Responses to reviewers

Summary

Biological Psychiatry: Cognitive Neuroscience and Neuroimaging

Tracking #: BPSC-D-24-00490

Title: "Distinct computational mechanisms of uncertainty processing explain opposing exploratory behaviors in anxiety and apathy"

We sincerely thank all reviewers for their time and thoughtful feedback on our manuscript. We have carefully considered each comment and have made substantial revisions to address their concerns. In this detailed response, we comprehensively explain our revisions and include additional analyses where requested.

For clarity, we have structured this response as follows:

- Each reviewer comment is marked with an arrow (→) and shown in **blue**
- Our responses follow each comment with detailed explanations
- Revised text additions to the manuscript are shown in **purple**
- All changes have been tracked in both the Main Text and Supplementary Information

We have provided extensive details and additional analyses to address each concern raised during the review process thoroughly.

Additionally, in the HBI framework, the parameters were estimated in an unconstrained space following normal distributions. During the model fitting with HBI, volatility, stochasticity parameters were transformed using a sigmoid function, inverse temperature was transformed using exponential function. Previously, we incorrectly used an exponential transformation to transfer all parameters. We have corrected the parameter transformation (for volatility, stochasticity) implementation to use the sigmoid function throughout our analysis consistently. This correction aligns our reporting with the original model fitting procedure. All figures and tables have been updated accordingly (tracked changes). Importantly, this correction did not affect any of our statistical results or conclusions.

Reviewer 1

→1 While the study is statistically robust, the effect sizes reported are small. It is crucial to address the clinical relevance of these findings more explicitly, particularly in terms of real-world applicability. Specifically, could the authors discuss how these small effects translate into meaningful behavioral or clinical outcomes, especially in neuropsychiatric populations?

We appreciate that more discussion about our effect sizes is needed to aid the readers in interpreting our findings.

While 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.

Revisions in the paper:

1. In this revision, we integrated these methodological considerations into an opinion/mini-review in revised SI, Text S1.

2. We have added the following text to the last but one paragraph in Discussion (Main Text):

"The observed effect sizes in the current study ($r \approx 0.10$-0.16) align with recent recalibrations of effect size interpretation in individual differences research, where $r = 0.10$ is considered meaningful (Gignac and Szodorai, 2016). These effects are comparable to well-established phenomena like the effectiveness of antihistamines on allergy symptoms ($r=0.11$) (Funder & Ozer, 2019). While modest in isolation, such effects on exploratory decisions can accumulate substantially over time, potentially contributing to the maintenance of anxiety or apathetic behaviors through persistent influence on decision-making."

Here, we explain further:

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 ≈ 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.

<u>Benchmarking effect sizes</u>

Studies by Fan, Gershman and Phelps (Fan et al., 2022) and Scholl and colleagues (Scholl et al., 2022) found similar or even smaller effects when examining how emotions influence decision-making.
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:

-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)

(Richard et al., 2003)

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

-Effectiveness of antihistamines on allergy symptoms (r = 0.11)

-Pain relief from nonsteroidal anti-inflammatory drugs (r = 0.14)

(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.

<u>Cumulative Effects: Even small effects can have substantial real-world impact</u>

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 $r = 0.09$ (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

While the extension of the current findings to clinical populations is the subject of ongoing work in our lab, we believe the potential clinical significance of these effects becomes evident when considering how they might impact daily functioning in neuropsychiatric populations. For example, anxiety is associated with heightened environmental scanning (Charpentier et al., 2022). This could create a cycle where overestimating volatility leads to increased environmental scanning and strategy changes, and ultimately contribute to the maintenance of anxiety symptoms. In contrast, for individuals with apathy: A subtle reduction in exploration might result in fewer novel experiences, reduced opportunity detection, and gradual withdrawal (Fahed & Steffens, 2021), 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 to 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.

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.

→2 The sample is drawn from an online platform (MTurk and Prolific), which may not accurately reflect individuals with clinical levels of anxiety and apathy. Therefore, the authors should acknowledge this limitation more explicitly and propose future work with clinically diagnosed participants to strengthen the translational impact.

