SUPPLEMENTARY INFORMATION

Distinct Computational Mechanisms of Uncertainty Processing Explain Opposing Exploratory Behaviors in Anxiety and Apathy

Yan *et al.*

Contents

Supplement Tables

Supplement Text

References

Supplement Methods

Method S1. Three-armed restless bandit task

Participants were free to choose between three targets for the potential to earn a reward of 1 point. Each target is associated with a hidden reward probability that randomly and independently changes throughout the task. We seeded each participant's reward probability randomly to prevent biases due to particular kinds of environments. Specifically, on each correct trial, there was a 67% chance that the reward probability for each target would either increase or decrease by 0.2, with these probabilities bounded between 0 and 1.0. Due to the variable and independent nature of the rewards, participants could only estimate the probabilities by actively sampling from the targets and accumulating their reward experiences over time.

Method S2. Hidden Markov Model

In current study, the observed choices (y) are "emissions" that are generated by an unobserved decision process that is in some latent, hidden state (z).
Latent states are defined by both the probability of making each choice (k, out of N_k possible options), and by the probability of transitioning from each state to every other state.
Our model consisted of two types of states, an *explore* state and the *exploit* states. The emissions model for the explore state was the maximum-entropy distribution for a categorical variable, a uniform distribution:

$$p(y_t = k|z_t = explore) = \frac{1}{N_k}$$

Where *N* is the number of stimuli that were presented (i.e. N = 3),
and *t* is the trial number.
Because exploitation involves repeated sampling of each option, exploit states only permitted choice emissions that matched one option. That is:

$$p(y_t = k|z_t = exploit_i, k = i) = 1$$

$$p(y_t = k|z_t = exploit_i, k \neq i) = 1$$

The latent states in this model are Markovian, meaning that the current state (z_t) depends only on the most recent state (z_{t-1}):

$$p(z_t|z_{t-1}, y_{t-1}, \dots, z_1, y_1) = p(z_t|z_{t-1})$$

This means that we can describe the entire pattern of dynamics in terms of a time-invariant transition matrix between **4 possible states** (three exploit states and one explore state). This matrix is a system of stochastic equations that describe the one-time-step probability of transitioning between every combination of past and future states (i, j).

$p(z_t = i|z_{t-1} = j)$

Because we had only a limited number of trials for each participant (300), parameters were tied across exploit states: each exploit state had the same probability of beginning (from exploring) and of sustaining itself. For conceptual reasons, the model also assumed that participants started in exploration and had to pass through exploration in order to start exploiting a new option, even if only for a single trial (Chen et al., 2021; Ebitz et al., 2018, 2019, 2020).
We have previously shown that models that lack these constraints by design tend to approximate them when fit to sufficiently large datasets (Chen et al., 2021; Ebitz et al., 2018).Because the emissions model was fixed, certain parameters were tied, the structure of the transmission matrix was constrained, and the initial state was specified, the final HMM had only two free parameters: one corresponding to the probability of exploring, given exploration on the last trial, and one corresponding to the probability of exploiting, given exploitation on the last trial. We have previously reported that this constrained model does not underperform an unconstrained model (Chen et al., 2021; Ebitz et al., 2018). And that unconstrained models tend to closely resemble to the constrained model when fit to large amounts of data (Ebitz et al., 2018).

The HMM was fit via expectation-maximization using the Baum Welch algorithm (Bilmes, 1998). This algorithm finds a (possibly local) maxima of the complete-data likelihood. Because the participants had no knowledge of the environment at the first trial, we assumed they began by exploring, rather than adding another parameter to the model here. The algorithm was reinitialized with random seeds 20 times, and the model that maximized the observed (incomplete) data log-likelihood across all the sessions for each participant was ultimately taken as the best. To decode latent states from choices, we used the Viterbi algorithm (Forney, 1973) to discover the most probable a posteriori sequence of latent states.

Method S3. Kalman filter

In the Kalman filter model for a multi-armed bandit task, *process noise* and *observation noise* refer to two distinct sources of uncertainty that affect the learning and decision-making process.

Process noise represents the uncertainty in the evolution of the hidden state (reward mean) over time. It accounts for how the true state evolves from one point in time to the next. In mathematical terms, process noise is part of the state transition equation in the Kalman Filter:

$$x_t = x_{t-1} + e_t$$

x_t is the state at time t

e_t is the process noise t, which is assumed to be drawn from a normal distribution with zeros mean and process noise variance v. Where the $e_t \sim N(0, v)$.

The process noise captures the idea that the reward-means for each arm can change from one trial to the next, even in the absence of any observations. A higher process noise variance v indicates a more volatile environment, where the reward means are expected to change more rapidly.

In contrast, observation noise represents the uncertainty in the observed rewards, given the current hidden state (reward mean). Which is assumed to be Gaussian with zero mean and a fixed variance σ^2.

The observation noise captures the idea that the observed rewards are noisy and can deviate from the true reward mean due to random fluctuations or measurement errors.

A higher measurement noise variance indicates a more stochastic environment, where the observed rewards are less reliable and informative about the underlying reward means.

The Kalman Filter operates optimally when the statistical properties of the process noise and the measurement noise are accurately known.

When observation noise variance (σ^2) is high relative to the process noise variance (v), the Kalman gain will be small, and the model will rely more on its prior beliefs and less on noisy observations. Conversely, when the observation noise variance (v), is high relative to the process noise variance (σ^2), the Kalman gain will be large, and the model will update its beliefs more strongly based on the observed rewards.

Method S4. Volatile kalman filter for three-armed bandit task

The key difference between a standard Kalman filter and a volatile Kalman filter (VKF) is the variance of the process noise, a stochastic variable that changes with time. In other words, the VKF introduces parameters to handle the volatility in the process noise. Specifically, it allows the process noise variance v to vary with the observed prediction errors, reflecting changes in environmental volatility.

Our approach here is essentially the same as that taken by Piray and Daw (Piray & Daw, 2020). Here, we briefly described the model details as follows.

Kalman gain:

$$k_t = (w_t + v_{t-1})/(w_{t-1} + v_{t-1} + \sigma^2)$$

Update for the reward means:

$$m_t = m_{t-1} + k_t(O_t - m_{t-1})$$

Update for posterior variance w_t:

$$w_t = (1 - k_t)(w_{t-1} + v_{t-1})w_{t-1,t} = (1 - k_t)w_{t-1}$$

Update for volatility:

$$v_t = v_{t-1} + \lambda((m_t - m_{t-1})^2 + w_{t-1} + w_t - 2w_{t-1,t} - v_{t-1})$$

Method S5. Rescorla-Wagner models

We also fitted the data to the classical Rescorla-Wagner model. Successful adaptation in a dynamic situation requires the appropriate feedback-based learning process where individuals integrate the feedback (reward or non-reward) into the stimulus-outcome association (Forstmann et al., 2016). The basic reinforcement learning model, the Rescorla-Wagner model can address this process well. So the first model (RW1) was defined as:

$$v_t = v_{t-1} + a \times (R_{t-1} - v_{t-1})$$

where v_t is the value of the option on trial *t*.
a represents the general learning rate from feedback.

To verify whether participants employed distinct or shared computational responses to positive and negative feedback, we built another model with two learning rates, one for positive feedback and the other for negative feedback (den Ouden et al., 2013). This model (RW2) can be defined as:

$$v_t = v_{t-1} + \alpha^{pos} \times (R_{t-1} - v_{t-1}), positive\ feedback$$
$$v_t = v_{t-1} + \alpha^{neg} \times (R_{t-1} - v_{t-1}), negative\ feedback$$

Where v_t is the value of the option on trial *t*. $\alpha^{pos}$and α^{neg} represent the learning rates from positive and negative feedback, respectively.
For these two models, $R_{t-1} \in \{0,1\}$ represents the feedback received in response to participants' choice on trial *t*-1. And $R_{t-1} - v_{t-1}$ represents prediction error in trial *t*-1.

We used a softmax choice function to map the value into choice. The softmax function for these four models can be defined as:

$$P^t = \frac{exp^{(\beta V_{,t})}}{exp^{(\beta V_{t,1} + \beta V_{t,2} + \beta V_{t,3})}}$$

Where the β represents the inverse temperature with choice value.

Method S6. Model fitting and comparison

Hierarchical Bayesian inference (HBI) is a powerful method for model fitting and comparison in group studies (Piray et al., 2019). Unlike traditional approaches such as maximum likelihood estimation (MLE) or maximum a posteriori (MAP) estimation, which fit models to each subject independently, HBI simultaneously fits models to all subjects while constraining individual fits based on group-level statistics (i.e., empirical priors). This approach yields more robust and reliable parameter estimates, particularly when individual subject data is noisy or limited.
In our study, we employed HBI to fit models to choice data. The method quantifies group-level mean parameters and their corresponding hierarchical errors. To ensure that parameter estimates remain within appropriate bounds during the fitting process, we used the sigmoid function to transform parameters bounded in the unit range or with an upper bound and the exponential function to transform parameters bounded to positive values. The initial parameters of all models were obtained using a MAP procedure, with the initial prior mean and variance for

all parameters set to 0 and 6.25, respectively, based on previous research (Piray & Daw, 2020). This initial variance allows parameters to vary widely without substantial influence from the prior. For model comparison, we used Bayesian model selection(Stephan et al., 2009), specifically employing the protected exceedance probability (PXP) to select the winning model. The PXP quantifies the probability that a given model is more frequent in the population than all other models under consideration while accounting for the possibility that the observed differences in model evidence may be due to chance. The model with the highest PXP is selected as the winning model. The detailed results of our model comparison, including PXP and BIC values for all models, can be found in Table S5.

