Mood dynamics are associated with learning and not choice

Bastien Blain^{1*}, Robb B Rutledge^{1,2}

¹Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London WC1B 5EH, United Kingdom

²Wellcome Centre for Human Neuroimaging, University College London, London WC1N 3BG, United Kingdom

* b.blain@ucl.ac.uk

Updating predictions about which stimuli are associated with reward is an important aspect of adaptive behaviour believed to relate to prediction errors, the difference between experienced and predicted outcomes. Behavioural sensitivity to prediction errors flexibly adapts to environmental statistics. Prediction errors also influence affective states during risky choice tasks that do not require learning, but the relationship between emotions and adaptive behaviour is unknown. Here, using computational modelling we found that mood dynamics, like behaviour, are sensitive to learning-relevant model variables (i.e., probability prediction error). Unlike behaviour, mood dynamics are not sensitive to model variables that influence choice (i.e., expected value), and increasing volatility does not reduce how many trials influence affective state. Finally, depressive symptoms reduce overall mood more in volatile than stable environments. Our findings suggest that mood dynamics are selective for variables relevant to adaptive behaviour and suggest a greater role for mood in learning than choice.

Introduction

Decisions are guided by beliefs about states of the world. Some states are directly observable, like the potential prize for a bet. Other states, like the probability of winning, may not be directly observable but can be inferred from past events. Thus, learning from experience is essential for adaptive behaviour. In the standard theoretical framework, learning is driven by how unexpected the outcome is (i.e., by the prediction error: the difference between outcome and prediction)¹. Sensitivity to the prediction error (i.e., the learning rate) flexibly adapts to environmental statistics. The decisions of both humans and non-human primates are consistent with a higher learning rate in more volatile environments, and subjects are more likely to stay on the same option after positive compared to negative feedback when reward probabilities change more frequently²⁻⁶.

Emotions are widely believed to play a role in adaptive behaviour⁷, but no computational framework exists to link the two. Unexpected outcomes influence affective states, so that bad outcomes feel worse when unexpected than when expected, but good outcomes feel better when unexpected than when expected^{8,9}. It has recently been shown that reward expectations and reward prediction errors (RPEs), the difference between experienced and predicted rewards, can explain changes in affective state in the context of decision making under uncertainty when learning is not required^{10,11}. Mood has been proposed to represent environmental momentum, whether an environment is getting better or worse, which could be a useful variable for adaptive behaviour^{12,13}.

Impairments in reward and emotion processing are associated with depression⁹. Learning in depression has been extensively studied, but there is not consistent evidence for a specific deficit¹⁴⁻²⁵ (reviewed in ref. ²⁶, meta-analysis in ref. ²⁷). The ability of individuals to appropriately adjust learning rates to environmental volatility is associated with anxiety symptoms⁵. Individuals with high trait anxiety show reduced ability to adjust updating of outcome expectancies for aversive outcomes to the volatility of the environment compared to individuals with low trait anxiety. Failure to appropriately adjust learning to environmental volatility has not been established in depression. Affective states reflect subjective estimates of uncertainty, which predict the dynamics of subjective and physiological stress responses²⁸. Overall mood during risky decision-making tasks is reduced with increasing depression severity, both in the laboratory and using remote smartphone-based data collection²⁹.

Here, our goal was to adapt an established reinforcement learning task with two levels of environmental volatility to manipulate learning rates (figure 1) and to determine how behavioural sensitivity to prediction errors relates to mood dynamics. We addressed the following questions: (1) Is mood more sensitive to learning- or choice-relevant variables in established reinforcement learning models? (2) Do mood dynamics adjust to environmental volatility? (3) Are mood dynamics affected by depression in the context of learning?

Results

To investigate the relationship between mood dynamics and adaptive behaviour, we adapted a reinforcement learning task used in previous studies^{2,6}. Participants repeatedly made decisions between two cars racing against each other (corresponding to a one-armed bandit because only one car can win each race). The reward magnitude associated with each car was explicit and assigned randomly on each trial, whereas the outcome probability was implicit and had to be inferred from the outcomes of previous trials. Participants ($n=75$) were exposed to two conditions (counterbalanced across participants). In the stable environment, the ‘best’ car won 80% of the races on average. In the volatile environment, the ‘best’ car had an 80% probability of winning and the ‘best’ car switched every 20 trials (see figure 1). Participants were asked to report their current happiness after every 3-4 trials. Unlike in previous studies in humans using a similar task^{2,5}, participants were informed that they would complete stable and volatile conditions. Which environment they were in was indicated at the start of each condition. Subjects were given no guidance as to how they should use this information, although we expected this manipulation to increase the difference in behavioural sensitivity between environments relative to previous studies.

Figure 1. Experimental design

Subjects ($n=75$) performed a one-armed bandit reinforcement learning task, choosing repeatedly between two cars. They were instructed to maximize their cumulative points. In the stable environment (80 trials), the probability to win for the best car was 80%. In the volatile environment (80 trials), reward probabilities switched between 80% for one car and 80% for the other car every 20 trials. Environment order was counterbalanced across subjects (see the Methods). The reward available for each car was randomly determined on each trial and unrelated to the probability of winning. Every three to four trials, subjects were asked to report “How happy are you right now?” by moving a cursor on a line. Each trial started with a fixation symbol in the centre of the screen. Then, the stimuli were displayed but choice was not permitted. The potential reward for each car was then displayed and participants were free to choose an option without any time constraints. The chosen option was outlined by a yellow frame. Finally, the outcome was displayed. Both the car and the reward magnitude frames were green if the chosen car won the race (example shown). The car frame was red and crossed out if the chosen car lost (example shown in inset).

