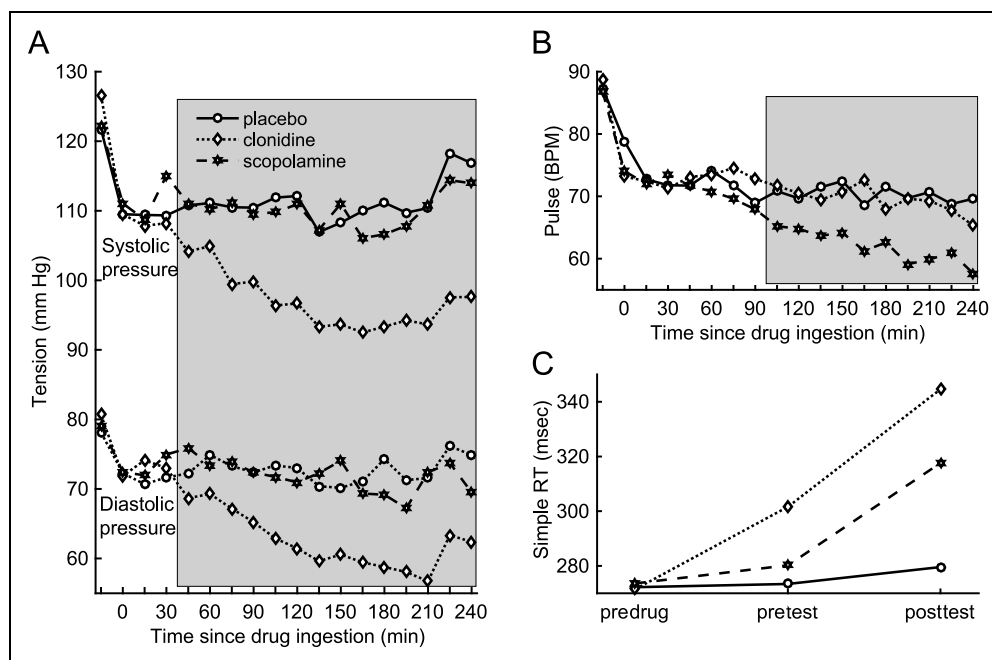
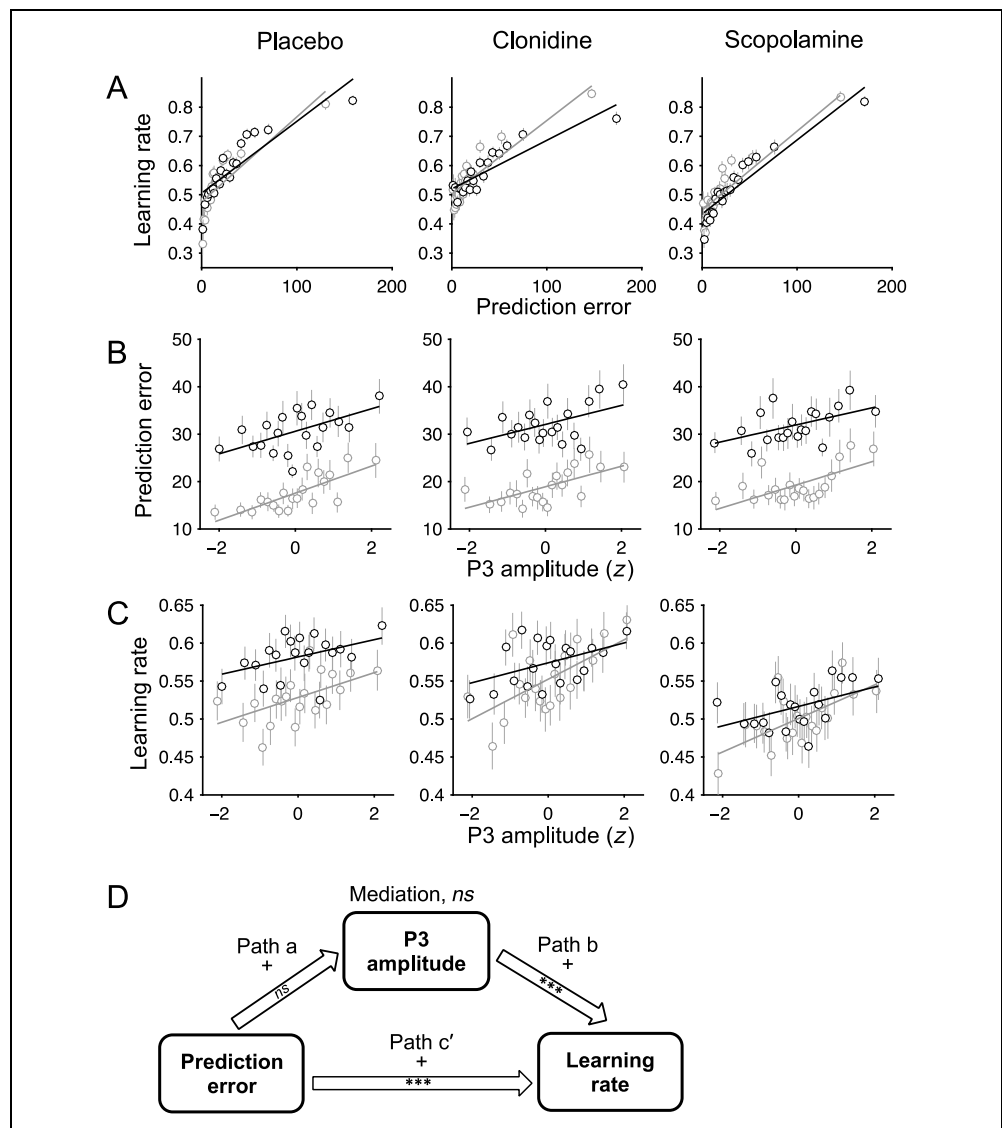


Figure 3. Relationships between prediction error, P3 amplitude, and learning rate. Note that we show the raw data here, although our statistical analyses were rank-based. (A) Learning rate as a function of prediction error, *SD* of the generative distribution and treatment. Learning rates are sorted into 20 bins according to single-trial prediction error, pooled across participants. The lines show the linear fit to the unbinned single-trial data. (B) Prediction error as a function of P3 amplitude, *SD* of the generative distribution and treatment. Prediction errors are sorted into 20 bins according to single-trial P3 amplitude, pooled across participants. Single-trial P3 amplitudes were z-scored separately for each participant and session, before sorting. Although we tested for effects of prediction error on P3 amplitude in our statistical analyses, we here plot P3 amplitude on the *x* axis to allow for a comparison of its relationships with prediction error and learning rate (B and C). (C) Learning rate as a function of P3 amplitude, *SD* of the generative distribution and treatment. Plotting procedure was the same as above. (D) Mediation model and results. *** $p < .001$, *ns* = not significant.



that clonidine and scopolamine reliably differed from placebo during both the pretest and the posttest (all $ps < .04$).

Behavioral and Electrophysiological Results

As in our previous study (Jepma et al., 2016), learning rate increased with increasing prediction error ($t(16) = 5.9$, bootstrap $p < .001$; Figure 3A). In addition, the effect of prediction error on learning rate was stronger in the low-noise than the high-noise block (Prediction error \times Noise interaction, $t(16) = 1.8$, bootstrap $p = .03$). The significance of these effects in a rank-based regression analysis suggests that they were not driven by a small number of outlier trials.