The better model performance further confirmed that the Kalman filter's formulation aligns well with theories of how the individuals might perform inference and learning under uncertainty, making it particularly suitable for our study of affective influences on these processes. And the KF model has the ability to dissociate uncertainty, which allows us to separately estimate volatility (process noise variance) and stochasticity (observation noise variance). This distinction is crucial for our research questions about how anxiety and apathy influence perceptions of different types of uncertainty

Method S7. Parameter recovery for HMM

First, we simulated a dataset with 100 subjects and repeated this process 50 times. Wilson and Collins (2019)(Wilson & Collins, 2019) proposed to adjust the input values of simulations to empirical obtained behavioral results. Therefore, we randomly selected 100 sets of transition probabilities from our empirically fitted parameters and added small amounts of random noise (10% of the parameter's standard deviation) to create true parameter values for simulation. This approach ensured that our simulated parameters maintained realistic distributions while introducing some variability. For each parameter set, we generated synthetic choice sequences (300 trials each) using the true parameters, then we applied our HMM fitting procedure to recover the parameters from these sequences, and finally, we compared the recovered parameters to the true generating parameters. The recovery rates were robust, we found the mean recovery rate for P(explore | explore)= 0.644, and the mean recovery rate for P(exploit | exploit)= 0.793.

Method S8. Model validation

We validated our modeling procedure using two approaches. First, we assessed parameter recovery by refitting data simulated from the winning model and comparing the resulting parameter estimates to their ground truth. We simulated 50 agents' choices and observations, repeating this process 50 times. Data were generated based on the Kalman filter (Eqs 1-3). Specifically, the observation on each trial was generated by random walk: there was a 67% chance that the reward probability for each target would either increase or decrease by 0.2, with these probabilities bounded between 0 and 1.0. The choice data were generated randomly by applying the softmax as the response model with parameter β. Following the standard procedure in original Kalman Filter model shared by Piray & Daw (Piray & Daw, 2020), we conducted parameter recovery analyses using synthetic data generated from the Kalman Filter model (Eqs 1-3). For each simulation, we generated data for 100 agents, with each subject completing three sequences of 300 trials (3 different cues). We obtained reasonable parameter recovery correlations. The mean Pearson correlations were 0.707, 0.671, 0.973 for v, σ^2, and β, respectively (Table S6).

Secondly, we tested the accuracy of the model prediction. We calculated the correlation between behavioral output predicted by model and real choices. We conducted this analysis with a randomly selected subset of 50 participants from the full dataset of 1001 and demonstrated strong correlations between observed behaviors and model predictions across three cues: r = 0.719 for cue1, r = 0.730 for cue2, and r = 0.775 for cue3. All correlations were statistically significant, $p < 0.0001$; see Figure S1)

Method S9. Split-half reliability

To assess the split-half reliability of our task, we examined the consistency of overall choices and model parameters from the winning model between the first and second halves of trials. For overall choice proportion, we employed Pearson's correlations to calculate reliability. For model parameters, we utilized a more sophisticated approach, calculating model-derived estimates of Pearson's r values from the parameter covariance matrix. This method, which estimates first- and second-half parameters within a single model, has recently been validated for accurate parameter reliability estimation (Waltmann et al., 2022). We interpreted indices of reliability based on conventional values of <0.40 as poor, 0.4–0.6 as fair, 0.6–0.75 as good, and >0.75 as excellent reliability (Fleiss, 2011). Overall choice proportion showed fair-to-good reliability (r=0.65, r=0.64, r=0.53 for v, σ^2, and β, respectively, see Figure S2). The model parameters showed good-to-excellent reliability (r = [0.79, 0.78, 0.69] after Spearman-Brown correction).

Method S10. Mediation analyses

Mediation analysis is a statistical method used to examine the underlying mechanisms by which an independent variable influences a dependent variable through one or more mediator variables (Hayes, 2017). In our study, we employed the bootstrapping method to estimate the mediation effect of volatility and stochasticity on the relationship between affective states (apathy and anxiety) and exploration. Bootstrapping is a nonparametric approach to effect-size estimation and hypothesis testing that is increasingly recommended for many types of analyses, including mediation (Mackinnon et al., 2004). This method involves repeatedly resampling from the available data to generate an empirical approximation of the sampling distribution of the indirect effect (i.e., the effect of the independent variable on the dependent variable through the mediator). We used this distribution to calculate p-values and construct confidence intervals based on 5,000 resamples. Bootstrapping is preferred over other methods, such as the Sobel test because it does not assume the normality of the sampling distribution and provides more accurate confidence intervals that are bias-corrected and accelerated (Hayes, 2017; Mackinnon et al., 2004). This approach offers a robust and powerful way to test mediation hypotheses, particularly in cases where the sample size is relatively small or the data violate assumptions of normality (Preacher & Hayes, 2008).

Method S11. Dimensionality reduction method

Popular and valid dimensionality reduction techniques to reveal manifolds include t-distributed stochastic neighborhood embedding (t-SNE) (Hinton & Roweis, 2002), uniform manifold approximation and projection (UMAP) (McInnes et al., 2018), and Principal component analysis (PCA) (Jolliffe & Cadima, 2016). However, t-SNE suffers from limitations, including slow computation time and loss of global data structure, and it is not a deterministic algorithm (Rhys, 2020). The main drawback of PCA is that it is highly affected by outliers in the dataset (Jolliffe & Cadima, 2016). In contrast, UMAP is a deterministic and efficient algorithm, it also preserves both local and global structure of original high-dimensional data. UMAP was implemented in the R language. The eight-dimensional datasets from all participants were passed into the R package *umap*, version 0.2.8.0, available at https://cran.r-project.org/web/packages/umap/) with default parameter setting as n_component = 2, n_neighbors = 15, min_dist = 0.3, metric = ‘Euclidean’. For reproducibility reasons, we fixed the random_state in this algorithm. The

hyperparameter *n_neighbors* decide the radius of the search region. Larger values will include more neighbors, thus forcing this algorithm to consider more global structure of original n-dimension data. Another important hyper-parameter, *min_dist* determines the allowed minimum distance apart for cases in lower-dimensional space. *metric* defines the way that UMAP is used to measure distances along the manifold.

Method S12. Turning point to divide the manifold into monotonically increasing and decreasing group

To divide the manifold into monotonically increasing and decreasing groups, we sorted the scores for dimension 1 in ascending order. Initially, we fitted a linear model using the first three data points located on the upper left of the manifold. We then expanded this model by sequentially including one additional data point from dimension 1, continuing this process until we incorporated the last score (i.e., the maximum dimension 1 score, situated on the upper right of the manifold). Throughout this procedure, we monitored the t-statistic of the dimension1 coefficient to assess the statistical significance of dimension1 as a predictor. Notably, a dimension 1 score of -0.671 marked the most significant negative coefficient, after which the relationship between dimension 1 and dimension2 gradually shifted to become positive (see Figure S11)

Method S13. False discovery rate correction

The Benjamini-Hochberg False Discovery Rate (FDR) procedure (Benjamini & Hochberg, 1995) with q = 0.05 was chosen for our study due to its superior performance in managing multiple comparisons while maintaining statistical power (Riffenburgh, 2014). This method is particularly well-suited for our research, which involves multiple correlations testing related hypotheses. The FDR procedure effectively balances the need to discover true effects while controlling false positives, making it more appropriate than traditional family-wise error rate controls such as the Bonferroni correction (Storey & Tibshirani, 2003). Unlike the Bonferroni method, which can be overly conservative and lead to an increased risk of Type II errors (false negatives), the FDR approach offers a better control of false discoveries (Glickman et al., 2014). Furthermore, by setting q = 0.05, we ensure that the expected proportion of false discoveries among all rejected null hypotheses is controlled at 5%, providing a reasonable balance between identifying true effects and limiting erroneous conclusions.

Method S14. Data quality check

We have attention checks include checking the consistency of forward and reverse scored survey responses and the face validity of direct questioning including “answer with the color of grass”. And participants must also meet a score threshold of 42% and an exploration threshold of 2 unique selections during the 25 practice bandit trials.