Learning rates change with environmental volatility

Participants chose the option with the higher expected value (i.e., choice accuracy) more often than chance (see figure 2A) in the stable environment ($82.1 \pm 1.1\%$ (mean $\pm$ SEM), $t_{74} = 28.3$, $P < 10^{-40}$) and in the volatile environment ($73.3 \pm 1.0\%$, $t_{74} = 22.6$, $P < 10^{-34}$). Participants chose the higher probability option more often than chance in the stable environment ($80.8 \pm 1.0\%$

(mean $\pm$ SEM), $t_{74} = 30.8$, $P < 10^{-43}$) and in the volatile environment ($62.6 \pm 1.1\%$, $t_{74} = 11.3$, $P < 10^{-16}$). To ensure that participants incorporated reward magnitude into their decisions, we considered only trials where the car with the lower outcome probability had the higher expected value. Subjects chose the low probability car in these trials more often than chance (stable: $61.2 \pm 3.0\%$, $t_{74} = 3.7$, $P < 0.001$; volatile: $75.6 \pm 2.7\%$, $t_{74} = 9.5$, $P < 10^{-13}$). These results are consistent with participants integrating both probability and reward to make their decisions.

Multiplicative and additive models that integrate probability and reward in different ways have been widely used to explain behaviour across stable and volatile environments^{2,4-6}. Both models include a learning component updating the probability estimate for each car to win (equations 1 and 2), based on the probability prediction error (PPE), the difference between the outcome (0 or 1) and the estimated probability of winning:

$$P_{car\ 1\ wins}(t) = P_{car\ 1\ wins}(t-1) + \alpha PPE(t), \quad (1)$$

$$PPE(t) = Outcome(t) - P_{car\ 1\ wins}(t-1), \quad (2)$$

where $P_{car\ 1\ wins}(t)$ is the estimated probability for car 1 winning on trial t , and α is the learning rate. For choices to car 2, a similar equation applies. Multiplicative and additive models differ regarding implementation of choice predictions with probability and reward magnitude either integrated multiplicatively^{2,5} or additively^{4,6,30}.

The multiplicative selector resembles the maximisation of expected utility common to economic decision models³¹:

$$EU_{car1}(t) = \max [\min [\eta(P_{car1\ wins}(t) - 0.5) + 0.5, 1], 0] \times Reward(t), \quad (3)$$

$$P_{choosing\ car1}(t) = \frac{1}{1 + \exp(-\beta(EU_{car1} - EU_{car2}))}, \quad (4)$$

where η is a free parameter related to the level of risk aversion and β is the inverse temperature (i.e., choice stochasticity or precision). The multiplicative model captured choice data in both stable (pseudo- $r^2 = 0.54 \pm 0.02$, mean $\pm$ SEM) and volatile environments (pseudo- $r^2 = 0.41 \pm 0.02$).

In contrast, the additive selector is implemented as follows:

$$P_{car\ 1\ chosen}(t) = \frac{1}{1 + \exp(-\beta(\phi\Delta Probability + (1-\phi)\Delta Reward))} \quad (5)$$

Where $\Delta Probability$ corresponds to the difference in estimated probability between the options and $\Delta Reward$ corresponds to the difference in reward magnitude between the options, and ϕ represents the relative weight of probability and magnitude on choice. The additive model captured choice data in the stable (pseudo- $r^2 = 0.62 \pm 0.02$) and volatile environments (pseudo- $r^2 = 0.45 \pm 0.02$, see figure 2A).

Model comparison demonstrated that the additive model better explained choice data with the same number of parameters as the multiplicative model in both the stable ($\Delta BIC = 690$) and the volatile environment ($\Delta BIC = 3328$, see table 1). Our results are consistent with similar findings obtained in highly trained non-human primates using a task design in which changes between stable and volatile environments were signalled⁶ as in the present study.

We next examined whether model fits were consistent with subjects integrating both potential reward magnitudes and probabilities to make decisions. The relative weight of probability and reward magnitude in the additive model (ϕ) was balanced on average (stable: $\phi = 0.57 \pm 0.018$ (mean $\pm$ SEM); volatile: $\phi = 0.47 \pm 0.029$), suggesting that subjects integrated both

probabilities and reward magnitudes to make decisions. Omitting ϕ and evaluating simpler models that considered only probabilities (α and β) or reward magnitudes (β only) resulted in worse fits (see table 1). Lower BIC for additive and multiplicate models confirmed that subjects considered both probability and reward magnitude when making decisions.

Model	Number of parameters	Stable pseudo- r^2	Volatile pseudo- r^2	Stable BIC	Volatile BIC	Stable Δ BIC	Volatile Δ BIC
Additive	3	0.62	0.45	4134	5594	0	0
Multiplicative	3	0.54	0.41	4824	5922	690	328
Probability only	2	0.35	0.16	6104	7632	1969	2038
Magnitude only	1	0.14	0.23	7473	6738	3339	1144

Table 1. Choice model comparison results

The ‘Additive’ model refers to a model implementing a weighted sum of probability difference and reward magnitude difference when making decisions^{4,6,30,32,33}. The ‘Multiplicative’ model refers to the model first used to describe behaviour in this task^{2,5}. The ‘Probability only’ model includes only the probability component of the additive model and the ‘Magnitude only’ model includes only the magnitude component of the additive model. Δ BIC refers to the Bayesian Information Criterion computed for each model compared to the additive model, the preferred model.

We then asked whether greater environmental volatility was associated with higher learning rates as observed in previous studies^{2,5,6}. A simple prediction for standard reinforcement learning models is that subjects should stay more on the same option after winning than losing (figure 2B). In line with this prediction, subjects had stayed more on the same option after winning than losing for the stable environment (difference in choice proportion, high probability car: $4.1 \pm 1.9\%$, mean $\pm$ SEM, $t_{74} = 2.1$, $P < 0.05$; low probability car: $1.3 \pm 3.2\%$, $t_{74} = 0.40$, $P = 0.70$). Subjects stayed on the same option more after winning than losing in the volatile environment (difference in choice proportion, high probability car: $22.1 \pm 2.3\%$, $t_{74} = 9.5$, $P < 10^{-13}$, low probability car: $21.4 \pm 3.7\%$, $t_{74} = 5.8$, $P < 10^{-6}$). Subjects stayed on the same option after winning compared to losing more in volatile than stable environments (difference volatile – stable, high probability car: $18.1 \pm 2.7\%$, $t_{74} = 6.6$, $P < 10^{-8}$, low probability car: $19.8 \pm 4.1\%$, $t_{74} = 4.7$, $P < 10^{-4}$).