There were no significant main effects of Prediction error or noise and no Prediction error \times Noise interaction effect on P3 amplitude (all bootstrap $ps > .5$; Figure 3B). However, P3 amplitude was increased on trials when

prediction error was exactly 0 ($t(16) = 1.9$, bootstrap $p = .02$), possibly reflecting the rewarding nature and/or atypical consequence (i.e., no possibility of updating the next prediction) of perfectly predicted outcomes. Note that we reported a significant positive effect of prediction error on P3 amplitude in our previous study (Jepma et al., 2016). This apparent inconsistency between our two studies can be explained by the fact that we used standard linear regression in our previous study but rank-based regression in the current study. Indeed, standard linear regression on the data from the current study did reveal a significant effect of prediction error on P3 amplitude. That this effect disappeared when using a rank-based analysis suggests that it was driven by a subset of trials with very large prediction errors, whereas there is no clear relationship between prediction error and P3 amplitude across the lower range of prediction errors.

Below, we report three sets of analyses and results. First, we report the within-subject relationships between

trial-to-trial fluctuations in prediction error magnitude, P3 amplitude, and learning rate. Second, we report how P3 amplitude and learning rate relate to trial-to-trial fluctuations in two latent variables, change point probability and relative uncertainty, that drive learning rate according to a previously established normative model (Nassar et al., 2010, 2012). The aim of these two analyses was to examine if we could replicate our recent findings that P3 amplitude predicts learning rate and mediates the effect of prediction error magnitude, change point probability, and relative uncertainty on learning rate (Jepma et al., 2016). Third, we report the effects of our clonidine and scopolamine manipulations on learning rate and P3 amplitude.

P3 Amplitude Predicts Learning Rate

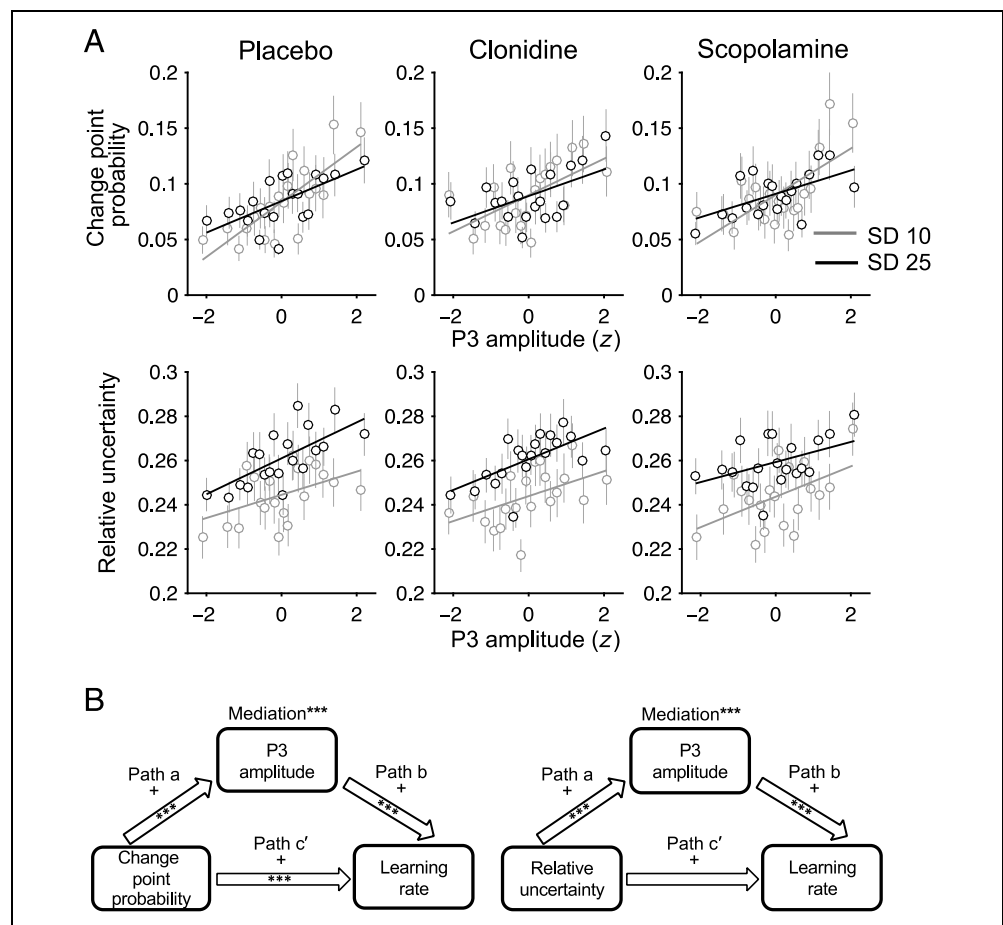
We used multilevel rank-based mediation to assess whether P3 amplitude mediated the effect of prediction error on learning rate (Figure 3D). In this mediation model, there was no significant effect of prediction error on P3 amplitude (path *a*, $p = .07$), consistent with the results above. Importantly, however, larger P3 amplitudes predicted higher learning rates when controlling for prediction error (path *b*, $p < .001$; see Figure 3C for the relationship between P3 amplitude and learning rate, not

controlled for prediction error). P3 amplitude did not formally mediate the effect of prediction error on learning rate ($a * b, p = .32$), which was to be expected given the absence of a significant path *a*.

P3 Amplitude Mediates the Effects of Surprise and Belief Uncertainty on Learning Rate

In two additional rank-based mediation analyses, we examined the relationships between trial-by-trial fluctuations in two latent variables that drive learning rate according to a normative model, P3 amplitude, and learning rate. The normative model assumes that participants use the observed sequence of outcomes to compute two latent variables on each trial—change point probability and relative uncertainty—which together determine learning rate. Change point probability approximates the posterior probability that a change point has occurred on the most recent trial, given all previous outcomes; hence, it reflects the unexpectedness of the most recent outcome. Relative uncertainty reflects the uncertainty about the mean of the outcome distribution *before* a new outcome is observed, which depends inversely on the number of prior observations attributable to the current environmental state.

Figure 4. Relationships between change point probability and relative uncertainty, P3 amplitude and learning rate. (A) Both change point probability and relative uncertainty predict P3 amplitude. Change point probability and relative uncertainty estimates are sorted into 20 bins according to z-scored single-trial P3 amplitude, pooled across participants. The lines show the linear fit to the unbinned single-trial data. (B) Mediation models and results. $***p < .001$.



In line with our previous findings (Jepma et al., 2016), both change point probability and relative uncertainty were positive predictors of P3 amplitude (path a , both $ps < .001$; Figure 4A), corroborating the idea that P3 amplitude reflects both the unexpectedness of an outcome and the preexisting belief uncertainty. Furthermore, P3 amplitude formally mediated the effect of both change point probability and relative uncertainty on learning rate (path $a * b$, both $ps < .001$; Figure 4B). That change point probability predicted P3 amplitude, whereas prediction error did not, may seem counterintuitive given the strong relationship between prediction error and change point probability. However, change point probability also depends on the current level of uncertainty about the next outcome, which is determined by the amount of random outcome variability (noise) and the observer's ignorance about the mean of the outcome distribution (estimation uncertainty). Specifically, the same absolute prediction error signals a larger change point probability when the noise level and/or estimation uncertainty are lower. Therefore, our finding that P3 amplitude tracks change point probability, but not prediction error in itself, suggests that P3 amplitude is sensitive to the context in which prediction errors occur, that is, to how unexpected an outcome is given the current level of uncertainty.