Moreover, we implemented a three-step screening procedure to ensure data quality in our questionnaire responses (Zorowitz et al., 2023). First, we calculated the inter-item standard deviation (ISD) across all items for each participant (7 items for GAD-7 and 18 items for AMI), noting that low ISD values were not considered problematic as they might reflect genuinely consistent anxiety or apathy levels (e.g., not anxious or apathetic at all). Second, we detected extreme alternation patterns by examining consecutive item responses (e.g., 3-0-3-0-3-0), calculating absolute differences between adjacent responses, and determining the proportion of extreme jumps (defined as differences ≥ 3 points). Finally, we established flagging criteria where responses were only considered suspicious if they met both conditions: an ISD > 2 (indicating very high response variation) AND extreme alternations in more than 60% of responses. Using this fine-grained approach to maintain data quality while respecting the clinical nature of the GAD-7 and AMI measures, we found that 0% of responses showed potentially problematic patterns.

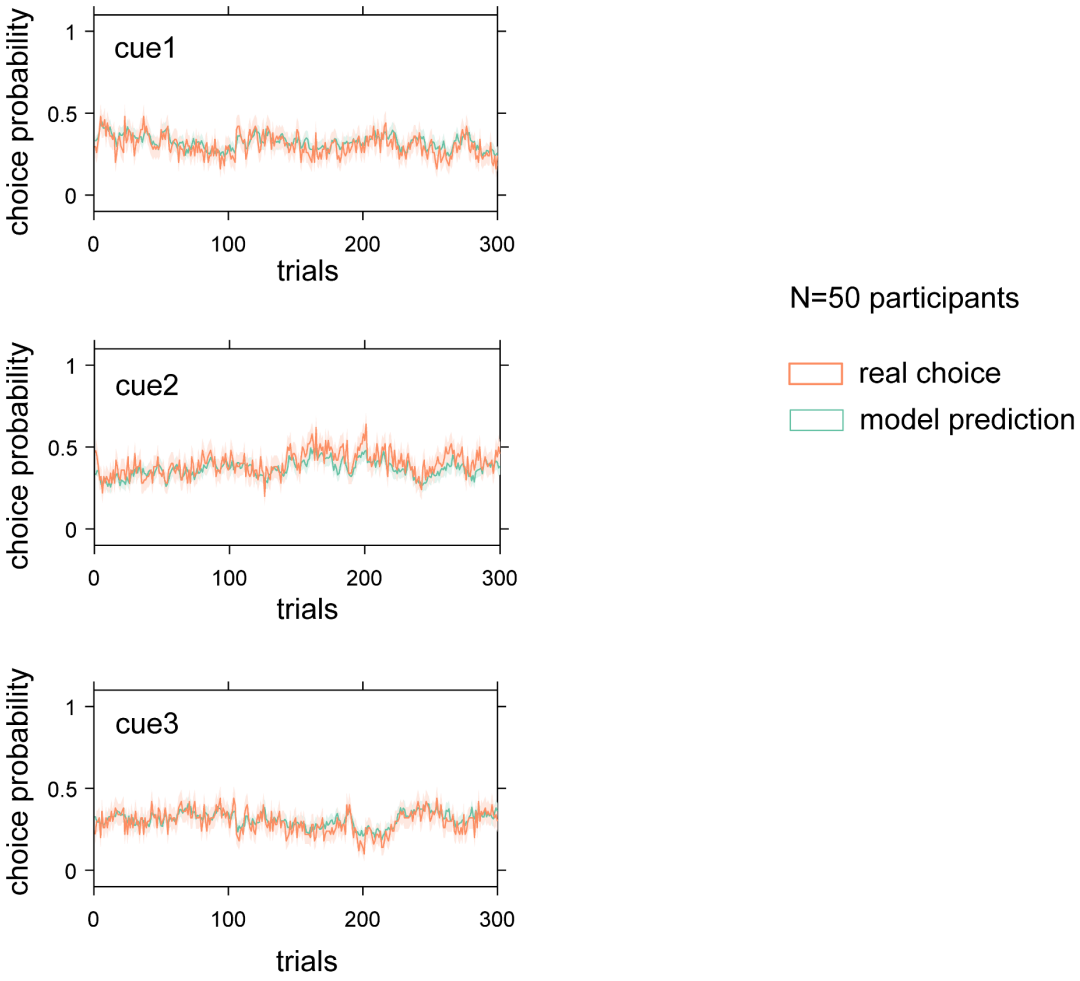

Figure S1. Posterior predictive check for model validity. This analysis was conducted with a randomly selected subset of 50 participants from the full dataset of 1001 and demonstrated strong correlations between observed behaviors and model predictions across three cues: r = 0.719 for cue1, r = 0.730 for cue2, and r = 0.775 for cue3. All correlations were statistically significant, $p < 0.0001$)

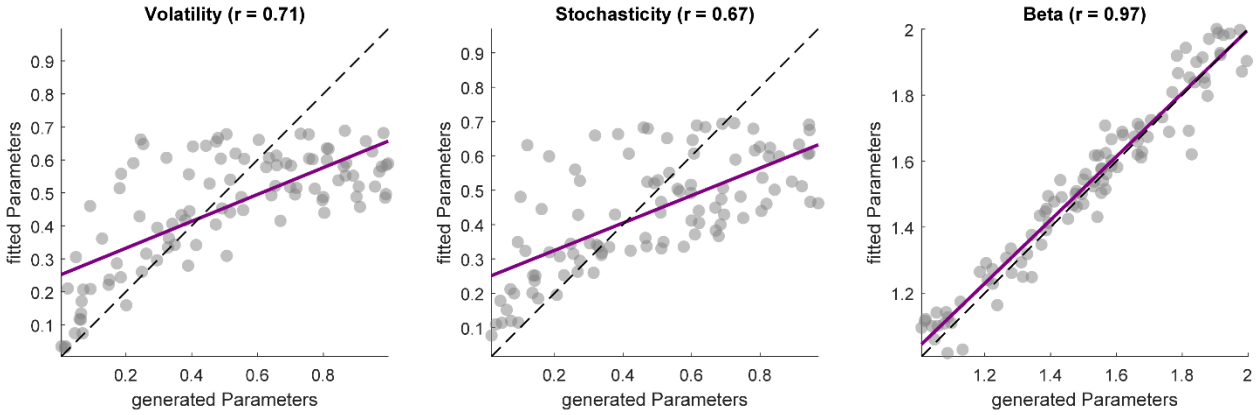

Figure S2. Parameter recovery. Following the standard procedure (Piray & Daw, 2020), we conducted parameter recovery analyses using synthetic data generated from the Kalman Filter model. For each simulation, we generated data for 100 agents, with each subject completing three sequences of 300 trials (3 different cues). We ran 50 simulations per agent and analyzed recovery using Pearson correlations between true parameters and averaged fitted parameters. We obtained reasonable parameter recovery correlations. Pearson correlations were for v = 0.707, σ^2= 0.671, and β = 0.973.

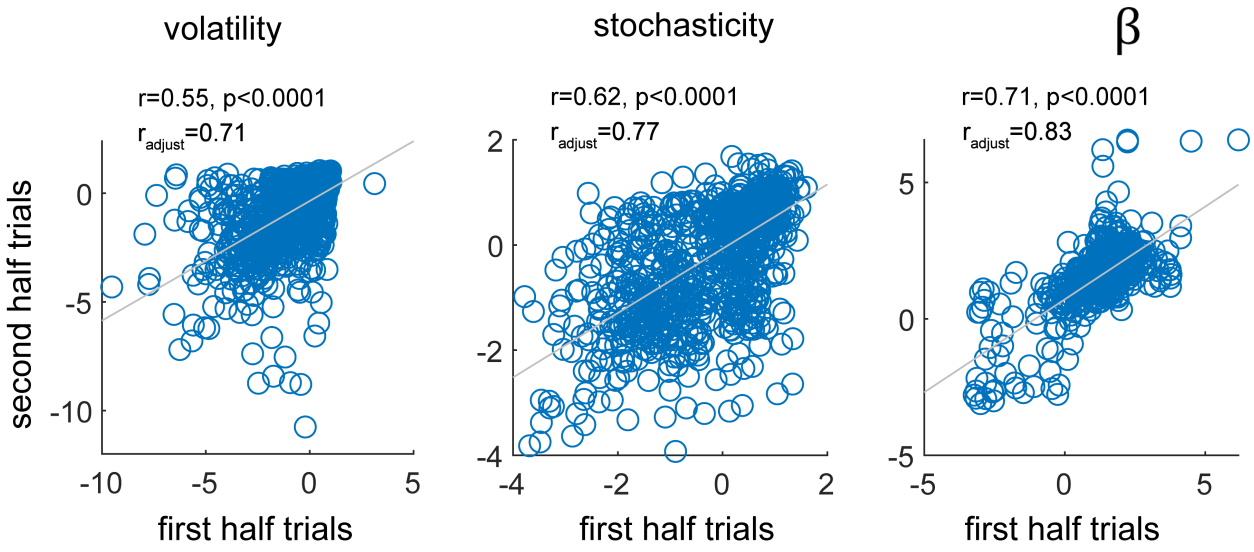

Figure S3. Split-half reliability of the task.