This pattern of results is captured by a reinforcement learning model fit separately to each environment (predicted change in choice proportion in stable, high probability car: $8.8 \pm 1.6\%$, $t_{74} = 5.6$, $P < 10^{-6}$, low probability car: $6.2 \pm 2.6\%$, $t_{70} = 2.3$, $P < 0.05$; in volatile, high probability car: $21.5 \pm 2.2\%$, $t_{74} = 10.1$, $P < 10^{-14}$, low probability car: $21.2 \pm 2.7\%$, $t_{73} = 7.8$, $P < 10^{-10}$; difference volatile – stable: $12.8 \pm 2.4\%$, $t_{74} = 5.3$, $P < 10^{-6}$). We found that learning rates (figure 2C) were substantially higher in volatile than stable environment (stable $\alpha = 0.16 \pm 0.02$, mean $\pm$ SEM; volatile $\alpha = 0.48 \pm 0.03$; difference volatile – stable, high probability car: $12.7 \pm 2.5\%$, $t_{74} = 5.2$, $P < 10^{-5}$, low probability car: $14.3 \pm 3.6\%$, $t_{69} = 3.8$, $P < 0.001$). Overall, these results demonstrate that the learning rate increases substantially in the volatile compared to the stable environment, in line with previous studies^{2,5,6}.

Figure 2. Learning rate adapts to environmental volatility

(A) Participants chose the option with the highest expected value 82% of the time in the stable environment (blue curve, left panel) and 73% of the time in the volatile environment (orange curve, right panel). The additive model containing three parameters (a learning rate determining the sensitivity to prediction error, an inverse temperature reflecting choice stochasticity, and a relative weight for probability and reward magnitude in choice) fitted choice data well (black dashed lines) in the stable environment (mean pseudo- $r^2 = 0.62$) and the volatile environment (mean pseudo- $r^2 = 0.45$).

(B) Participants chose more often the option with the higher probability in the stable environment compared to the volatile environment. Critically, participants stayed on the same option more often if choosing that option resulted in the car winning (light orange) compared to the car losing (dark orange) in the volatile environment compared to the stable environment (light blue and dark blue represent staying after winning and losing, respectively). This suggests that participant behaviour was more sensitive to feedback in the volatile than stable environment, as an agent with a higher learning rate would be. Additive model predictions show a similar difference in feedback sensitivity across environments (purple).

(C) Learning rates were higher in the volatile environment (orange) compared to the stable environment (blue). This was true for participants that experienced the stable environment before (stable 1) or after (stable 2) the volatile environment. Error bars represent SEM. * $P < 0.05$, *** $P < 0.001$

Happiness is more sensitive to learning-relevant than choice-relevant variables

We next examined how happiness changes over time. Subjects varied their happiness ratings in both the stable ($SD = 24.2 \pm 1.1$, mean $\pm$ SEM) and volatile ($SD = 25.0 \pm 1.1$) environments. They were happier on average after winning than after losing (stable: 63.8 ± 1.9 vs 34.5 ± 2.0 , $t_{74} = 13.1$, $P < 10^{-20}$; volatile: 61.9 ± 1.8 vs 33.6 ± 2.0 , $t_{74} = 13.7$, $P < 10^{-21}$, figure 3A). Participants were happier on average in the stable environment than in the volatile environment (stable: 55.0 ± 1.7 , volatile: 49.5 ± 1.6 , $t_{74} = 3.4$, $P < 0.01$). This effect may be at least partly

due to lower choice accuracy in volatile environments and the difference in average happiness was correlated with the difference in choice accuracy (Spearman $\rho = 0.24$, $P < 0.05$).

Figure 3. Happiness is associated with probability and probability prediction error

(A) Most participants were happier when the chosen car was winning than when it was losing (97% of participants in the stable environment, 96% in the volatile environment, in the left and right panel, respectively).

(B) Momentary happiness was best explained by a model (black dotted lines) including both the chosen probability estimate and the probability prediction error (PPE) derived from the additive choice model in addition to a forgetting factor and a baseline mood parameter, for both the stable (mean $r^2 = 0.58$) and the volatile (mean $r^2 = 0.63$) environment.

(C) The chosen probability and the PPE parameters were significantly different from 0 for both environments. Both variables are significantly associated with changes in affective state over time. Error bars represent SEM. *** $P < 0.001$

Previous studies have reported that momentary happiness in response to outcomes in a probabilistic reward task were explained by recent RPEs when maximizing cumulative reward does not require learning^{10,11}. In our task, maximizing cumulative reward requires learning the outcome probability. In this context, PPEs (depending on whether the outcome was a win or a loss and the subjective probability of that outcome) are relevant to learning but RPEs (depending on the magnitude of the reward received and the expected value of the chosen option) are not relevant to learning or future behaviour. Reward information was choice relevant and choices were driven by the (additive) integration of the estimated outcome probability and the potential reward magnitude. Therefore, we tested whether happiness was more strongly associated with the PPEs used for learning or alternatively by RPEs that incorporate learning-irrelevant reward magnitudes. We compared two models that both use the subjective probability as estimated in the additive choice model to compute prediction errors:

$$Happiness(t) = w_0 + w_{\widehat{PPE}} \sum_{j=1}^t \gamma^{t-j} \widehat{PPE}, \quad (6)$$

where $\widehat{PPE}$ refers to the PPE, defined as the difference between the outcome (1 for win, 0 for loss) and the subjective probability estimated from the additive choice model, w_0 is a constant term, $w_{\widehat{PPE}}$ is a weight capturing the influence of past PPEs, and $0 \leq \gamma \leq 1$ is a forgetting factor that makes events in more recent trials more influential than those in earlier trials;

$$Happiness(t) = w_0 + w_{\widehat{RPE}} \sum_{j=1}^t \gamma^{t-j} \widehat{RPE}, \quad (7)$$

where $\widehat{RPE}$ is the difference between reward magnitude and the expected value of the chosen option computed based on the subjective probability estimated from the additive choice model. Reward magnitudes were rescaled from 0 to 1.