Responses:

We agree with the reviewer that the translation of our findings to clinical populations remains to be established. In the revised manuscript, we have addressed this limitation **in the last paragraph in Discussion**.

Revisions in the paper:

"While our findings provide valuable insights into uncertainty processing mechanisms, we acknowledge that our sample may not fully represent individuals with clinical levels of anxiety and apathy, potentially limiting generalizability to diagnosed populations. Future research with clinical samples will be crucial to validate and extend these findings, strengthening their translational impact."

→3 The manuscript utilizes multiple computational models, which might be difficult for non-specialist readers to follow. It would be helpful if the authors could provide a clearer explanation of how the Kalman filter and Hidden Markov Model results complement each other in the main text, possibly reducing the reliance on the supplementary methods section.

Responses:

Thank you for this suggestion, it is indeed important to clarify how the Kalman filter and Hidden Markov Model results complement each other.

Kalman filter model is a process model, while the Hidden markov model is a latent state model.

Process models, such as reinforcement learning models, or the Kalman Filter model we use here, seek to explain the algorithm that a decision-maker uses when making choices Specifically, the Kalman Filter model, estimates how individuals weigh sources of noise/uncertainty in updating values and making choices. Latent state models, in contrast, are designed to infer the underlying states that make certain choice patterns more or less likely by learning the statistical structure of choice sequences. The HMM we use is a latent state model that identifies **trial-by-trial differences** in states of exploration and exploitation. As we have previously demonstrated (C. S. Chen et al., 2021; Ebitz et al., 01/2018, 2019; Kaske et al., 2022), combining these two kinds of models allows us to examine individual differences in the process of decision making (how fast do participants adapt their behavior? how sensitive is that adaptation to forms of noise?) and the underlying states that constrain the kinds of choices people make (exploratory choices or exploitative choices).

Revisions in the paper:

1.We have added the following paragraph in the Method's section

Complementary computational approaches: process model and latent space model

To comprehensively characterize decision-making under uncertainty, we employed two complementary computational approaches: a Kalman filter process model capturing the individual differences in uncertainty processing and learning, and a Hidden Markov Model revealing the trial-by-trial differences in states of exploration and exploitation across individuals. These models provide distinct but complementary insights (more details see **Text S2**).

2. We added SI Text S2 to give more details.
Though we tried to reduce the reliance on the supplementary methods section, due to the word limits in BPCNNI (4000 words) and other additions to the Discussion in response to reviewers, we moved the explanations about process and latent space model into Text S2.

Reviewer 2:

1. Conceptual clarity

→Anxiety and apathy are defined primarily through behavioral tendencies in response to uncertainty, but the theoretical link between these affective states and their computational correlates (volatility and stochasticity) could be clarified. For instance, explain more fully why apathy is conceptualized as an "overweighting" of stochasticity rather than as an absence of perceived control. This task does not allow for control of the environment since participants' actions do not change the reward probabilities (e.g.; in many cases all three options might have a low reward probability so participants might "feel stuck" even when exploring).

Responses:

We appreciate your insightful comment that highlights our need to clarify the theoretical relationship between perceived control and computational parameters.

The weighting of stochasticity (relative to volatility) is indeed how perceived control is formalized in the computation model. While perceived control is a subjective experience that would require self-report to quantify, whether subjects respond to noise in the environment as if it were learnable or random provides objective evidence of the degree to which participants believe they can exert control over their outcomes. As the reviewer observes, in this task, that control is exercised by modifying choices, rather than modifying the environment. Since the absence of perceived control is indeed a crucial aspect of apathy, we hypothesized that apathy would be associated with an overweighting of stochasticity.

→More specifically, this sentence could be clarified a bit more: "Apathy, characterized by a lack of motivation and goal-directed behavior (4, 5), is an affective state associated with imprecise beliefs about action outcomes (6) and a tendency to persist with previous choices rather than explore (7). This suggests that apathetic individuals may view outcomes as primarily stochastic, attributing events more to chance than controllable variables"

To us this doesn't follow trivially from imprecise beliefs and tendency to persist with previous choice. These links need to be explained theoretically.