We assessed the reliability of our task measures using a split-half approach. The scatter plots comparing parameters (before transformation) from the first and second halves of the task are presented, along with their corresponding reliability estimates (Pearson's r values). Reliability estimates for the computational measures from the winning computational model were computed by fitting split-half parameters within a single model and then using the parameter covariance matrix to derive Pearson's correlation coefficients for each parameter across halves. Reliability estimates are reported as unadjusted values (r) and after adjusting for reduced number of trials via Spearman-Brown correction (r_{adjust}). Statistics reported here based on the correlation between transformed parameters between these two halves. Grey lines show lines-of-best-fit. Overall choice proportion showed fair-to-good reliability (r=0.55, r=0.62, r=0.71 for volatility, stochasticity, and β, respectively). The model parameters showed good-to-excellent reliability (r = [0.71, 0.77, 0.83] after Spearman-Brown correction).

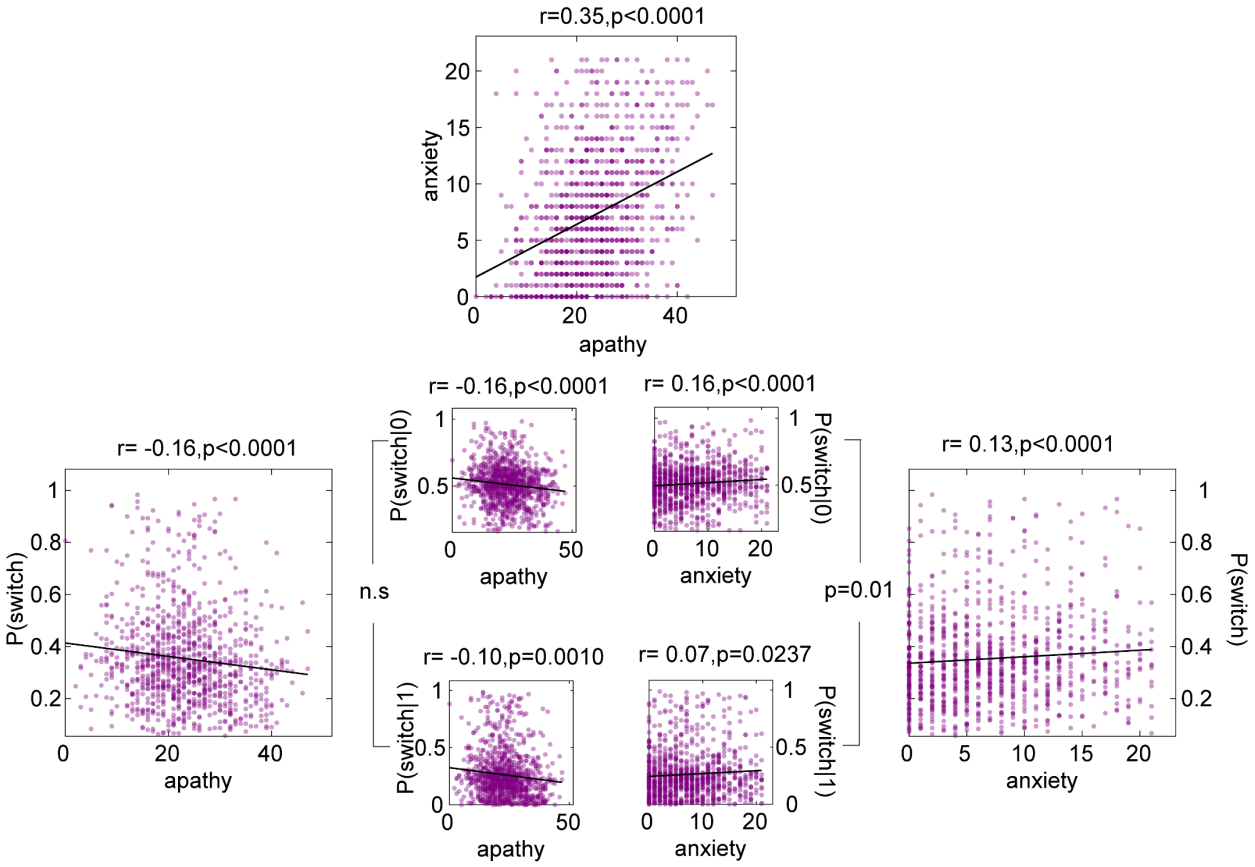

Figure S4. Distinct behavioral patterns associated with apathy and anxiety (plots with all data points). (Top panel). Apathy and anxiety correlated positively. (Bottom panels). Apathy correlated negatively with switch behaviors, while anxiety correlated positively with switch behaviors. Anxious individuals were more sensitive to undesired feedback (no reward) and exhibited more switch behaviors than reward feedback.

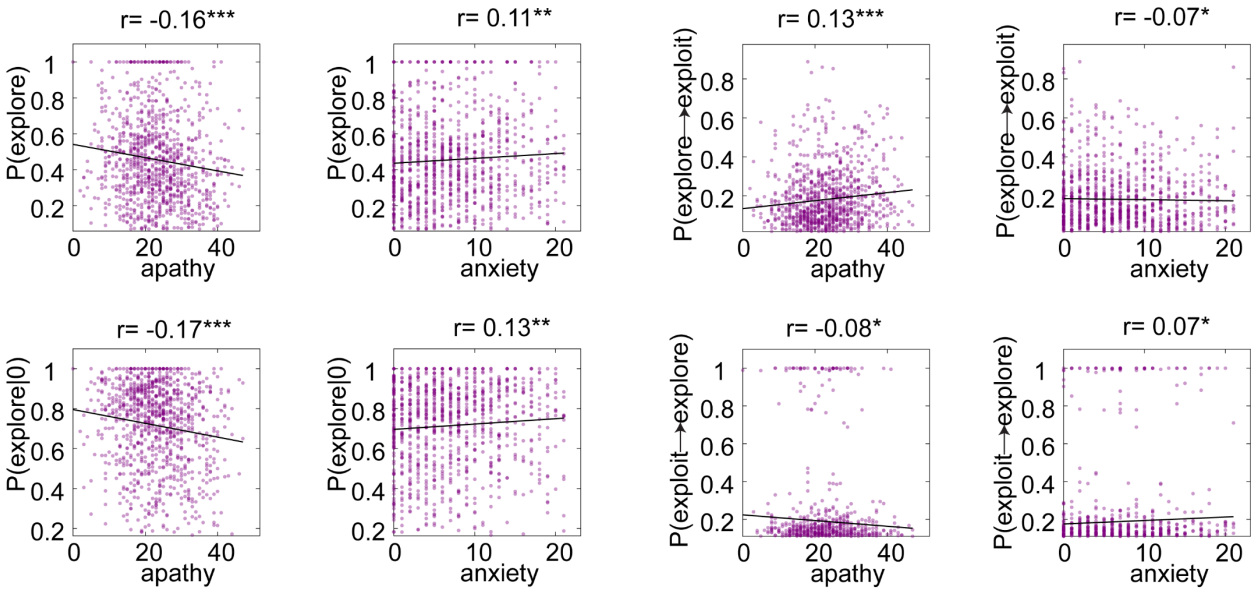

Figure S5. Apathy and anxiety have opposing relationships with exploration and explore and exploit state dynamics (plots with all data points)

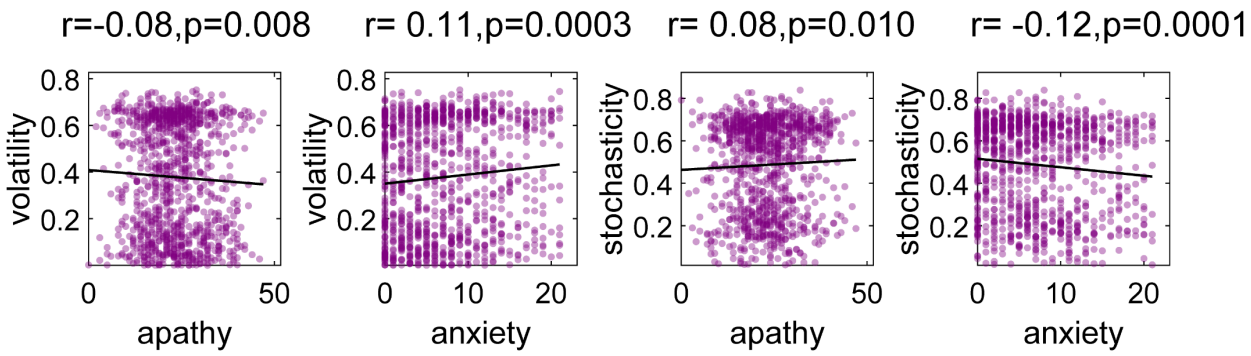

Figure S6. Apathy was positively correlated with stochasticity but negatively correlated with volatility estimation. Conversely, anxiety showed a negative correlation with stochasticity and a positive correlation with volatility (plots with all data points).