In the previous analysis, we obtained trial-by-trial estimates of change point probability and relative uncertainty by running the normative model over the observed sequence of outcomes. A potential problem with this approach is that the predictions made—and thus the prediction errors experienced—by the model may not always perfectly match those of the participant. If this is the case, participants' subjective change point probability and relative uncertainty values will differ from those computed by the model. To address this issue, we obtained subjective measures of change point probability and uncertainty, using a method recently developed by Nassar et al. (2016; see Methods), and repeated the above-described mediation analyses using the ranks of these subjective values as predictor variables. These analyses showed that subjective estimates of both change point probability and relative uncertainty were related to P3 amplitude as well (path a , p values are .03 and $< .001$, respectively). In addition, P3 amplitude mediated the effects of both subjective change point probability and subjective relative uncertainty on learning rate (p values are .02 and .009, respectively).

Baseline-dependent Effects of Clonidine and Scopolamine on Learning Rate following Nonobvious Change Points

Both clonidine and scopolamine reduced P3 amplitude ($t(16) = 2.6, p = .02$ and $t(16) = 3.6, p = .003$, respectively; Figure 5). However, neither clonidine nor scopolamine affected learning rate (both $ps > .15$), and there




Figure 5. Grand-averaged outcome-evoked P3 (mean signal across a cluster of five centroparietal electrodes centered on location Pz) in the placebo, clonidine, and scopolamine session. Shaded areas indicate SEM. The blue and red horizontal lines indicate the time points at which signal amplitude in the clonidine and scopolamine session, respectively, was significantly different compared with the placebo session ($p < .05$, cluster-corrected). There were no reliable amplitude differences between the clonidine and scopolamine traces after cluster correction.

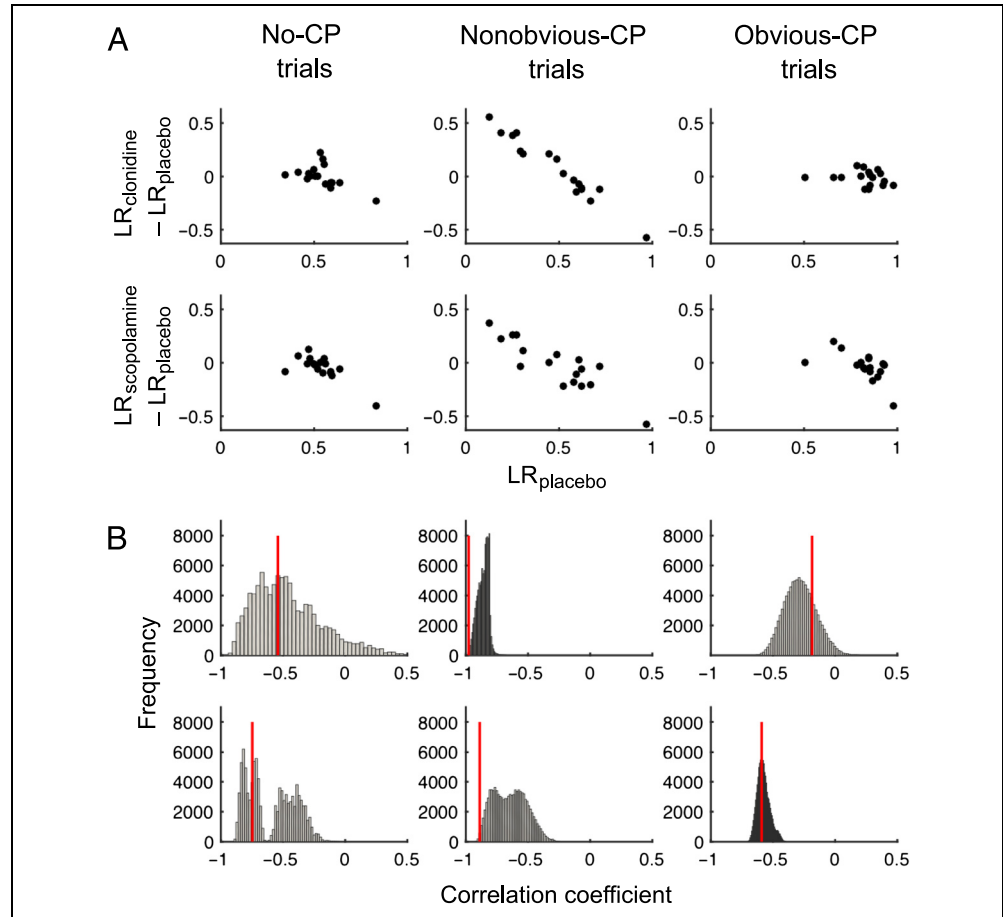
were no Treatment \times Prediction error or Treatment \times Noise interactions on learning rate or P3 amplitude (all $ps > .11$). The effects of the model-based variables change point probability and relative uncertainty on learning rate and P3 amplitude did not interact with either of the treatments (all $ps > .12$). Finally, the hazard rate parameter in the clonidine and scopolamine sessions, obtained by fitting the normative model to each participant's predictions, did not differ from the hazard rate in the placebo session (mean estimated hazard rate = 0.40 ($SD = 0.20$), 0.36 ($SD = 0.17$), and 0.38 ($SD = 0.20$), respectively; $ps > .6$). As in previous studies (Jepma et al., 2016; Nassar et al., 2010), the model-estimated hazard rates were higher than the actual hazard rate (0.09 in this study; see Methods), suggesting that human observers on this task consistently overestimate the frequency at which change points occur.

Thus, both drugs had an overall suppressive effect on the P3 amplitude but did not affect task performance at the group mean level. Importantly, previous studies have shown that the effects of catecholaminergic drugs depend on individuals' arousal state or baseline level of catecholamine activity (Gibbs, Bautista, Mowlem, Naudts, & Duka, 2014; de Rover et al., 2012; Cools et al., 2009; Luksys, Gerstner, & Sandi, 2009; Cools & Robbins, 2004; Coull, 2001). Consistent with this, we previously found that the NE transporter (NET) blocker atomoxetine increased learning rates in participants who normally (in a placebo session) used low learning rates but decreased learning rates in participants who normally used high learning rates (Jepma et al., 2016). This baseline-dependent atomoxetine effect was stronger than predicted by regression to the mean for change point trials, but not for trials on which no change point occurred. To test for similar baseline-dependent drug effects in

this study, we computed the across-subject correlation between mean learning rate in the placebo session (LR_{placebo}) and the effect of each drug on learning rate ($LR_{\text{clonidine}} - LR_{\text{placebo}}$ and $LR_{\text{scopolamine}} - LR_{\text{placebo}}$, in separate analyses). As in our previous study, we separately computed these correlations for the trials on which no change point occurred, the trials on which an obvious change point occurred (change point outcome > 2 SDs from previous mean), and the trials on which a less obvious change point occurred (change point outcome < 2 SDs from previous mean). These analyses revealed that individual differences in learning rate in the placebo session negatively correlated with the clonidine and scopolamine effects on learning rate (Figure 6A). Thus, both drugs were associated with increased learning rates in participants with low baseline (i.e., placebo) learning rates and with decreased learning rates in participants with high baseline learning rates. For both drugs, these negative correlations were strongest for the nonobvious change point trials. Importantly, however, regression to the mean likely contributed to these negative correlations. We therefore performed two additional analyses to test for baseline-dependent drug effects above and beyond regression to the mean.