Mood fluctuations were better explained by a model including past PPEs than by a model including past RPEs, both in the stable ($BIC_{ppe} = -5958$, $BIC_{rpe} = -5635$, $\Delta BIC = 323$) and in the volatile ($BIC_{ppe} = -5620$, $BIC_{rpe} = -5340$, $\Delta BIC = 280$) environments (see table 2 and figure 4A). This result holds for a broader model space including other definitions of the prediction error terms (see table 2):

$$Happiness(t) = w_0 + w_{PPE} \sum_{j=1}^t \gamma^{t-j} PPE, \quad (8)$$

where PPE refers to the objective PPE defined as the difference between the outcome sign and the objective probability of the chosen option (0.2 or 0.8), and also:

$$Happiness(t) = w_0 + w_{RPE} \sum_{j=1}^t \gamma^{t-j} RPE, \quad (9)$$

where RPE is computed by taking the difference between the reward magnitude and the objective expected value (potential reward multiplied by the objective probability of the chosen option) as above.

We then tested whether mood fluctuations were additionally sensitive to current expectations, as shown in risky choice tasks that do not require learning (e.g., ref. ^{10,11}). Again, two types of expectations may explain mood fluctuations: a subjective probability relevant to learning, or expected values that incorporate learning-irrelevant reward magnitudes. We compared the following models:

$$Happiness(t) = w_0 + w_{\hat{p}} \sum_{j=1}^t \gamma^{t-j} (\hat{p} - 0.5) + w_{\widehat{PPE}} \sum_{j=1}^t \gamma^{t-j} \widehat{PPE}, \quad (10)$$

where $\hat{p}$ is the probability estimated with the additive choice model, and:

$$Happiness(t) = w_0 + w_{\widehat{EV}} \sum_{j=1}^t \gamma^{t-j} (\widehat{EV} - \overline{EV}) + w_{\widehat{PPE}} \sum_{j=1}^t \gamma^{t-j} \widehat{PPE}, \quad (11)$$

259 where $\widehat{EV}$ is the product between $\widehat{P}$ and the reward magnitude and this term is mean-centred.

260

261 **Figure 4. Happiness is more strongly associated with learning than choice**

262 (A) Comparison between the r^2 for the happiness model including a PPE term (denoted $\widehat{PPE}$) estimated
 263 in the additive choice model (y axis) and the r^2 for the happiness model including a RPE term instead
 264 (denoted $\widehat{RPE}$). Both models had the same number of parameters. The PPE model accounted for more
 265 variance in mood ratings on average in both stable (blue) and volatile (orange) environments. Dots above
 266 the dashed line correspond to subjects for whom more variance in happiness is explained by the PPE
 267 compared to the RPE model.

268 (B) The $\widehat{PPE}$ model including the chosen estimated probability (denoted $\widehat{P}$ and estimated from the
 269 additive choice model) better explained happiness ratings than a $\widehat{PPE}$ model including expected value
 270 (denoted $\widehat{EV}$) for both the stable (blue) and volatile (orange) environments with both models having the
 271 same number of parameters. Dots above the dashed line correspond to subjects where more variance in
 272 happiness is explained by the $\widehat{P} + \widehat{PPE}$ compared to the $\widehat{EV} + \widehat{PPE}$ model.

273

274 The model including choice probability better explained happiness ratings (stable: mean $r^2 =$
 275 0.58 ; volatile: mean $r^2 = 0.63$) than the model including the expected value in the stable
 276 ($BIC_{p+ppe} = -6134$, $BIC_{ev+ppe} = -5944$, $\Delta BIC = 189$) and in the volatile ($BIC_{p+ppe} = -6214$,
 277 $BIC_{ev+ppe} = -5764$, $\Delta BIC = 450$) environments (see table 2 and figure 4B). The probability and
 278 PPE error weights were significantly different from 0 at the group level in both the stable ($w_p =$
 279 0.21 ± 0.03 , $t_{74} = 7.4$, $P < 10^{-9}$; $w_{ppe} = 0.34 \pm 0.024$, $t_{74} = 14.3$, $P < 10^{-22}$) and the volatile
 280 environments ($w_p = 0.25 \pm 0.04$, $t_{74} = 6.5$, $P < 10^{-8}$; $w_{ppe} = 0.30 \pm 0.02$, $t_{74} = 14.9$, $P < 10^{-23}$,
 281 figure 3C).

282

283

284

Model	Number of parameters	Stable mean r^2	Volatile mean r^2	Stable BIC	Volatile BIC	Stable Δ BIC	Volatile Δ BIC
$\widehat{PPE}$	3	0.49	0.44	-5958	-5620	176	594
$\widehat{RPE}$	3	0.40	0.38	-5635	-5340	499	874
PPE	3	0.48	0.43	-5901	-5505	232	709
RPE	3	0.39	0.37	-5596	-5315	538	899
$\hat{P} + \widehat{PPE}$	4	0.58	0.63	-6134	-6214	0	0
$\widehat{EV} + \widehat{PPE}$	4	0.55	0.54	-5945	-5764	189	450
$\widehat{EV} + \widehat{RPE}$	4	0.46	0.49	-5572	-5493	561	720

Table 2. Happiness model comparison results

PPE is probability prediction error, RPE is reward prediction error, P is the probability, and EV is the expected value. The hat over a variable refers to that variable incorporating the trial-by-trial choice probability estimated from the additive model. Δ refers to the comparison of the model scores (using the Bayesian Information Criterion) to the preferred model.