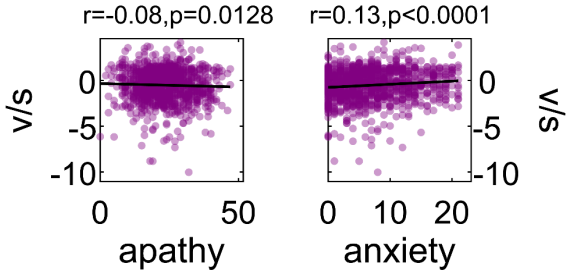

Figure S7. Distinctions in apathy and anxiety on the ratio of volatility to stochasticity.

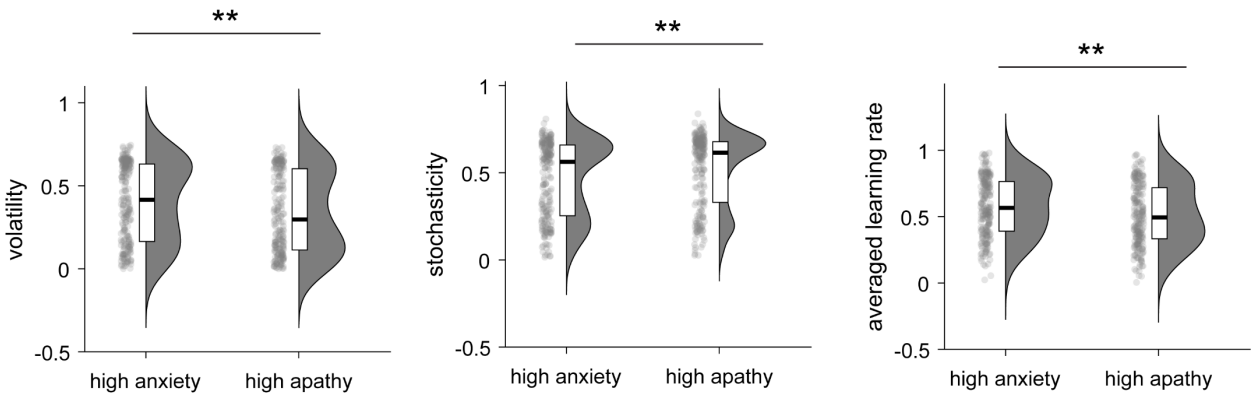

Figure S8. Distinct computational processes between highly anxious and highly apathetic individuals. Further comparisons showed that high anxious individuals had higher volatility estimates than those with high apathy (t(449) = 2.98, p=0.003). In contrast, high apathetic individuals had higher stochasticity estimates than their high anxiety counterparts (t(449) = 2.69, p=0.007), resulting in a higher learning rate among the high anxiety group (t(449) = 3.04, p=0.002). ***p*<0.01;

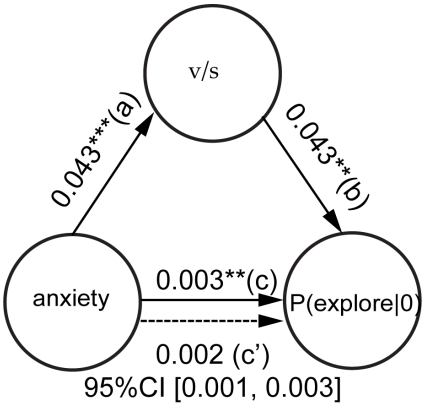

Figure S9. The ratio of volatility to stochasticity significantly partially mediated the relationship between anxiety and exploration after undesired feedback. ***p*<0.01; ****p*<0.001;

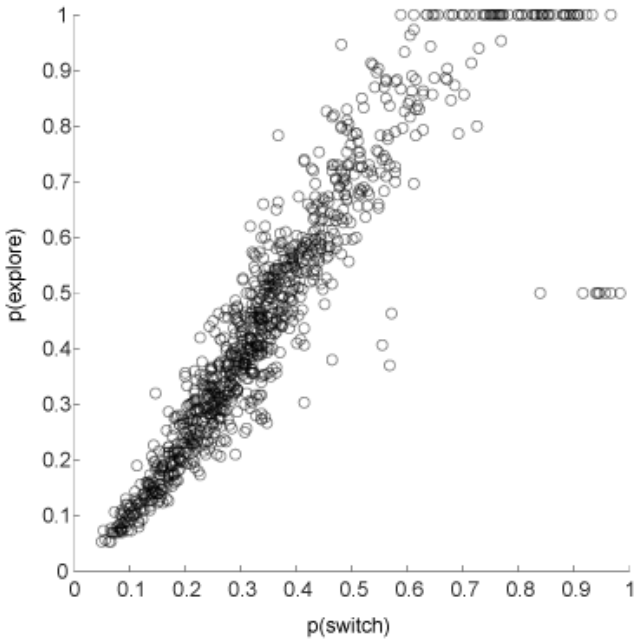

Figure S10. The strong correlation between P(explore) and P(switch) (r=0.916, p<0.0001). They measure separate but related constructs. P(switch) is a model-free measure that is sensitive to all changes in choice, regardless of context. P(explore) derives from the HMM, and is constrained by the inferred states, which are sensitive to the temporal structure of choices (unlike P(switch)). We expect a high degree of switching in the explore state, and a low degree in exploit, but not all switch decisions are necessarily labeled exploratory and not all repeat choices are labeled exploit. P(explore) is grounded in a theory of the latent states of exploratory decision-making and mathematically related to other properties of the HMM.

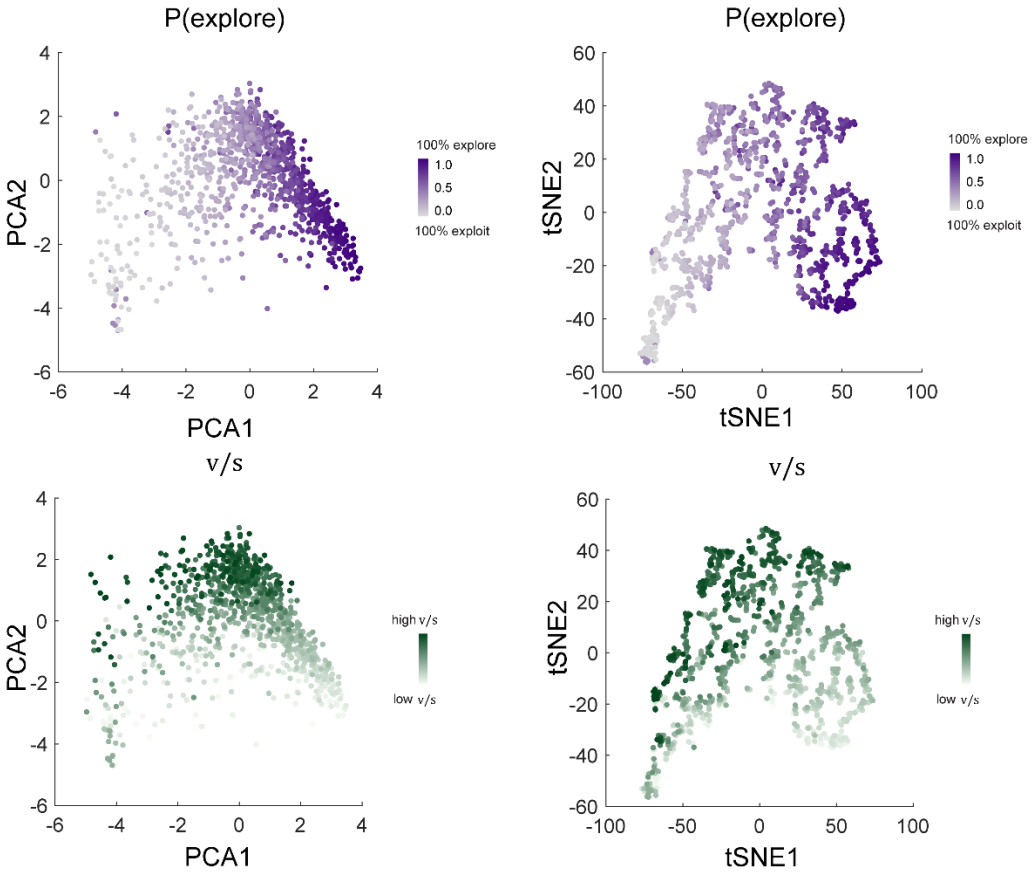

Figure S11. The low-dimensional space from PCA and tSNE, and the representation for P(explore), the ratio of volatility to stochasticity.

We conducted t-Distributed Stochastic Neighbor Embedding (t-SNE) and Principal Component Analysis (PCA) to confirm the manifold. The same eight-dimensional datasets from all participants were passed into the R package Rtsne, version 0.16 (available at https://cran.r-project.org/web/packages/Rtsne/index.html) with default parameter setting as n_component = 2, perplexity = 30, min_iter = 1000, metric = 'Euclidean'. We showed a similar manifold shape as UMAP found. Like in the main text, we also mapped model-free indices, as well as parameters from HMM onto t-SNE manifolds. The meaning of gradient change here is the same as with the UMAP manifold. The low-dimensional space from PCA is also quite similar to the manifold from UMAP and t-SNE. All other correlation results based on t-SNE and PCA scores can be found in the tables below.