First, we compared our observed correlation coefficients against the distribution of correlation coefficients predicted exclusively by regression to the mean, using permutation testing. Specifically, we computed the above-described across-subject correlation 100,000 times, each time randomly assigning “placebo” versus “drug” labels to each participant’s placebo and drug session (separately for the placebo vs. clonidine and placebo vs. scopolamine sessions). This analysis revealed that the baseline-dependent effects of both clonidine and scopolamine on learning rate were stronger than predicted by regression to the mean for the nonobvious change point trials (proportion of permutation distribution below the observed correlation = .002 and .005 for clonidine and scopolamine, respectively). However, for the trials on which no change point occurred and on which an obvious change point occurred, the observed negative correlations did not differ from those predicted by regression to the mean (proportion of permutation distribution below the observed correlation $> .3$ for all comparisons). To further examine the effect of change point probability on baseline-dependent drug effects, we repeated the permutation analysis for different levels of subjective change point probability (Figure A1). The

Figure 6. Baseline-dependent drug effects on learning rate. (A) Individual differences in the clonidine and scopolamine effect on learning rate as a function of participants’ learning rate in the placebo session, separately for the trials with no change point, nonobvious change points, and obvious change points. Note that the apparent correlations for the no-CP and obvious-CP trials and part of the correlation for the nonobvious-CP trials are driven by regression to the mean (see B). (B) Permutation distributions of the correlation coefficients predicted by regression to the mean (the plots corresponds to those in A). The red lines indicate our observed correlations.



results of this additional analysis suggest that clonidine had baseline-dependent effects on learning rate above and beyond regression to the mean for change point probabilities between .20 and .40. For scopolamine, this analysis only revealed (marginal) evidence for baseline-dependent effects for change point probabilities below .025. It must be noted, however, that change point probability was far from equally distributed, as the vast majority of trials had a change point probability below .10. There were especially few trials with change point probabilities between .40 and .80, which may have prevented the detection of baseline-dependent effects in this range. Thus, the results of this additional analysis must be considered with caution.

Second, baseline-dependent drug effects on learning rate would produce lower across-subject variance in learning rate in the drug sessions than in the placebo session, but regression to the mean would not (Kelly & Price, 2005). We tested this prediction using Pitman's test of equality of variance in paired samples (Pitman, 1939). For the nonobvious change point trials, the across-subject variance in learning rate was indeed lower in the clonidine (0.007) and scopolamine (0.012) sessions than in the placebo session (0.049; $t(15) = 5.9, p < .001$ and $t(15) = 3.0, p = .008$, respectively). For the trials on which no change point occurred and for the obvious change point trials, the variance in learning rate did not differ between sessions (all $ps > .2$). Together, the results from these two analyses strongly suggest that both clonidine and scopolamine had baseline-dependent effects on learning rate, but selectively following nonobvious change points.

STUDY 2

Having established that the effects of clonidine and scopolamine—like those of atomoxetine in our previous study—on learning rate depend on individual differences in baseline learning rate, we sought to explore the genetic basis of these individual differences in learning rate. We hypothesized that such differences in learning rate might be related to interindividual variation in genes controlling noradrenergic function. Although previous research has examined in detail how dopaminergic genes modulate the efficacy of particular neural computations (Frank & Fossella, 2011), very little is known about the potential role of noradrenergic genes in specific neural computations. The aim of Study 2 was to conduct a first exploration of this important question by examining how learning rate may be affected by interindividual variation in genes coding for the synthesis of NE (*DBH*), for metabolizing NE (*COMT*), for removing NE from the synaptic cleft (*NET*), and for adrenergic receptor sensitivity (*ADRA2A*, *ADRB1*). To this end, we administered the predictive interference task to a group of 151 young adults who were genotyped for nine single-nucleotide polymorphisms (SNPs) in genes assumed to affect noradrenergic

function. We envisaged that the resulting data set would reveal promising candidate genes that may be the focus of future confirmatory studies.

Methods

Participants

Saliva samples were collected from 151 healthy, highly educated young adults, who completed the predictive interference task described above. Five participants were excluded from analyses because they fully updated their predictions to the most recent outcome on nearly all trials, suggesting a misunderstanding of the task structure. The data from nine other participants were discarded because no high-quality DNA could be extracted from the saliva sample. The remaining sample hence consisted of 137 participants (mean age = 21.5 years, range = 18–28 years; 110 women). The study was approved by the medical ethics committees of the Vrije Universiteit Amsterdam and Leiden University, and written informed consent was obtained from all participants.

Genotyping and Statistical Analyses

A saliva sample was collected from each participant using the Saliva DNA Collection, Preservation and Isolation Kit (Norgen Biotek Corporation). DNA was extracted from saliva using the Oragene kits (DNA Genotek, Inc.). In total, 137 participants were genotyped using Sanger Sequence technology for three variants in the NE transporter gene (*NET*: rs5569, rs2242446, rs28386840), three variants in the dopamine beta-hydroxylase gene (*DBH*: rs1108580, rs1611115, *DBH5'* ins/del), and variants in the α 2A receptor gene (*ADRA2A*: rs1800544), beta-1 receptor gene (*ADRB1*: rs1801253), and *COMT* gene (rs165599). Because of missing genotypes, six samples had to be discarded from the analyses of the *ADRB1* (rs1801253) SNP and *NET* (rs28386840) SNP. Three samples had to be discarded from the analysis of the *ADRA2A* (rs1800544) SNP, two from the analysis of the *DBH5'* ins/del SNP, and one from the analysis of the *NET* (rs5569) SNP. For all SNPs, the genotypes were in Hardy–Weinberg equilibrium ($ps > .05$), with the exception of the *COMT* SNP ($p = .0059$), which was therefore discarded from further analysis. The genotype frequencies for the remaining eight SNPs are reported in Table 1.

The effect of each SNP on learning rate was examined with a general linear model analysis that included genotype (two levels, see below) as a categorical between-subject variable, trial type as a repeated-measures variable (no change point, nonobvious change point, obvious change point; see Study 1), and age as nuisance covariate. In each analysis, individuals homozygous for the ancestral allele were contrasted with the combined group of heterozygotes and individuals homozygous for the derived allele.

Table 1. SNPs, Genotype Frequencies, and Genotype Effects on Learning Rate

| Gene | NET | | NET | | NET | | DBH | | DBH | | DBH | | ADRA2A | | ADRB1 | |
|----------|--------|-------|-----------|-------|------------|-------|-----------|-------|-----------|-------|---------------|--------------|-----------|-------|-----------|-------|
| | rs5569 | | rs2242446 | | rs28386840 | | rs1108580 | | rs1611115 | | DBH*5 INS/DEL | | rs1800544 | | rs1801253 | |
| Allele | GG | GA/AA | TT | CT/CC | AA | TA/TT | GG | AG/AA | CC | TC/TT | ins/ins | del carriers | CC | GC/GG | CC | GC/GG |
| <i>N</i> | 62 | 60/14 | 69 | 56/12 | 60 | 51/20 | 35 | 77/25 | 86 | 45/6 | 43 | 58/34 | 66 | 55/13 | 78 | 46/7 |
| Mean LR | 0.58 | 0.56 | 0.56 | 0.59 | 0.56 | 0.60 | 0.56 | 0.58 | 0.58 | 0.57 | 0.58 | 0.58 | 0.60 | 0.56 | 0.57 | 0.59 |
| <i>F</i> | 0.1 | | 3.0 | | 4.6 | | 0.9 | | 0.0 | | 0.0 | | 6.7 | | 0.4 | |
| <i>p</i> | .73 | | .08 | | .03* | | .36 | | .84 | | .96 | | .01* | | .55 | |

Mean LR = mean learning rate. Bonferroni-corrected $\alpha = .00625$.

* $p < .05$ (uncorrected).

Results

Table 1 reports the mean learning rates and the statistical main effect of genotype on learning rate for each of the eight SNPs that were included in the analyses. For most of the SNPs, there was no association with learning rate. However, the *ADRA2* ($p = .011$) and one of the *NET* SNPs (rs28386840, $p = .03$) showed suggestive evidence for an association with learning rate. As the associations did not survive the Bonferroni correction for multiple tests (corrected $\alpha = .00625$), the findings should be regarded as hypothesis-generating and need confirmation in a follow-up study. Mean learning rates were 0.496 ($SD = 0.127$) for trials on which no change point occurred, 0.480 ($SD = 0.165$) for nonobvious change point trials, and 0.759 ($SD = 0.128$) for obvious change point trials. There were no interactions between genotype and trial type.