Sensitivity of mood dynamics to learning variables does not depend on volatility

We found that the learning rate (i.e., behavioural sensitivity to PPE) was approximately three times higher in the volatile compared to the stable environment. Furthermore, mood dynamics were highly sensitive to PPE. We next focused on whether sensitivity of mood dynamics to PPE changed with environmental volatility. The probability and PPE weights estimated from happiness ratings were not different between environments (stable – volatile: $\Delta w_p = -0.04 \pm 0.05$, $t_{74} = -0.96$, $P = 0.34$; $\Delta w_{ppe} = 0.03 \pm 0.02$, $t_{74} = 1.9$, $P = 0.07$; see figure 3C). Moreover, PPE weights were not correlated with the learning rate in the stable (Spearman $\rho = -0.11$, $P = 0.36$) and volatile (Spearman $\rho = 0.03$, $P = 0.77$) environments nor was the difference of PPE weights across environments related to the difference in learning rate (Spearman $\rho = -0.001$, $P = 0.94$). Instead, PPE weights were consistent across environments (Spearman $\rho = 0.66$, $P < 10^{-10}$, see figure 5A).

The happiness model forgetting factor γ determines how many previous trials influence current affective state. When γ is equal to 1, mood is equally influenced by all previous trials, when γ is equal to 0, mood is influenced by only the most recent trial. If the change in forgetting factor mirrors the observed change in learning rate, forgetting factors should be lower in the volatile than in the stable environment, reflecting integration over fewer trials and consistent with the higher learning rates observed in volatile environments. Instead, the forgetting factor was higher in the volatile environment (stable: $\gamma = 0.57 \pm 0.03$, volatile $\gamma = 0.65 \pm 0.03$, corresponding to current happiness being influenced by 6-7 previous trials on average in both environments, stable – volatile: $\Delta\gamma = 0.08 \pm 0.04$, $t_{74} = 2.36$, $P = 0.02$; see figure 5B). Furthermore, the change in happiness forgetting factor was not correlated with the parameter difference between environments (Spearman $\rho = -0.13$, $P = 0.26$). A linear regression with ten previous probability prediction errors as independent variables confirmed this model-based result (see supplementary figure 1). To further test for a relationship between the forgetting factor and the learning rate, we switched the learning rates estimated from stable and volatile conditions in each individual before re-fitting happiness data (i.e., we used the ‘wrong’ learning rate to

estimate probabilities and PPEs before fitting the happiness model and estimating a forgetting factor). This did not substantially affect estimates of the forgetting factor (stable $\gamma = 0.56 \pm 0.03$, volatile $\gamma = 0.63 \pm 0.03$, stable – volatile: $\Delta\gamma = 0.07 \pm 0.04$, $t_{74} = 1.97$, $P = 0.05$). The resulting forgetting factor estimates were highly correlated with forgetting factors estimated using the actual learning rates (stable: $r = 0.96$, $P < 10^{-41}$; volatile: $r = 0.93$, $P < 10^{-33}$). The happiness forgetting factor was consistent across environments (Spearman $\rho = 0.34$, $P < 0.01$, see figure 5C), suggesting that the number of previous trials that affective state depends on may be a trait-like feature of individuals.

Figure 5. Forgetting factors are consistent across stable and volatile environments

(A) Weights for PPEs in determining happiness were consistent across environments.

(B) The happiness forgetting factor did not change between stable (blue) and volatile (orange) environments, regardless of testing order.

(C) Happiness forgetting factors were consistent across environments. Error bars represent SEM. ** $P < 0.01$, *** $P < 0.001$

Depressive symptoms are associated with reduced happiness in volatile environments

Previous studies have linked learning rates to anxiety^{5,34} and individuals with high trait anxiety showed less ability to appropriately adjust updating of outcome expectancies between stable and volatile environments. We found that depressive symptoms (PHQ) were uncorrelated with choice accuracy (stable: Spearman $\rho = -0.056$, $P = 0.63$; volatile: Spearman $\rho = -0.09$, $P = 0.46$; volatile – stable: Spearman $\rho = 0.04$, $P = 0.73$) and all parameters estimated in the additive choice model (α , stable: Spearman $\rho = 0.08$, $P = 0.51$; volatile: Spearman $\rho = 0.10$, $P = 0.38$; volatile – stable: Spearman $\rho = 0.19$, $P = 0.1$; ϕ , stable: Spearman $\rho = 0.20$, $P = 0.09$; volatile: Spearman $\rho = 0.01$, $P = 0.94$; volatile – stable: Spearman $\rho = -0.20$, $P = 0.09$; β , stable: Spearman $\rho = -0.04$, $P = 0.76$; volatile: Spearman $\rho = -0.07$, $P = 0.53$; volatile – stable: Spearman $\rho = 0.04$, $P = 0.75$). In the stable environment, where uncertainty and volatility are low, average happiness did not correlate with depressive symptoms (Spearman $\rho = 0.07$, $P = 0.58$; figure 6A, left panel). In the volatile environment, where uncertainty is high and volatility is high, average happiness was correlated with depressive symptoms, with lower happiness associated with higher depressive symptoms (Spearman $\rho = -0.23$, $P < 0.05$; figure 6A, central panel). This result is consistent with the finding that depressive symptoms are associated with reduced happiness in a context of risk decision-making²⁹. Finally, the difference between average happiness between volatile and stable environments was also correlated with depressive symptoms (volatile – stable: Spearman $\rho = -0.28$, $P < 0.05$; figure 6A, right panel). Baseline mood parameters

estimated in our happiness model showed the same relationship to depressive symptoms (stable: Spearman $\rho = -0.07$, $P = 0.54$; volatile: Spearman $\rho = -0.35$, $P < 0.01$; volatile – stable: Spearman $\rho = -0.34$, $P < 0.01$; see figure 6B). No other happiness model parameters were correlated with depressive symptoms (w_P , stable: Spearman $\rho = 0.20$, $P = 0.09$; volatile: Spearman $\rho = 0.07$, $P = 0.55$; volatile – stable: Spearman $\rho = -0.08$, $P = 0.52$; w_{PPE} , stable: Spearman $\rho = 0.09$, $P = 0.45$; volatile: Spearman $\rho = 0.07$, $P = 0.55$; volatile – stable: Spearman $\rho = -0.06$, $P = 0.61$; γ , stable: Spearman $\rho = 0.06$, $p = 0.59$; volatile: Spearman $\rho = -0.044$, $P = 0.71$; volatile – stable: Spearman $\rho = -0.10$, $P = 0.38$).