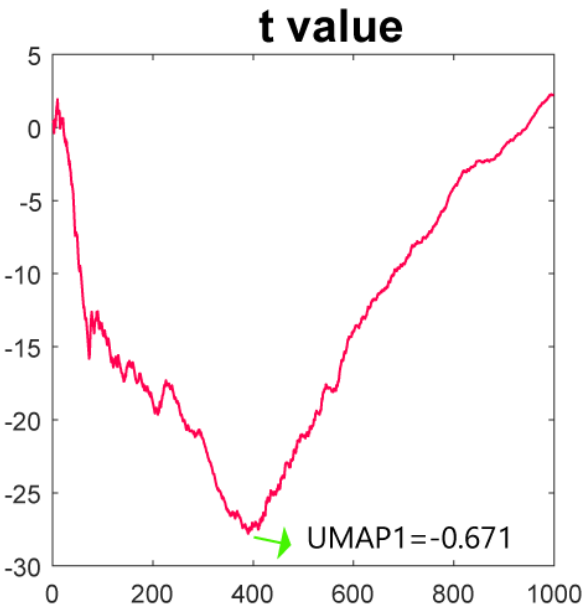

Figure S12. Turning point to divide the manifold into monotonically increasing and decreasing group.

A dimension 1 score of -0.671 marked the most significant negative coefficient, after which the relationship between dimension 1 and dimension2 gradually shifted to become positive

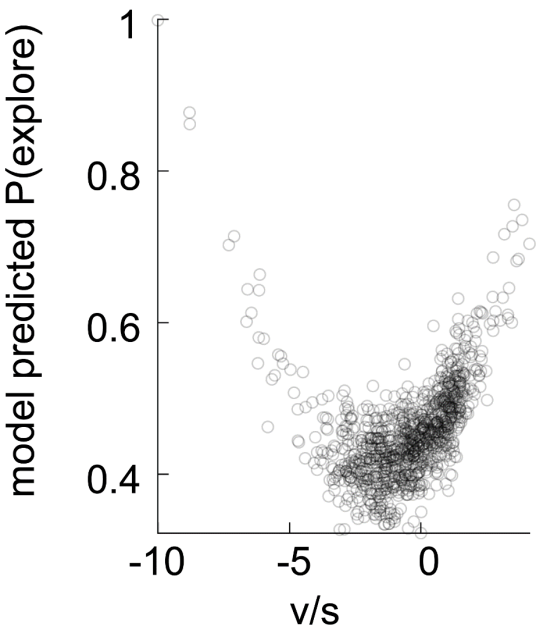

Figure S13. Non-linear relationship between the ratio of volatility to stochasticity and exploration. The results revealed that both the linear and quadratic terms are significant (linear term, coefficient = 0.03, SE = 0.005, t(996)=5.69, $p<10^{-8}$; quadratic term, coefficient = 0.009, SE = 13×10^{-4}, t(996)=6.948, $p<10^{-11}$), indicating a complex, non-linear relationship between the ratio of volatility to stochasticity and exploration. Linear and quadratic relationship between Ratio of volatility to stochasticity and exploration.

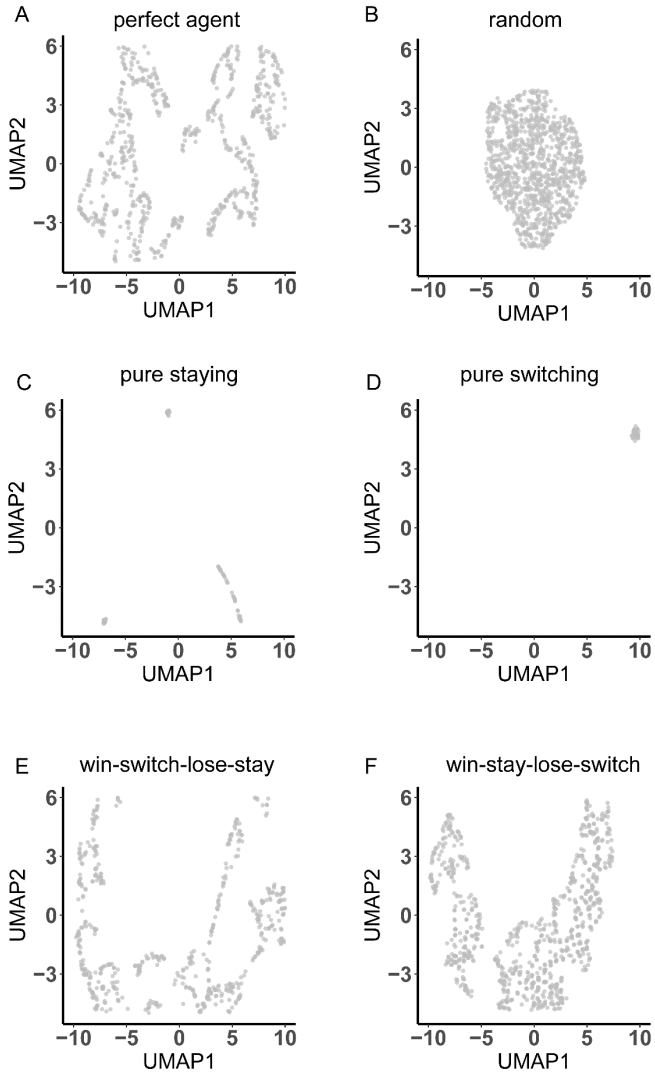

Figure S14. UMAP structures for simulated decision strategies. Low-dimensional representations of different simulated decision strategies reveal distinct patterns. (A) Perfect agent. (B) Random choice. (C) Pure staying. (D) Pure switching shows systematic alternation between options. (E) Win-switch-lose-stay (F) Win-stay-lose-switch. These qualitatively different UMAP structures validate that our dimensionality reduction approach captures meaningful variations in decision-making strategies rather than imposing artificial structure on the data.

TableS1. Detailed exclusion criteria table.

Exclusion criterion	Number excluded	% of initial sample
Incomplete questionnaires	54	3.57%
Incomplete task data	457	30.22%
Total excluded	511	33.8%
Final sample	1001	66.2%

Note: Some participants met multiple exclusion criteria. Numbers represent first criterion met in sequential screening.

Table S2a. Descriptive statistics for questionnaires (total score for GAD-7 and subscale scores for AMI)

	GAD-7	Apathy	Apathy-BA	Apathy-SM	Apathy-ES
Mean	7.17	30.46	10.54	12.69	7.23
SD	5.54	9.32	5.04	4.80	4.22
Minimum	0	4	0	0	0
Maximum	21	64	24	24	24

[a]GAD-7 = General Anxiety Disorder Screener; Apathy-BA = Apathy behavioral activation; Apathy-ES = Apathy emotional sensitivity; Apathy-SM = Apathy social motivation.

Table S2b. Descriptive statistics for questionnaires (mean score for GAD-7 and mean score for AMI and its subscales)

	GAD-7	Apathy	Apathy-BA	Apathy-SM	Apathy-ES
Mean	1.02	1.69	1.75	2.12	1.21
SD	0.79	0.52	0.84	0.80	0.70

Table S3. Descriptive statistics for model-free indices

	Model-free indices			
	Stay %	Switch %	Win. Stay	Lose.shift
Mean	0.646	0.353	0.861	0.732
SD	0.180	0.180	0.210	0.204

Table S4. Descriptive statistics for HMM indices

	Model-based indices from HMM					
	Exploitation%	Exploration %	Exploit-Exploit	Exploit-Explore	Explore-Explore	Explore-Exploit
Mean	0.544	0.455	0.811	0.188	0.818	0.181
SD	0.243	0.243	0.215	0.215	0.146	0.146

Table S5. Model performance

All models	loglikelihood	PXP (protected exceedance probabilities)	BIC
RW1	-2.1456e+05	0	440529.648
RW2	-2.1378e+05	0	444689.509
KF	**-1.7346e+05**	**1**	**364039.334**
VKF	-1.8786e+05	0	398548.485

Table S6. Statistical details to support Figure 3D and Figure 3E

measurements	**High vs Low anxiety**	**High vs Low apathy**
volatility	t= 2.82, p=0.005	t= -3.04, p=0.002
stochasticity	t= -2.99, p=0.002	t= 3.20, p=0.001
learning rate	t= 3.21, p=0.001	t= -3.11, p=0.002

Table S7. Results from tSNE and PCA

	t-SNE1	t-SNE2	PCA1	PCA2
Switch%	-0.914***	-0.323***	-0.967***	-0.139***
Stay%	0.914***	0.323***	0.967***	0.139***
Win.stay	0.713***	0.608***	0.832***	0.513***
Lose.shift	-0.631***	0.482***	-0.665***	0.603***
Exploration %	-0.912***	-0.250***	-0.937***	-0.016***
Exploitation %	0.912***	0.250***	0.937***	0.016
Exploit-exploit	0.536***	0.434***	0.683***	0.357***
Explore-explore	-0.484***	0.111***	-0.449***	0.199***
Explore-exploit	0.484***	-0.111***	0.449***	-0.199***
Exploit-explore	-0.536***	-0.434***	-0.683***	-0.357***
v/s	-0.151***	0.688**	-0.146***	0.807***

[a]all significant P-values reported here survive FDR correction.