DISCUSSION

In the present research, we report three main findings. First, we replicated our recent findings that trial-to-trial variability in centroparietal P3 amplitude—a putative index of phasic catecholamine release in the cortex (Polich, 2007; Nieuwenhuis et al., 2005; Pineda et al., 1989)—predicts learning rate and mediates the effects of outcome-evoked surprise and belief uncertainty on learning rate (Jepma et al., 2016). Replication of these effects in a new group of participants reinforces the idea that these are true and robust effects. Second, pharmacological suppression of either NE or ACh activity produced baseline-dependent effects on learning rate following nonobvious task changes, but not following obvious task changes or during periods of stability. Third, we identified association of two SNPs, located in genes coding for the NE transporter and the $\alpha 2A$ receptor, with learning rate, thus providing first evidence for a genetic underpinning of this computational variable. At the moment, the genetic findings should be regarded as hypothesis-generating and provide promising candidates for future studies in-

vestigating the genetic basis of individual differences in belief updating under uncertainty.

Our P3 results contribute to the discussion about the functional role of this event-related brain potential. That single-trial P3 amplitude was sensitive to outcome-evoked surprise and predicted learning rate supports the prevalent context-updating theory of P3 function, which assumes a role for the centroparietal P3 process in the updating of one's internal model of a task context (Donchin & Coles, 1988). Interestingly, two recent studies found that model-based estimates of belief updating correlated with P3 amplitude measured over frontocentral (P3a), but not centroparietal (P3b), scalp regions (Bennett, Murawski, & Bode, 2015; Kolossa et al., 2015). The reason for the more anterior scalp distributions in these studies remains to be examined; these may be related to the different paradigms used in these studies and/or to the fact that these studies focused on Bayesian, model-based indices of belief updating while we used a direct behavioral measure of learning rate. Centroparietal P3 amplitude also reflected the uncertainty about the outcome-generating process (relative uncertainty, which is strongly related to estimation uncertainty). In contrast, P3 amplitude did not reflect the noisiness of the generative process, which varied between task blocks. Estimation uncertainty and noise are both considered forms of expected uncertainty, which make it harder to detect environmental change. However, these two types of expected uncertainty differently affect learning rate and are represented in partially different brain regions (Kobayashi & Hsu, 2017; Payzan-LeNestour et al., 2013). Importantly, whereas uncertainty due to noise is irreducible, estimation uncertainty decreases with each additional observation during a period of stability. In terms of precision, noise and estimation uncertainty can be considered as the inverse precisions of observations and beliefs, respectively (cf. Kwisthout, Bekkering, & van Rooij, 2017): Noise determines how precisely outcomes can be predicted under a given generative distribution (corresponding to the standard deviation of the mean), and estimation uncertainty reflects

the imprecision of the observer's internal model of, or belief about, the generative distribution (corresponding to the standard error of the mean). Our finding that P3 amplitude is not sensitive to noise needs to be verified using tasks in which noise levels vary within task blocks.

Given the evidence for a link between the centroparietal P3 and the phasic LC–NE response (De Taeye et al., 2014; Nieuwenhuis, 2011; Polich, 2007; Nieuwenhuis et al., 2005; Pineda et al., 1989), our P3 results provide indirect support for a role of the catecholamine systems in the adjustment of learning rate based on ongoing estimates of change and uncertainty. Importantly, however, the P3 results cannot dissociate the specific contributions of NE and DA activity because (i) LC activity results in the corelease of DA from noradrenergic terminals (Devoto & Flore, 2006) and (ii) there are bidirectional projections between the dopaminergic nucleus ventral tegmental area and the LC (Sara, 2009). Indeed, dopaminergic drugs have been shown to affect the centroparietal P3 to unexpected and novel stimuli (Rangel-Gomez, Hickey, van Amelsvoort, Bet, & Meeter, 2013; Kahkonen et al., 2002; Hansenne, 2000). In addition, P3 amplitude can also be modulated by pharmacological manipulations of the cholinergic system (Brown et al., 2016; Brown, van der Wee, et al., 2015; this study), presumably reflecting mutual interactions between the basal forebrain and the LC (Acquas, Wilson, & Fibiger, 1998; Adams & Foote, 1988; Egan & North, 1985), as will be discussed below.

Both clonidine and scopolamine were observed to have baseline-dependent (i.e., normalizing) effects on learning rate, but specifically following nonobvious change points. Baseline-dependent drug effects are consistent with evidence that the relationship between neuromodulatory activity—in particular catecholamine activity—and neurocognitive function is not monotonic but follows an inverted U-shape (Gibbs et al., 2014; de Rover et al., 2012; Cools et al., 2009; Luksys et al., 2009; Aston-Jones & Cohen, 2005; Cools & Robbins, 2004; Coull, 2001). But why would these baseline-dependent drug effects be specific to the nonobvious change points? Nonobvious change point trials are likely associated with high uncertainty about whether or not a change occurred; hence, higher-level processes, such as the explicit attribution of prediction errors to change versus noise, may not provide sufficient guidance on whether or not beliefs should be updated on these trials. Consequently, stimulus-driven effects on learning rate, mediated by catecholamine-induced increases in neural gain (Aston-Jones & Cohen, 2005; Servan-Schreiber, Printz, & Cohen, 1990), may have a relatively large impact on belief updating following nonobvious change points, which could explain the specificity of the drug effects to these trials. The different effects we found for obvious and nonobvious change points stress the importance of dissociating between clear and ambiguous task changes when studying the role of neuromodulatory systems in adaptive behavior. A remaining question is why, in our previous

study, atomoxetine had baseline-dependent effects on learning rate following both obvious and nonobvious change points (Jepma et al., 2016). A possible reason for this more general effect of atomoxetine is that atomoxetine-induced increase of catecholamine levels results in activation of all types of catecholamine receptors, whereas the effects of clonidine are specific to the α_2 receptor. In addition, the mechanisms through which drug-induced decreases and increases in baseline NE activity—the presumed effects of clonidine and atomoxetine, respectively—can both produce baseline-dependent effects remain to be explored.

Thus, our clonidine findings may suggest that the α_2 receptor plays an important role in belief updating following ambiguous changes. Importantly, antagonism of the muscarinic ACh receptor by scopolamine, however, produced a highly similar baseline-dependent effect on learning rate following nonobvious change points. Furthermore, clonidine and scopolamine had a similar suppressive effect on P3 amplitude during the predictive inference task and similarly affected performance and ERPs during tasks measuring temporal attention and phasic alertness (Brown et al., 2016; Brown, Tona, et al., 2015; Brown, van der Wee, et al., 2015). It seems unlikely that the strikingly similar effects of clonidine and scopolamine across a range of cognitive and neural measures were realized through two separate neural mechanisms specifically involving noradrenergic and cholinergic neuromodulation. Instead, we would argue that it is more likely that the similar effects of clonidine and scopolamine were caused by interactions between the noradrenergic and cholinergic systems (Briand, Gritton, Howe, Young, & Sarter, 2007). Animal studies have provided evidence that ACh stimulation increases LC activity and that scopolamine reduces noradrenergic baseline activation by antagonizing muscarinic receptors in the LC (Adams & Foote, 1988; Egan & North, 1985). In addition, clonidine has been shown to inhibit cortical ACh release (Acquas et al., 1998), probably by stimulating α_2 receptors in the basal forebrain (cf. Dringenberg & Vanderwolf, 1998). Thus, both clonidine and scopolamine may have reduced NE as well as ACh activity, which could explain their similar effects. Unfortunately, this means that the specific contributions of NE and ACh to learning rate regulation cannot be dissociated based on our pharmacological manipulations. Future studies using more specific manipulations and measures, perhaps in animal models, are required to dissociate the functional roles of these two systems. Our findings underline the importance of studying interactions between different neuromodulatory systems in regulating cognitive function. They also highlight the importance of manipulating more than one neuromodulatory system to prevent premature conclusions regarding the specificity of observed effects to the system under investigation.