Responses:

We agree that the connection between apathy, imprecise beliefs, persistence with previous choices, and viewing outcomes as stochastic needs a more thorough theoretical explanation in the manuscript.

Revisions in the paper:

We rewrote the paragraph (3rd paragraph in Introduction) as:

"Apathy, characterized by a lack of motivation and goal-directed behavior (4, 5), is an affective state associated with imprecise beliefs about action outcomes (6) and a tendency to persist with previous choices rather than explore (7). These features are mechanistically linked: imprecise outcome beliefs increase uncertainty about new actions, potentially leading individuals to choose familiar options. This computational bias self-reinforces as reduced engagement limits action-outcome learning and restricted exploration prevents exposure to diverse outcomes. Building on these observations, we hypothesize that apathetic individuals may perceive outcomes as primarily stochastic rather than controllable, potentially perpetuating a cycle of reduced exploration and helplessness (8)"

2. Participants and scores

→How was the sample size determined?

<u>Revisions in the paper:</u>

We apologize for not providing more justification for our sample size. We have added these details to the Methods section

"Sample size was determined through a priori power analysis. To detect correlations of r = 0.1 (typical for individual differences research, see Text S1) with 80% power at α = 0.05, we required a minimum sample of 782 participants. We recruited 1500 participants to account for expected exclusions based on previous large-sample online studies (Fan et al., 2022; Scholl et al., 2022) and our own pilot work, expecting to achieve a sample of between 900-1100 participants, thus allowing for a buffer above the minimum sample size. Our final sample of 1001 participants provided 98% power to detect r = 0.1."

To explain further:

<u>Larger sample sizes provide more precise estimates and enable meaningful clinical implications</u>

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 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 (Funder & Ozer, 2019). **This shift is crucial because smaller effect sizes (e.g., 0.1-0.2), 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 capture systematic variance of brain-affective symptom relationships robustly.

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.**

→Please also specify the exclusion criteria a bit more since about a third of the participants were excluded. Was this due to many participants not completing the task?

Revisions in the paper:

We have added these details to the Methods section
"We recruited a sample of 1512 participants via Prolific (Prolific. co); exclusion criteria included current or history of neurological and psychiatric disorders. Participants were excluded if they did not complete all questionnaires (3.57% of initial sample) or they did not complete the bandit task (30.22% of initial sample) (TableS1). 1001 participants completed all questionnaires and the bandit task (age range 18-54, mean ± SD = 28.446 ± 10.354 years; gender, 493 female). All participants were compensated for their time in accordance with minimum wage."

Table S1. 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.

→Please also mention straightaway in the abstract and intro that this is a non-clinical sample.

Responses:

Thank you for this suggestion. In the revised manuscript, we mentioned that our sample is a non-clinical population.

Revisions in the paper:

1. Abstract, 2nd paragraph

“Methods

Participants (N = 1001, non-clinical sample) completed a restless three-armed bandit task that was analyzed using both latent state and process models. ”

2. Introduction, last paragraph

“To address these questions, we recruited 1001 participants from a non-clinical population with anxiety and apathy measurement.”

3. HMM modelling

Revisions in the paper:

We appreciate your close reading required for such detailed comments and suggestions about HMM. We agree the original manuscript did not provide enough information to clearly describe how the data was fitted by the HMM. **So we expanded our Method S2. Hidden Markov Model and added substantial details**.

Below we provided more details to address each points raised by you.

→The states in the HMM model were defined as exploit and explore. Were these states inferred in a data driven way? If so, why are they labelled as such?