Figure 6. Baseline mood decreases with depressive symptoms in volatile environments

(A) Average happiness during the task was not correlated with depressive symptoms (PHQ) in the stable environment (left panel, blue) but decreased with depressive symptoms in the volatile environment (middle panel, orange). The difference in happiness between stable and volatile environments was also significantly related to depression (right panel).

(B) Baseline mood parameters showed the same relationship to depressive symptoms as average happiness with lower parameters in volatile than stable environments. * $P < 0.05$, ** $P < 0.01$

Discussion

We found that subjects tracked outcome probabilities and made decisions by integrating both learned probability and explicit reward magnitudes. Learning rates adapted to environmental volatility, with a higher learning rate in the more volatile environment consistent with previous studies^{2,5,6}. That behaviour was consistent with an additive choice model^{4,6,30,32} is consistent with recent empirical evidence^{33,35} suggesting that humans and non-human primates adopt a

377 multiplicative strategy under risk when probabilities are known, but both species adopt an
378 additive strategy under uncertainty when probabilities must be learned.

379 Mood dynamics were more closely related to learning-relevant variables than choice-relevant
380 variables. Mood was sensitive to the combined influence of past chosen subjective probabilities
381 and past PPEs. Parameters for probability, PPE, and forgetting factors estimated from happiness
382 ratings were correlated across stable and volatile environments. Finally, we found that although
383 choice accuracy and choice model parameters were not affected by depressive symptoms when
384 changes between safe and volatile environments are signalled, the decrease in happiness
385 observed in the volatile environment was correlated with symptom severity. The same pattern
386 was captured by the baseline mood parameter in the happiness model. Experiencing a stable
387 environment with low uncertainty and volatility could attenuate the expression of depressive
388 symptoms on mood. Risky decision tasks used in previous studies^{10,11,29} maximize irreducible
389 uncertainty (i.e., risky options had a 50% probability of each option), which is more comparable
390 to the volatile environment in the current study and may explain the previous finding of a link
391 between baseline mood parameters and depressive symptoms. Computational models that
392 capture ecologically relevant learning and decision processes may provide a critical advantage
393 for understanding the mechanisms that underlie psychiatric symptoms²⁶. Our findings suggest
394 that subjective feelings measured during tasks in depression-relevant domains may provide
395 additional information not captured by computational models of learning and decision-making.
396 One reason depression might reduce mood more in volatile than stable environments could be
397 an increase in the number of negative prediction errors experienced. Misestimation of the level
398 of uncertainty may also lead to a tendency for negative events to disproportionately affect
399 depressed individuals, and this uncertainty misestimation is believed to contribute to depression
400 and anxiety³⁴. That the learning rate difference between the volatile and the stable environment
401 did not correlate with anxiety symptom severity (STAI) is consistent with previous findings⁵
402 that anxiety is linked to learning deficits for aversive but not appetitive outcomes. In the
403 aversive domain, anxiety severity might be associated with differences in behavioural
404 adaptation to volatility changes as well as mood.

405 Previous studies using risky decision tasks where reward and probability were explicitly
406 represented^{10,11} showed that mood dynamics were explained by past expected values and RPEs.
407 The present design allows us to dissociate the impact of learning-relevant and learning-
408 irrelevant information for mood. Our results suggest that when goal attainment requires
409 adaptive behaviour, mood dynamics are selective for learning-relevant information. Consistent
410 with the results obtained from risky decision tasks used in previous studies^{10,11,29}, potential
411 rewards were a key determinant of behaviour in our task. However, rewards were not a
412 determinant of mood in the current study, in contrast to previous results in risky decision
413 tasks^{10,11,29}. This result might imply a role for mood in learning in line with influential
414 proposals^{12,13}. However, mood did not reflect all learning-relevant information that influenced
415 behaviour. Happiness forgetting factors were highly correlated across environments and acted
416 more as a stable trait that differed between individuals and did not adjust to environmental
417 volatility. What distinguishes individuals with high or low happiness forgetting factors is
418 unknown. Overall, our findings show that mood dynamics are sensitive to depressive symptoms
419 and reflect variables relevant to adaptive behaviour irrespective of environmental volatility.

Methods

Subjects

Seventy-five healthy subjects (age range 18-35, 24 males) took part in the experiment. Thirsty-seven participated in the stable environment first and volatile environment second. Subjects were paid £10 for their participation. All subjects gave informed consent and the Research Ethics Committee of University College London approved the study.

Procedure

Participants were first instructed about the task. They performed 20 practice trials before a test ensuring that they understood that both probability and magnitude mattered to maximize the number of points obtained. They were told that they would be exposed to two environments: an environment where one car is more likely to win for the entire session, and an environment where which car is more likely to win changes occasionally. Because we wanted to maximise the efficiency of the behavioural manipulation (i.e., the difference in learning rate between environments observed in previous studies^{2,5,6}) to study how mood dynamics varied with behavioural sensitivity, we explicitly signalled the environment by using different pairs of cars in the two environments. Moreover, before each condition, an instruction screen explained in which environment participants will be placed. Finally, a fixation symbol displayed in the centre of the screen in each trial was specific for each environment: ‘-’ for the stable environment and ‘~’ for the volatile environment. Participants were given no guidance as to how they should use information about environmental volatility. After completion of the task, participants completed three standard clinical questionnaires: Beck Depression Inventory (BDI-II)³⁶, Patient Health Questionnaire (PHQ-9)³⁷, and the State/Trait Anxiety Inventory (STAI)³⁸.