Our findings that (i) there was a positive trial-to-trial relationship between P3 amplitude and learning rate and (ii) clonidine and scopolamine suppressed P3 amplitude

at the group mean level seem at odds with our finding that (iii) clonidine and scopolamine did not produce an overall reduction in learning rate. The second and third finding may reflect the complexity of effects of pharmacological manipulations on behavior, such as nonlinear (inverted U shape) relationships between tonic neuromodulator levels and cognitive performance. Relatedly, the drugs may have suppressed both tonic (spontaneous) and phasic (stimulus-evoked) NE and ACh activity (Adams & Foote, 1988), and it is unknown whether learning rate is associated more with one or the other of these two components of activity or with the ratio between them. Individual differences in phasic-to-tonic ratio may also underlie our findings of baseline-dependent drug effects on learning rate, as the effects of pharmacological manipulations on this ratio may depend on someone's natural activity pattern. Future studies directly measuring both phasic and tonic neuromodulatory activity (Bari & Aston-Jones, 2013) and their relationships to belief updating are needed to test these ideas. It is also possible that the drug-induced reduction of outcome-evoked NE and ACh release—as reflected in a smaller P3—did have a suppressive effect on learning rate, but that this was compensated for by increased reliance on other, possibly higher-level, processes. Indeed, P3 amplitude was a partial rather than a complete mediator of the effect of prediction error on learning rate, suggesting that this effect was mediated by processes not reflected in P3 amplitude as well.

In Study 2, we explored to what extent individual differences in learning rate are associated with genetic

variation in specific genes that are presumed to play a role in the noradrenergic pathway by analyzing several SNPs known to affect the dynamics of the LC–NE system. None of the effects of these SNPs on learning rate were statistically significant after correction for multiple comparisons, so Study 2 does not warrant any strong conclusions. However, the study identified two SNPs as promising targets for future research with a larger sample size and greater power to find a statistical difference. At the moment, our findings are merely hypothesis-generating. Although learning rate did not show association with SNPs affecting the synthesis of NE (*DBH*) or sensitivity of the β 1 receptor (*ADRB1*), SNPs located in genes coding for the removal of NE from the synaptic cleft (*NET*) and for sensitivity of the α 2A receptor (*ADRA2A*) did. Notably, the two SNPs that reached suggestive significance can be linked to the two noradrenergic drugs that produced baseline-dependent effects on learning rate in our present and previous studies: atomoxetine blocks the NE transporter (NET), and clonidine has strongest binding affinity for the α 2A receptor. An effect on learning rate of *ADRA2A* would be particularly interesting in light of previous research that found this gene to be a key genetic determinant of P3 amplitude (Liu et al., 2009), which showed a strong relationship with learning rate in our studies. Future genetic association studies with increased statistical power are needed to examine the consistency of the relationship between learning rate and noradrenergic receptor genes, including those coding for the α 1, α 2B, and β 2 receptors.

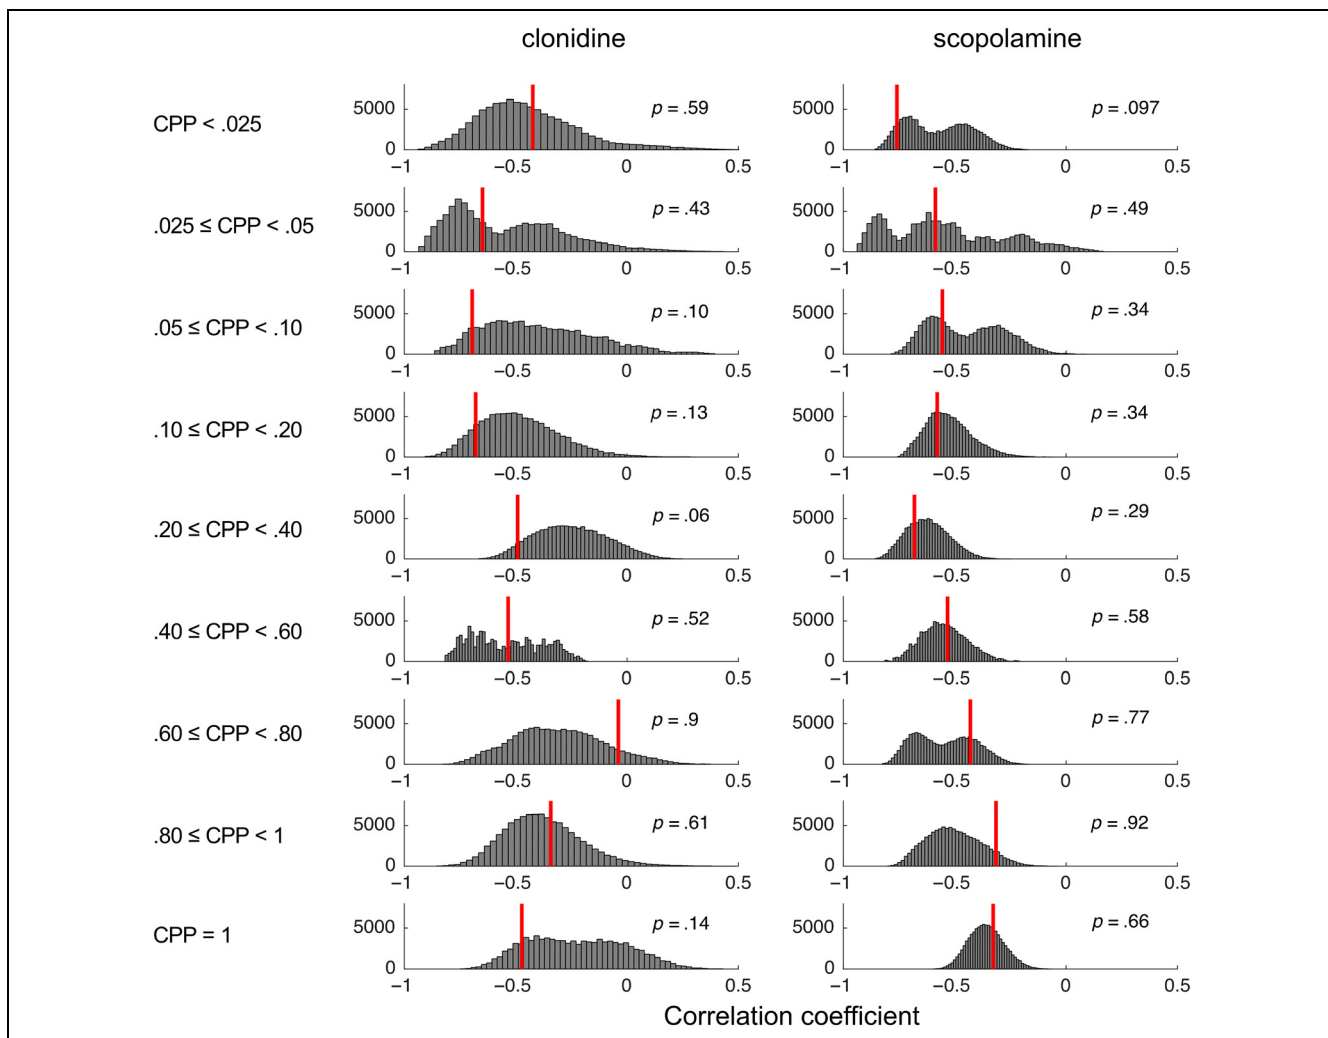


Figure A1. (This figure is related to Figure 6 in the main text.) Baseline-dependent drug effects on learning rate for different levels of change point probability. Trials were divided in nine bins, using the change point probabilities indicated on the left side of each plot as cutoff values. Note that, as change point probability was not equally distributed, the different bins contained different numbers of trials. The gray distributions are permutation distributions of the correlation coefficients predicted by regression to the mean in each bin, and the red lines indicate the observed correlations.