(Original Benjamini & Hochberg FDR procedure, q<0.05)

Table S8.

Linear models	**stats$_{quadratic,\ apathy}$**	**stats$_{quadratic,\ anxiety}$**
P(switch)~1+ apathy+anxiety+anxiety^2+apathy^2	p=0.430	p=0.191
P(exploration)~1+ apathy+anxiety+anxiety^2+apathy^2	p=0.354	p=0.087
v/s ~1+ apathy+anxiety+anxiety^2+apathy^2	p=0.182	p=0.849

There is no significant quadratic relationship between apathy, anxiety, and exploration, neither between these affective states nor the ratio of volatility to stochasticity.

Table S9. Volatility, stochasticity and their correlations with HMM indices

	P(explore)-HMM	P(exploit)-HMM	P(explore→exploit)	P(exploit→explore)
volatility	0.151***	-0.151***	-0.152***	-0.056
stochasticity	-0.147**	0.147**	0.112***	0.030

(we report correlation coefficients here)

*** P<0.01, ***P<0.001*

all significant P-values reported here survive FDR correction.

(Original Benjamini & Hochberg FDR procedure, q<0.05)

Text S1. Effect size and clinical relevance

While we acknowledge that the observed correlations ($r \approx 0.10$-0.16) may appear small based on traditional standards, recent methodological discussions in the field of psychological research have challenged conventional interpretations of effect sizes, particularly in individual differences research. We believe it's important to consider our findings within this evolving context.

Interpreting effect sizes in individual differences research

Recent methodological discussions have challenged conventional interpretations of effect sizes in psychological research, particularly in the field of individual differences. Gignac and Szodorai (Gignac & Szodorai, 2016) conducted a comprehensive meta-analysis that suggests a recalibration of effect size interpretation:

r = 0.10: small but typical

r = 0.20: medium

r = 0.30: relatively large

In light of this, our observed correlations ($r \approx 0.1$-0.16) fall within the expected and meaningful range for this field of study. We acknowledge that these effects may appear small based on traditional standards, but we believe they warrant careful consideration within the context of individual differences research.

Benchmarking effect sizes

To provide further context, we find it helpful to compare our results with well-established psychological phenomena, as suggested by Funder and Ozer (Funder & Ozer, 2019) in their impactful paper *"Evaluating Effect Size in Psychological Research: Sense and Nonsense".*

The idea behind using benchmarks to evaluate effect size is that the magnitude of a finding can be illuminated by comparing it with some other finding that is already well understood. Some relevant comparisons include:

Findings	Effect size
Scarcity increases perceived value of a commodity	r = 0.12)
People attribute failures to bad luck	r = 0.10
Communicators perceived as more credible are more persuasive	r = 0.10

Resources from (Richard et al., 2003)

Additionally, **clinical comparisons** can provide an intuitive understanding:

Findings	Effect size
Effectiveness of antihistamines on allergy symptoms	r = 0.11
Pain relief from nonsteroidal anti-inflammatory drugs	r = 0.14

Resources from (Meyer et al., 2001)

These comparisons illustrate that our effect sizes are consistent with many important and widely accepted findings in psychology and clinical practice.

Cumulative Effects: Even small effects can have substantial real-world impact

While individual effects may appear small, we believe it's important to consider their cumulative impact over time. As Funder and Ozer (Funder & Ozer, 2019) argued, seemingly small effects can have substantial real-world impact when considered cumulatively.

Consider a compelling example from a large-scale study that analyzed 2 million financial transactions across 2,000+ individuals. The researchers found that the correlation between extraversion and holiday shopping expenditure was merely r = .09 (Weston et al., 2019). While this effect size might seem negligible for a single consumer, its significance becomes evident when considering a department store during the holiday season with thousands of shoppers.

In our study, this manifests in several ways:

Single Decision: While the effect on exploration (r = 0.13) may seem small for a single decision, its impact compounds over time.

Daily Impact: Approximately 20 decisions could be affected.

Monthly Impact: Around 600 decisions might be influenced.

Annual Impact: Over 7000 decisions could be shaped by these computational differences

Clinical significance in neuropsychiatric populations

We believe the clinical significance of these effects becomes evident when considering their impact on daily functioning in neuropsychiatric populations. For example:

Anxiety: A small bias in volatility estimation could lead to increased environmental scanning, more frequent strategy changes, and ultimately contribute to the maintenance of anxiety symptoms.

Apathy: A subtle reduction in exploration might result in fewer novel experiences, reduced opportunity detection, and gradual social withdrawal, potentially reinforcing apathetic symptoms.

The Clinical significance of small effects: population Impact and service Implications

Recent research (Carey et al., 2023) on youth mental health during the COVID-19 pandemic illustrates how small statistical effects can translate into substantial clinical outcomes. A seemingly modest effect size of d = 0.14 in depression scores led to 160,870 additional cases of depression in a population of 10 million youth, resulting in approximately 64,000 new referrals to mental health services and a 16% increase in clinical caseload.

Larger sample size are necessary provide more precise estimates and meaningful clinical implications

As Schönbrodt and Perugini (Schönbrodt & Perugini, 2013) demonstrated through Monte Carlo simulations, a sample size approaching 250 is typically needed for stable effect size estimates. This aligns with the growing recognition that many published studies, particularly in fields like psychology and neuroscience, are underpowered.

The current incentive structure in academia often rewards statistically significant results, which can lead to p-hacking and the inflation of small effect sizes. However, a more robust approach would be to incentivize the collection of data from large samples and the honest reporting of

effect sizes, even when they are small. This shift is crucial because smaller effect sizes, when estimated from larger samples, are more likely to reflect true population parameters.

Feng et al. (Feng et al., 2022) provided compelling evidence for this in their meta-analysis of brain imaging studies. They found that published brain imaging measures accounted for an average of only 8% of the variance in affective symptoms, with a wide confidence interval (1.6%–23%). Importantly, they noted that this average effect size was likely inflated due to the prevalence of small sample sizes in the field. And their findings support the need for large-sample clinical studies to robustly capture systematic variance of brain-affective symptom relationships

These findings underscore the need for large-sample clinical studies, particularly in fields like neuropsychiatric research. Larger samples not only provide more precise effect size estimates but also allow for the detection of smaller, yet potentially clinically relevant, effects. Moreover, they enable more robust statistical modeling to capture the complex relationships between brain function and behavior.

Future Directions

To further establish clinical utility, we propose:

Longitudinal studies: Track how computational parameters predict symptom progression, examine treatment response patterns, and assess functional outcomes over time.

Clinical validation: Replicate findings in clinical populations, compare with standard clinical measures, and evaluate sensitivity to treatment interventions.

In conclusion, while we acknowledge that the effect sizes in our study may appear small at first glance, we respectfully suggest that their clinical relevance becomes apparent when considering cumulative effects, population-level impact, and the specific context of neuropsychiatric research. We believe these findings are robust and valid since it’s comparable with previous established psychological findings.

Text S2. Complementary computational approaches: process model and latent space model

Kalman filter model is process model.

Process models were designed to identify individual differences in uncertainty estimations. Specifically, Kalman filter model, which consider the noise in the environment into the learning rate and further value updating. The two key noises in Kalman filter is, process noise (volatility), which represents how quickly the underlying reward probabilities change over time. And observation noise (stochasticity), which represents the randomness in outcomes even when the underlying probability stays the same. Together, using process model like Kalman filter, we can identify **individual differences** in certain parameters, (1) differences in process noise estimation (volatility), and (2) differences in observation noise estimation (stochasticity).

Hidden markov model is a latent state model.