Experimental task

The task was implemented using the Cogent toolbox in MATLAB (MathWorks, Inc.). Subjects had to choose between two cars, each associated with a probability (20% or 80%) of winning. If the chosen car won, participants earned the corresponding amount of points. In the stable environment (80 trials), the probability to win for the best car was 80%. In the volatile environment (80 trials), reward probabilities switched between 80% for one car and 20% for the other car every 20 trials. The order was counterbalanced between subjects ($n = 38$ in stable-volatile order, $n = 37$ in volatile-stable order). The outcomes were locally pseudo-randomised, to ensure that every 10 trials (i.e., trials 1-10, 11-20), the car with the highest outcome probability won on exactly 8 of 10 trials. The possible pairs of rewards were 10-10, 10-40, 10-60, 10-80, 20-40, 40-10, 40-20, 40-40, 60-10, 80-10. Subjects were primed when the side of the screen for each car was swapped (every six to ten trials) with an explicit cue. They were also instructed that the car location was unrelated to the outcome probability and to the change in outcome probabilities in the volatile environment. Every three to four trials, subjects were asked to indicate “How happy are you right now?” on a scale from very unhappy to very happy. They were told to consider these extremes within the context of the experiment. Each trial started with a fixation screen for a duration varying between 0.9 and 1.9 s. Cars were displayed for 1.2 s without any information about reward magnitudes, and no choice was allowed in this phase. Subjects were free to choose the option they preferred without any time constraints as soon as the potential reward for each car was displayed. The chosen option was surrounded by a yellow frame for 1.5 s. Finally, the outcome was displayed for 2 s. Both the car and the reward magnitude frames were green if the chosen car won. They were red and the car was crossed out if the chosen car lost.

Data analyses

One-sample two-tailed t-tests were used to compare performance, proportion of win-stay/lose-shift, and model parameters between environments at the group level. Spearman rank correlations were used to test for relationships between parameters and to relate depression scores to behavioural and happiness parameters. All analyses were performed using MATLAB.

Computational models

All models were fitted to experimental data by minimizing the negative log likelihood of the predicted choice probability given different model parameters using the `fmincon` function in MATLAB (Mathworks Inc.). We used standard model comparison techniques^{39,40} to compare model fits. For each model fit in individual subjects, we computed the Bayesian Information Criterion (BIC), which penalizes for model complexity (i.e., number of parameters), and then summed BIC across subjects. The model with the lowest BIC is the preferred model. For all choice models, the learning rate was bounded between 0 and 1, the inverse temperature between 0 and 50 (to avoid ceiling effects), the probability-magnitude relative weight ϕ between 0 and 1, and the gamma risk aversion parameters between 0 and 10. For the happiness models, we first fitted each happiness model on both environments using the same parameters for both environments. Then each model was fitted for each environment separately, with the starting parameters determined by the joint model fit under standard constraints (the forgetting factor and the baseline mood parameter could only vary between 0 and 100).

Acknowledgments

We thank Sankalp Garud, Matilde Vaghi, Benjamin Chew, Paul Sharp, and Rachel Bedder for their help. R.B.R. is supported by a Medical Research Council Career Development Award (MR/N02401X/1) and a 2018 NARSAD Young Investigator Grant (27674) from the Brain and Behavior Research Foundation, P&S Fund. The Max Planck UCL Centre is a joint initiative supported by UCL and the Max Planck Society. The Wellcome Centre for Human Neuroimaging is supported by core funding from the Wellcome Trust (203147/Z/16/Z).

References

1. Barto, A. G. Adaptive critics and the basal ganglia. in *Models of Information Processing in the Basal Ganglia* 215–232 (The MIT Press, 1995).
2. Behrens, T. E. J., Woolrich, M. W., Walton, M. E. & Rushworth, M. F. S. Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–1221 (2007).
3. Mathys, C., Daunizeau, J., Friston, K. J. & Stephan, K. E. A bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* **5**, 39 (2011).
4. Donahue, C. H. & Lee, D. Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. *Nat. Neurosci.* **18**, 295–301 (2015).
5. Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X. & Bishop, S. J. Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nat. Neurosci.* **18**, 590–596 (2015).
6. Massi, B., Donahue, C. H. & Lee, D. Volatility facilitates value updating in the prefrontal cortex. *Neuron* **99**, 598–608.e4 (2018).
7. Fredrickson, B. L. The broaden-and-build theory of positive emotions. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **359**, 1367–1377 (2004).
8. Mellers, B. A., Schwartz, A., Ho, K. & Ritov, I. Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychol. Sci.* **8**, 423–429 (1997).
9. Shepperd, J. A. & McNulty, J. K. The affective consequences of expected and unexpected outcomes. *Psychol. Sci.* **13**, 85–88 (2002).