Acknowledgments

This research was supported by a starting grant of the European Research Council (S. N.) and by a VENI grant of the Netherlands Organization for Scientific Research (M. J.). We thank Saskia Heijnen and Eefje Poppelaars for assistance with data collection and Rachel Plak for technical help with the genetic analyses.

Reprint requests should be sent to Marieke Jepma, University of Amsterdam, Nieuwe Achtergracht 29B, 1001 NK, Amsterdam, Netherlands, or via e-mail: m.jepma@uva.nl.

REFERENCES

- Acquas, E., Wilson, C., & Fibiger, H. C. (1998). Pharmacology of sensory stimulation-evoked increases in frontal cortical acetylcholine release. *Neuroscience*, *85*, 73–83.
- Adams, L. M., & Foote, S. L. (1988). Effects of locally infused pharmacological agents on spontaneous and sensory-evoked activity of locus coeruleus neurons. *Brain Research Bulletin*, *21*, 395–400.
- Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*, 403–450.
- Atlas, L. Y., Bolger, N., Lindquist, M. A., & Wager, T. D. (2010). Brain mediators of predictive cue effects on perceived pain. *Journal of Neuroscience*, *30*, 12964–12977.
- Bari, A., & Aston-Jones, G. (2013). Atomoxetine modulates spontaneous and sensory-evoked discharge of locus coeruleus noradrenergic neurons. *Neuropharmacology*, *64*, 53–64.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*, 1214–1221.
- Bennett, D., Murawski, C., & Bode, S. (2015). Single-trial event-related potential correlates of belief updating(1,2,3). *eNeuro*, *2*.
- Briand, L. A., Gritton, H., Howe, W. M., Young, D. A., & Sarter, M. (2007). Modulators in concert for cognition: Modulator interactions in the prefrontal cortex. *Progress in Neurobiology*, *83*, 69–91.

- Brown, S. B., Slagter, H. A., van Noorden, M. S., Giltay, E. J., van der Wee, N. J., & Nieuwenhuis, S. (2016). Effects of clonidine and scopolamine on multiple target detection in rapid serial visual presentation. *Psychopharmacology (Berlin)*, *233*, 341–350.
- Brown, S. B., Tona, K. D., van Noorden, M. S., Giltay, E. J., van der Wee, N. J., & Nieuwenhuis, S. (2015). Noradrenergic and cholinergic effects on speed and sensitivity measures of phasic alerting. *Behavioral Neuroscience*, *129*, 42–49.
- Brown, S. B., van der Wee, N. J., van Noorden, M. S., Giltay, E. J., & Nieuwenhuis, S. (2015). Noradrenergic and cholinergic modulation of late ERP responses to deviant stimuli. *Psychophysiology*, *52*, 1620–1631.
- Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, *18*, 590–596.
- Chumbley, J. R., Flandin, G., Bach, D. R., Daunizeau, J., Fehr, E., Dolan, R. J., et al. (2012). Learning and generalization under ambiguity: An fMRI study. *PLoS Computational Biology*, *8*, e1002346.
- Cools, R., Frank, M. J., Gibbs, S. E., Miyakawa, A., Jagust, W., & D'Esposito, M. (2009). Striatal dopamine predicts outcome-specific reversal learning and its sensitivity to dopaminergic drug administration. *Journal of Neuroscience*, *29*, 1538–1543.
- Cools, R., & Robbins, T. W. (2004). Chemistry of the adaptive mind. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, *362*, 2871–2888.
- Coull, J. T. (2001). Modulation of attention by noradrenergic alpha2-agents varies according to arousal level. *Drug News & Perspectives*, *14*, 5–11.
- Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, *16*, 199–204.
- de Rover, M., Brown, S. B., Boot, N., Hajcak, G., van Noorden, M. S., van der Wee, N. J., et al. (2012). Beta receptor-mediated modulation of the late positive potential in humans. *Psychopharmacology (Berlin)*, *219*, 971–979.
- De Taeye, L., Vonck, K., van Bochove, M., Boon, P., Van Roost, D., Mollet, L., et al. (2014). The P3 event-related potential is a biomarker for the efficacy of vagus nerve stimulation in patients with epilepsy. *Neurotherapeutics*, *11*, 612–622.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*, 9–21.
- Devauges, V., & Sara, S. J. (1990). Activation of the noradrenergic system facilitates an attentional shift in the rat. *Behavioural Brain Research*, *39*, 19–28.
- Devoto, P., & Flore, G. (2006). On the origin of cortical dopamine: Is it a co-transmitter in noradrenergic neurons? *Current Neuropharmacology*, *4*, 115–125.
- Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating. *Behavioral and Brain Sciences*, *11*, 357–374.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, *15*, 495–506.
- Doya, K. (2007). *Bayesian brain: Probabilistic approaches to neural coding*. Cambridge, MA: MIT Press.
- Dringenberg, H. C., & Vanderwolf, C. H. (1998). Involvement of direct and indirect pathways in electrocorticographic activation. *Neuroscience and Biobehavioral Reviews*, *22*, 243–257.
- Egan, T. M., & North, R. A. (1985). Acetylcholine acts on m2-muscarinic receptors to excite rat locus coeruleus neurones. *British Journal of Pharmacology*, *85*, 733–735.
- Frank, M. J., & Fossella, J. A. (2011). Neurogenetics and pharmacology of learning, motivation, and cognition. *Neuropsychopharmacology*, *36*, 133–152.
- Friston, K. (2009). The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, *13*, 293–301.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*, 127–138.
- Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., et al. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, *8*, e1002327.
- Gibbs, A. A., Bautista, C. E., Mowlem, F. D., Naudts, K. H., & Duka, D. T. (2014). Catechol-O-methyltransferase val158met genotype determines effect of reboxetine on emotional memory in healthy male volunteers. *Journal of Psychiatry & Neuroscience*, *39*, E24–E31.
- Gratton, G., Coles, M. G., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology*, *55*, 468–484.
- Hansenne, M. (2000). The p300 cognitive event-related potential. I. Theoretical and psychobiologic perspectives. *Clinical Neurophysiology*, *30*, 191–210.
- Hasselmo, M. E. (1995). Neuromodulation and cortical function: Modeling the physiological basis of behavior. *Behavioural Brain Research*, *67*, 1–27.
- Hasselmo, M. E. (2006). The role of acetylcholine in learning and memory. *Current Opinion in Neurobiology*, *16*, 710–715.
- Hsieh, C. Y., Cruikshank, S. J., & Metherate, R. (2000). Differential modulation of auditory thalamocortical and intracortical synaptic transmission by cholinergic agonist. *Brain Research*, *880*, 51–64.
- Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E., et al. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, *80*, 519–530.
- Jepma, M., Murphy, P. R., Nassar, M. R., Rangel-Gomez, M., Meeter, M., & Nieuwenhuis, S. (2016). Catecholaminergic regulation of learning rate in a dynamic environment. *PLoS Computational Biology*, *12*, e1005171.
- Kahkonen, S., Ahveninen, J., Pekkonen, E., Kaakkola, S., Huttunen, J., Ilmoniemi, R. J., et al. (2002). Dopamine modulates involuntary attention shifting and reorienting: An electromagnetic study. *Clinical Neurophysiology*, *113*, 1894–1902.
- Kelly, C., & Price, T. D. (2005). Correcting for regression to the mean in behavior and ecology. *Am Nat*, *166*, 700–707.
- Kimura, F., Fukuda, M., & Tsumoto, T. (1999). Acetylcholine suppresses the spread of excitation in the visual cortex revealed by optical recording: Possible differential effect depending on the source of input. *European Journal of Neuroscience*, *11*, 3597–3609.
- Kobayashi, K., & Hsu, M. (2017). Neural mechanisms of updating under reducible and irreducible uncertainty. *Journal of Neuroscience*, *37*, 6972–6982.
- Kolossa, A., Kopp, B., & Fingscheidt, T. (2015). A computational analysis of the neural bases of Bayesian inference. *Neuroimage*, *106*, 222–237.
- Krugel, L. K., Biele, G., Mohr, P. N., Li, S. C., & Heekeren, H. R. (2009). Genetic variation in dopaminergic neuromodulation influences the ability to rapidly and flexibly adapt decisions. *Proceedings of the National Academy of Sciences, U.S.A.*, *106*, 17951–17956.
- Kwisthout, J., Bekkering, H., & van Rooij, I. (2017). To be precise, the details don't matter: On predictive processing, precision, and level of detail of predictions. *Brain and Cognition*, *112*, 84–91.