Responses:

We apologize for the misunderstanding. Following extensive previous work (C. S. Chen et al., 2021; Ebitz et al., 01/2018, 2019, 2020), we employed a HMM model with a designed structure that categorizes trials into states according to whether behavior is random or persistent, which we call exploration and exploitation, respectively. The random and persistent states captured by the model have been previously validated as reflecting behavioral patterns characteristic of the normative definitions of exploration as being reward-independent, whose purpose is learning about rewards, and exploitation as reward-driven. The emissions model for the explore state was the maximum-entropy distribution for a categorical variable, a uniform distribution:

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

Where N is the number of stimuli that were presented (i.e. N = 3). 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:

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

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

This labeling is consistent with the normative understanding of exploration as sampling from all options with equal probability, and exploitation as repeatedly sampling a preferred option.

→If HMM was fit to behavioral data sequences (choices of bandit A,B or C) then we should get a stochastic strategy in each latent state (percentage of time player chooses bandit A, B or C).

Responses:

The HMM was indeed fit to behavioral data sequences (choices of bandit A, B, or C). However, as explained above, our model structure constrained the emission probabilities for each state:

In the explore state, the probability of choosing each bandit was equal (1/3 for each).

In each exploit state, the probability of choosing the corresponding bandit was 1, and 0 for the others.

These constraints were imposed to clearly differentiate between exploratory and exploitative behavior, based on theoretical considerations and previous research (C. S. Chen et al., 2021; Ebitz et al., 01/2018, 2019, 2020).

We also addressed this point in **Method S2. Hidden Markov Model**

→After fitting the HMM, is it possible that the researchers interpret one of the inferred states as "explore" if it shows patterns associated with exploration, such as higher rates of switching or responses to changes in reward structure? This inference is not automatic then; it relies on manual labeling post-fitting based on the statistical properties of the actions within each hidden state. For example, states showing frequent switching or less adherence to previously rewarded options might be labeled as "explore," whereas states with consistent choices or lower switching rates could be labeled as "exploit."

Responses:

Again we apologize for the misunderstanding, which was due to a lack of clarify in the manuscript, which we have addressed above and in the manuscript by further explaining the pre-structured nature of the HMM. The trials were not labeled manually post-hock, but automatically according to our predefined state labels.

→For greater clarity and accuracy, the paper could improve by detailing the criteria used to label these states and discussing the limitations of using HMMs in this manner for defining specific cognitive states.
Also, please provide more information in the SI how the HMM was fitted.

Revisions in the paper:

We have addressed this in the manuscript, see **Method S2. Hidden Markov Model**

4. Distinguishability of the methodological approaches

One strength of the manuscript is that the authors use various methodological approaches to understand this large dataset. They write "The HMM state-model of exploration-exploitation and the Kalman filter process model of uncertainty estimation represent complementary ways of understanding adaptive behavior that our mediation results suggest are intrinsically related." Overall, it is quite difficult to understand throughout the text how independent these methods are by design.

→First, it would be good to theoretically specify upfront potential relationships (and das interdependence).

Revisions in the paper:

We have clarified in the manuscript how the Kalman filter and Hidden Markov Model approaches complement each other.

1. We have added one more section in Methods:

"Complementary computational approaches: process model and latent space model

To comprehensively characterize decision-making under uncertainty, we employed two complementary computational approaches: a Kalman filter process model capturing the individual differences on how to learn and process uncertainty, and a Hidden Markov Model revealing the trial-by-trial differences in exploration and exploitation across individuals. These models provide distinct but complementary insights (more details see **Text S2**)"

2. Due to the word limits in BP:CNNI (4000 words in main body of text), **we had detailed explanations about process and latent space model in Text S2.**

3. And to avoid any confusion, we revised the original text "The HMM state-model of exploration-exploitation and the Kalman filter process model of uncertainty estimation represent complementary ways of understanding adaptive behavior that our mediation results suggest are intrinsically related."

The revised version:

"The HMM state-model of exploration-exploitation and the Kalman filter process model of uncertainty estimation represent complementary ways of understanding adaptive behaviors (see Text S2)."

→Second, the authors compare a series of metrics to apathy and anxiety scores but they to not show how these metrics are empirically related to each other.