10. Rutledge, R. B., Skandali, N., Dayan, P. & Dolan, R. J. A computational and neural model of momentary subjective well-being. *Proc. Natl. Acad. Sci.* **111**, 12252–12257 (2014).
11. Rutledge, R. B., Skandali, N., Dayan, P. & Dolan, R. J. Dopaminergic modulation of decision making and subjective well-being. *J. Neurosci.* **35**, 9811–9822 (2015).
12. Eldar, E. & Niv, Y. Interaction between emotional state and learning underlies mood instability. *Nat. Commun.* **6**, (2015).
13. Eldar, E., Rutledge, R. B., Dolan, R. J. & Niv, Y. Mood as representation of momentum. *Trends Cogn. Sci.* **20**, 15–24 (2016).
14. Blanco, N. J., Otto, A. R., Maddox, W. T., Beevers, C. G. & Love, B. C. The influence of depression symptoms on exploratory decision-making. *Cognition* **129**, 563–568 (2013).
15. Chase, H. W. *et al.* Approach and avoidance learning in patients with major depression and healthy controls: relation to anhedonia. *Psychol. Med.* **40**, 433–440 (2010).
16. Cella, M., Dymond, S. & Cooper, A. Impaired flexible decision-making in major depressive disorder. *J. Affect. Disord.* **124**, 207–210 (2010).
17. Herzallah, M. M. *et al.* Learning from negative feedback in patients with major depressive disorder is attenuated by SSRI antidepressants. *Front. Integr. Neurosci.* **7**, 67 (2013).
18. Kunisato, Y. *et al.* Effects of depression on reward-based decision making and variability of action in probabilistic learning. *J. Behav. Ther. Exp. Psychiatry* **43**, 1088–1094 (2012).
19. Mueller, E. M., Pechtel, P., Cohen, A. L., Douglas, S. R. & Pizzagalli, D. A. Potentiated processing of negative feedback in depression is attenuated by anhedonia. *Depress. Anxiety* **32**, 296–305 (2015).
20. Pechtel, P., Dutra, S. J., Goetz, E. L. & Pizzagalli, D. A. Blunted reward responsiveness in remitted depression. *J. Psychiatr. Res.* **47**, 1864–1869 (2013).
21. Taylor Tavares, J. V. *et al.* Neural basis of abnormal response to negative feedback in unmedicated mood disorders. *NeuroImage* **42**, 1118–1126 (2008).
22. Thoma, P., Norra, C., Juckel, G., Suchan, B. & Bellebaum, C. Performance monitoring and empathy during active and observational learning in patients with major depression. *Biol. Psychol.* **109**, 222–231 (2015).
23. Vrieze, E. *et al.* Reduced reward learning predicts outcome in major depressive disorder. *Biol. Psychiatry* **73**, 639–645 (2013).
24. Robinson, O. J., Cools, R., Carlisi, C. O., Sahakian, B. J. & Drevets, W. C. Ventral striatum response during reward and punishment reversal learning in unmedicated major depressive disorder. *Am. J. Psychiatry* **169**, 152–159 (2012).
25. Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A. & Daw, N. D. Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *eLife* **5**, e11305 (2016).
26. Scholl, J. & Klein-Flügge, M. Understanding psychiatric disorder by capturing ecologically relevant features of learning and decision-making. *Behav. Brain Res.* **355**, 56–75 (2018).
27. Huys, Q. J., Pizzagalli, D. A., Bogdan, R. & Dayan, P. Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biol. Mood Anxiety Disord.* **3**, 12 (2013).
28. de Berker, A. O. *et al.* Computations of uncertainty mediate acute stress responses in humans. *Nat. Commun.* **7**, 10996 (2016).
29. Rutledge, R. B. *et al.* Association of neural and emotional impacts of reward prediction errors with major depression. *JAMA Psychiatry* **74**, 790–797 (2017).
30. Farashahi, S. *et al.* Metaplasticity as a neural substrate for adaptive learning and choice under uncertainty. *Neuron* **94**, 401–414.e6 (2017).
31. Kahneman, D. & Tversky, A. Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 363–391 (1979).
32. Rouault, M., Drugowitsch, J. & Koehlin, E. Prefrontal mechanisms combining rewards and beliefs in human decision-making. *Nat. Commun.* **10**, 1–16 (2019).
33. Farashahi, S., Donahue, C. H., Hayden, B. Y., Lee, D. & Soltani, A. Flexible combination of reward information across primates. *Nat. Hum. Behav.* **3**, 1215–1224 (2019).
34. Pulcu, E. & Browning, M. The misestimation of uncertainty in affective disorders. *Trends Cogn. Sci.* **23**, 865–875 (2019).
35. Koehlin, E. Human decision-making beyond the rational decision theory. *Trends Cogn. Sci.* (2019).

36. Beck, A. T., Steer, R. A., Ball, R. & Ranieri, W. F. Comparison of beck depression inventories-ia and-ii in psychiatric outpatients. *J. Pers. Assess.* **67**, 588–597 (1996).
37. Kroenke, K., Spitzer, R. L. & Williams, J. B. The PHQ-9: validity of a brief depression severity measure. *J. Gen. Intern. Med.* **16**, 606–613 (2001).
38. Spielberger, C. D. State-trait anxiety inventory for adults. (1983) doi:10.1037/t06496-000.
39. Schwarz, G. Estimating the dimension of a model. *Ann. Stat.* **6**, 461–464 (1978).
40. Burnham, K. P. & Anderson, D. R. *Model Selection and Multi-model Inference: Practical Information-theoretic Approach*. (Springer, 2004).

Supplementary figure 1

Happiness is influenced by multiple past probability prediction errors

Each bar corresponds to the influence of the past trials on the current happiness rating in the stable (blue) and volatile (orange) environments. (A) The left panel shows the influence of probability prediction error estimated from the additive learning model ($\hat{PPE}$). (B) The right panel shows the influence of the objective probability prediction error (PPE). Error bars represent SEM. *P < 0.05, **P < 0.01, ***P < 0.001

Supplementary method

To estimate the influence of the past trials on the current happiness without any assumption regarding the shape of the influence decay, we fitted a general linear model including for each rating the previous ten probability prediction errors using the Matlab glmfit function. Each value was then tested against 0 at the group level using one-sample two-tailed t-tests.

Parameter	STAI			PHQ		
	Stable	Volatile	Volatile - Stable	Stable	Volatile	Volatile - Stable
α	0.07	0.09	0.11	0.08	0.10	0.19 [†]
β	0.08	-0.08	-0.12	-0.04	-0.07	0.04
ϕ	0.11	-0.03	-0.14	0.20 [†]	0.01	-0.20 [†]

601 **Supplementary table 1. Choice parameters are unrelated to symptom severity**
602
603 State Trait Anxiety Inventory (STAI) and Patient Health Questionnaire (PHQ) scales were
604 collected for all subjects. Numbers are Spearman rank correlation coefficients. [†]P <0.1