- Lapiz, M. D., Bondi, C. O., & Morilak, D. A. (2007). Chronic treatment with desipramine improves cognitive performance of rats in an attentional set-shifting test. *Neuropsychopharmacology*, *32*, 1000–1010.
- Lapiz, M. D., & Morilak, D. A. (2006). Noradrenergic modulation of cognitive function in rat medial prefrontal cortex as measured by attentional set shifting capability. *Neuroscience*, *137*, 1039–1049.
- Liu, J., Kiehl, K. A., Pearson, G., Perrone-Bizzozero, N. I., Eichele, T., & Calhoun, V. D. (2009). Genetic determinants of target and novelty-related event-related potentials in the auditory oddball response. *Neuroimage*, *46*, 809–816.
- Luksys, G., Gerstner, W., & Sandi, C. (2009). Stress, genotype and norepinephrine in the prediction of mouse behavior using reinforcement learning. *Nature Neuroscience*, *12*, 1180–1186.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*, 177–190.
- Marshall, L., Mathys, C., Ruge, D., de Berker, A. O., Dayan, P., Stephan, K. E., et al. (2016). Pharmacological fingerprints of contextual uncertainty. *PLoS Biology*, *14*, e1002575.
- McGaughy, J., Ross, R. S., & Eichenbaum, H. (2008). Noradrenergic, but not cholinergic, deafferentation of prefrontal cortex impairs attentional set-shifting. *Neuroscience*, *153*, 63–71.
- McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, *84*, 870–881.
- Moran, R. J., Campo, P., Symmonds, M., Stephan, K. E., Dolan, R. J., & Friston, K. J. (2013). Free energy, precision and learning: The role of cholinergic neuromodulation. *Journal of Neuroscience*, *33*, 8227–8236.
- Nassar, M. R., Bruckner, R., Gold, J. I., Li, S. C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, *7*, 11609.
- Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasley, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, *15*, 1040–1046.
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, *30*, 12366–12378.
- Newman, L. A., Darling, J., & McGaughy, J. (2008). Atomoxetine reverses attentional deficits produced by noradrenergic deafferentation of medial prefrontal cortex. *Psychopharmacology (Berlin)*, *200*, 39–50.
- Nieuwenhuis, S. (2011). Learning, the P3, and the locus coeruleus–norepinephrine system. In J. S. R. Mars, M. Rushworth, & N. Yeung (Eds.), *Neural basis of motivational and cognitive control* (pp. 209–222). Oxford: Oxford University Press.
- Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus–norepinephrine system. *Psychological Bulletin*, *131*, 510–532.
- O'Reilly, J. X., Schuffelgen, U., Cuell, S. F., Behrens, T. E., Mars, R. B., & Rushworth, M. F. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *110*, E3660–E3669.
- Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J. P. (2013). The neural representation of unexpected uncertainty during value-based decision making. *Neuron*, *79*, 191–201.
- Pineda, J. A., Foote, S. L., & Neville, H. J. (1989). Effects of locus coeruleus lesions on auditory, long-latency, event-related potentials in monkey. *Journal of Neuroscience*, *9*, 81–93.
- Pitman, E. J. G. (1939). A note on normal correlation. *Biometrika*, *31*, 9–12.
- Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, *118*, 2128–2148.
- Rangel-Gomez, M., Hickey, C., van Amelsvoort, T., Bet, P., & Meeter, M. (2013). The detection of novelty relies on dopaminergic signaling: Evidence from apomorphine's impact on the novelty N2. *PLoS One*, *8*, e66469.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extraclassical receptive-field effects. *Nature Neuroscience*, *2*, 79–87.
- Sara, S. J. (2009). The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience*, *10*, 211–223.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, *249*, 892–895.
- Set, E., Saez, I., Zhu, L., Houser, D. E., Myung, N., Zhong, S., et al. (2014). Dissociable contribution of prefrontal and striatal dopaminergic genes to learning in economic games. *Proceedings of the National Academy of Sciences, U.S.A.*, *111*, 9615–9620.
- Seu, E., Lang, A., Rivera, R. J., & Jentsch, J. D. (2009). Inhibition of the norepinephrine transporter improves behavioral flexibility in rats and monkeys. *Psychopharmacology (Berlin)*, *202*, 505–519.
- Silvetti, M., Seurinck, R., van Bochove, M. E., & Verguts, T. (2013). The influence of the noradrenergic system on optimal control of neural plasticity. *Frontiers in Behavioral Neuroscience*, *7*, 160.
- Svensson, T. H., Bunney, B. S., & Aghajanian, G. K. (1975). Inhibition of both noradrenergic and serotonergic neurons in brain by the alpha-adrenergic agonist clonidine. *Brain Research*, *92*, 291–306.
- Tait, D. S., Brown, V. J., Farovik, A., Theobald, D. E., Dalley, J. W., & Robbins, T. W. (2007). Lesions of the dorsal noradrenergic bundle impair attentional set-shifting in the rat. *European Journal of Neuroscience*, *25*, 3719–3724.
- Vossel, S., Bauer, M., Mathys, C., Adams, R. A., Dolan, R. J., Stephan, K. E., et al. (2014). Cholinergic stimulation enhances Bayesian belief updating in the deployment of spatial attention. *Journal of Neuroscience*, *34*, 15735–15742.
- Wager, T. D., van Ast, V. A., Hughes, B. L., Davidson, M. L., Lindquist, M. A., & Ochsner, K. N. (2009). Brain mediators of cardiovascular responses to social threat, part II: Prefrontal-subcortical pathways and relationship with anxiety. *Neuroimage*, *47*, 836–851.
- Wager, T. D., Waugh, C. E., Lindquist, M., Noll, D. C., Fredrickson, B. L., & Taylor, S. F. (2009). Brain mediators of cardiovascular responses to social threat: Part I: Reciprocal dorsal and ventral sub-regions of the medial prefrontal cortex and heart-rate reactivity. *Neuroimage*, *47*, 821–835.
- Yu, A. J., & Dayan, P. (2002). Acetylcholine in cortical inference. *Neural Networks*, *15*, 719–730.
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*, 681–692.