HMM is a latent state model to identify **trial-by-trial differences** in exploration and exploitation. Briefly, our approach is based on certain statistical observations about the structure of choice sequences. We build on these observations to model exploration and exploitation as the latent, unobserved states underlying choices in HMM. Across several published papers, we have demonstrated that the HMM offers excellent face validity, better distinguishes quantitatively distinct behavioral and neural regimes than any previous method, and ensures that we avoid common pitfalls like conflating exploration with errors of reward maximization or noise estimation (Chen et al., 2021; Wilson et al., 2021). The HMM also delivers several useful quantitative measures that we have found to be related to individual differences. First, the HMM allows us to determine whether individual choices are due to the flexible computations associated with exploration or the rigid computations associated with exploitation (Ebitz et al., 2018). This has allowed us to identify individual differences in the rate of exploration with sex (Chen et al., 2021) and chronic stress (Kaske et al., 2022). Second, fitting the HMM involves estimating two interpretable free parameters. We have previously found that these parameters change under catecholamine neuromodulator modulation (Chen et al., 2023; Ebitz et al., 2019). Furthermore, these parameters explain individual differences in clinical symptoms. Third, because the fitted HMM represents a system of equations describing the evolution of behavior over time, we can use mathematical analyses to precisely characterize individual differences in the dynamics of exploration and exploitation, such as the stability or “stickiness” of each of

these behavioral states. This gives an intuitive and succinct “fingerprint” of an individual’s long-term tendency to explore, exploit, and switch between the two strategies

Text S3. Additional analyses for high and low anxiety/apathy group

To ensure our findings were not dependent on arbitrary cutoffs, we also validated our results using a standard deviation approach (±1SD from the mean):
High anxiety (n=186) vs Low anxiety (n=176)
High apathy (n=172) vs Low apathy (n=142)
We constructed linear regression models with three dependent variables (volatility, stochasticity, and learning rate), using group level (high vs. low) and the other affective state as predictors. For example, for anxiety analyses: volatility ~ High/Low group + Apathy.
Consistently, apathetic individuals overestimated stochasticity ($t(311) = 2.785$, $p=0.005$), underestimated the volatility($t(311) = -2.803$, $p=0.005$) , and had lower learning rate($t(311) = -2.817$, $p=0.005$) compared to those with low apathy. In contrast, individuals with high anxiety levels tended to overestimate volatility ($t(359) = 2.377$, $p=0.017$), underestimate stochasticity ($t(359) = -2.522$, $p=0.012$), and had higher learning rates compared to those with low anxiety $t(359) = 2.876$, $p=0.004$) .

References

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society*, *57*(1), 289–300.

Bilmes, J. A. (1998). A Gentle Tutorial of the EM Algorithm and its Application to Parameter Estimation for Gaussian. *Mixture and Hidden Markov Models*, *4*.

Carey, E. G., Ridler, I., Ford, T. J., & Stringaris, A. (2023). Editorial Perspective: When is a "small effect" actually large and impactful? *Journal of Child Psychology and Psychiatry, and Allied Disciplines*, *64*(11), 1643–1647.

Chen, C. S., Knep, E., Han, A., Ebitz, R. B., & Grissom, N. M. (2021). Sex differences in learning from exploration. *ELife*, *10*. https://doi.org/10.7554/eLife.69748

Chen, C. S., Mueller, D., Knep, E., Ebitz, R. B., & Grissom, N. M. (2023). Dopamine and norepinephrine differentially mediate the exploration-exploitation tradeoff. *BioRxiv : The Preprint Server for Biology*. https://doi.org/10.1101/2023.01.09.523322

den Ouden, H. E. M., Daw, N. D., Fernandez, G., Elshout, J. A., Rijpkema, M., Hoogman, M., Franke, B., & Cools, R. (2013). Dissociable effects of dopamine and serotonin on reversal learning. *Neuron*, *80*(6), 1572.

Ebitz, R. B., Albarran, E., & Moore, T. (2018). Exploration Disrupts Choice-Predictive Signals and Alters Dynamics in Prefrontal Cortex. *Neuron*, *97*(2), 450-461.e9.

Ebitz, R. B., Sleezer, B. J., Jedema, H. P., Bradberry, C. W., & Hayden, B. Y. (2019). Tonic exploration governs both flexibility and lapses. *PLoS Computational Biology*, *15*(11), e1007475.

Ebitz, R. B., Tu, J. C., & Hayden, B. Y. (2020). Rules warp feature encoding in decision-making circuits. *PLoS Biology*, *18*(11), e3000951.

Feng, C., Thompson, W. K., & Paulus, M. P. (2022). Effect sizes of associations between neuroimaging measures and affective symptoms: A meta-analysis. *Depression and Anxiety*, *39*(1), 19–25.

Fleiss, J. L. L. (2011). *Design and analysis of clinical experiments*. Wiley-Interscience. https://books.google.com/books?hl=en&lr=&id=PwCi4e8CxSsC&oi=fnd&pg=PR11&dq=The+Design+and+Analysis+of+Clinical+Experiments&ots=IKHHaJaKth&sig=qszDsBel1bJ6Bw1wU8UnUcZ1NN4

Forney, G. D. (1973). The viterbi algorithm. *Proceedings of the IEEE. Institute of Electrical and Electronics Engineers*, *61*(3), 268–278.

Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, *67*(1), 641–666.

Funder, D. C., & Ozer, D. J. (2019). Evaluating effect size in psychological research: Sense and nonsense. *Advances in Methods and Practices in Psychological Science*, *2*(2), 156–168.

Gignac, G. E., & Szodorai, E. T. (2016). Effect size guidelines for individual differences researchers. *Personality and Individual Differences*, *102*, 74–78.

Glickman, M. E., Rao, S. R., & Schultz, M. R. (2014). False discovery rate control is a recommended alternative to Bonferroni-type adjustments in health studies. *Journal of Clinical Epidemiology*, *67*(8), 850–857.

Hayes, A. F. (2017). *Introduction to mediation, moderation, and conditional process analysis: a regression-based approach. xvii* (Vol. 507). Guilford Press.

Hinton, G. E., & Roweis, S. (2002). Stochastic neighbor embedding. *Advances in Neural Information Processing Systems*, *15*. https://cs.nyu.edu/~roweis/papers/sne_final.pdf

Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: a review and recent developments. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, *374*(2065), 20150202.

Kaske, E. A., Chen, C. S., Meyer, C., Yang, F., Ebitz, B., Grissom, N., Kapoor, A., Darrow, D. P., & Herman, A. B. (2022). Prolonged Physiological Stress Is Associated With a Lower Rate of Exploratory Learning That Is Compounded by Depression. *Biological Psychiatry. Cognitive Neuroscience and Neuroimaging*. https://doi.org/10.1016/j.bpsc.2022.12.004

Mackinnon, D. P., Lockwood, C. M., & Williams, J. (2004). Confidence Limits for the Indirect Effect: Distribution of the Product and Resampling Methods. *Multivariate Behavioral Research*, *39*(1), 99.

McInnes, L., Healy, J., Saul, N., & Großberger, L. (2018). UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software*, *3*(29), 861.

Meyer, G. J., Finn, S. E., Eyde, L. D., Kay, G. G., Moreland, K. L., Dies, R. R., Eisman, E. J., Kubiszyn, T. W., & Reed, G. M. (2001). Psychological testing and psychological assessment: A review of evidence and issues. *The American Psychologist*, *56*(2), 128–165.

Piray, P., & Daw, N. D. (2020). A simple model for learning in volatile environments. *PLoS Computational Biology*, *16*(7), e1007963.

Piray, P., Dezfouli, A., Heskes, T., Frank, M. J., & Daw, N. D. (2019). Hierarchical Bayesian inference for concurrent model fitting and comparison for group studies. *PLoS Computational Biology*, *15*(6), e1007043.

Preacher, K. J., & Hayes, A. F. (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods*, *40*(3), 879–891.

Rhys, H. I. (2020). *Machine Learning with R, the tidyverse, and mlr*. Simon and Schuster.

Richard, F. D., Bond, C. F., Jr, & Stokes-Zoota, J. J. (2003). One hundred years of social psychology quantitatively described. *Review of General Psychology: Journal of Division 1, of the American Psychological Association*, *7*(4), 331–363.

Riffenburgh, R. H. (2014). *Statistics in Medicine* (3rd ed.). Academic Press. https://books.google.com/books?hl=en&lr=&id=Pd4KCgJeXeEC&oi=fnd&pg=PP1&dq=Statistics+in+Medicine&ots=9YwsKYsxUj&sig=DHlsuVnduwnFl3SeAY-zL2_Jw78

Schönbrodt, F. D., & Perugini, M. (2013). At what sample size do correlations stabilize? *Journal of Research in Personality*, *47*(5), 609–612.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*(4), 1004–1017.

Storey, J. D., & Tibshirani, R. (2003). Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences of the United States of America*, *100*(16), 9440–9445.

Waltmann, M., Schlagenhauf, F., & Deserno, L. (2022). Sufficient reliability of the behavioral and computational readouts of a probabilistic reversal learning task. *Behavior Research Methods*, *54*(6), 2993–3014.

Weston, S. J., Gladstone, J. J., Graham, E. K., Mroczek, D. K., & Condon, D. M. (2019). Who are the scrooges? Personality predictors of holiday spending. *Social Psychological and Personality Science*, *10*(6), 775–782.

Wilson, R. C., Bonawitz, E., Costa, V. D., & Ebitz, R. B. (2021). Balancing exploration and exploitation with information and randomization. *Current Opinion in Behavioral Sciences*, *38*, 49–56.

Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *ELife*, *8*. https://doi.org/10.7554/eLife.49547

Zorowitz, S., Solis, J., Niv, Y., & Bennett, D. (2023). Inattentive responding can induce spurious associations between task behaviour and symptom measures. *Nature Human Behaviour*, *7*(10), 1667–1681.