Responses:

Thank you for this valuable suggestion about showing the relationships between different metrics. We have now added a new Table in SI, **Table S9. Volatility, stochasticity and their correlations with HMM indices**, which systematically examines these relationships.

Revisions in the paper:

See **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)

→5. Optimal behavior in the task

Could the authors derive the optimal the as a benchmark, i.e., the set of parameters in the winning model that perform best in this task?

Responses:

While we agree that optimality (and distance from optimality) are important aspects of behavior, we were specifically interested in examining how decision-making strategies differ in ways orthogonal to optimality. Our task was specifically designed to allow for a range of viable behavioral strategies that can (and do) achieve similar performance levels, and there is no closed-form (or easily specified) solution for an optimal strategy that we know of. This design feature is important because:

1. In real-world decision-making under uncertainty, there often isn't a single "optimal" strategy, but rather multiple strategies that can be equally effective depending on context and individual preferences.
2. Our research focus is not on how participants deviate from optimal behavior, but rather on understanding individual differences in computational strategies and their relationship with neuropsychiatric symptoms.
3. Previous studies using similar paradigms (Chakroun et al., 2020; C. S. Chen et al., 2021; Fan et al., 2022; Kaske et al., 2022) have demonstrated that different combinations of exploration-exploitation strategies can lead to comparable reward rates, making it difficult and potentially misleading to define a single optimal benchmark.

6. Various methods improvements

→While the paper applies FDR correction, it would be helpful to specify the correction method in greater detail and why it was chosen over others.

Responses:

Thank you for suggesting more detail about our multiple comparison correction approach. We used the Benjamini-Hochberg False Discovery Rate (FDR) procedure with q = 0.05 for the following reasons:

Appropriateness for current study

1. Multiple correlations testing related hypotheses

2. Interest in discovering true effects while controlling false positives (Benjamini & Hochberg, 1995)

3. Maintenance of reasonable statistical power (Riffenburgh, 2014)

Advantages over alternative methods

1. Better suited for correlated tests than Bonferroni correction (Glickman et al., 2014)

2. Balances Type I and Type II errors effectively (Storey & Tibshirani, 2003).

Revisions in the paper:

Due to word limits (4000 words) in main text, **we expanded FDR correction in Method S13** as follows:

"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."

→Moreover, given the large number of correlations, is there a reason why effect sizes (e.g., reporting standardized regression coefficients) were not reported? It might offer insight into the practical significance of findings.

Responses:

Thank you for your attention to our effect size reporting. We would like to clarify that our analyses are primarily correlational in nature, and we have reported correlation coefficients (r) throughout the paper, which are themselves standardized effect size measures (Nakagawa & Cuthill, 2007). These r-values directly indicate both the magnitude and direction of relationships between our variables of interest, representing effect sizes on a standardized scale from -1 to +1. Since our analyses focus on bivariate relationships, the correlation coefficient (r) is mathematically equivalent to the standardized regression coefficient (β) in simple regression with a single predictor (P. Y. Chen & Popovich, 2002). To enhance the interpretability of these effect sizes, we have now added practical interpretations of effect magnitudes; please see **Text S1**, which aims to address all effect size and sample size-related issues.

→The Kalman filter model was selected as the optimal model for describing participant behaviors. A brief justification as to what sets it apart from the other models and why this one was best would be helpful.

Responses:

Thank you for this question about model selection. We have detailed our model comparison and selection process in the **Methods section, Model fitting and comparison, and Supplement Method S6, Table S5**, where we report that the Kalman filter model was selected based on:

(1) Protected exceedance probability (PXP = 1); (2) Lower BIC value (364039.334) compared to alternative models: RW1: 440529.648; RW2: 444689.509; VKF: 398548.485. These quantitative metrics demonstrate the Kalman filter model's superior fit to participant behavior.

(2) The Kalman filter's formulation also aligns well with theories of how 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. While the RW1 (Rescorla-wagner model with general learning rate) and RW2 (with positive and negative learning rate) only quantify prediction errors through fixed learning rates, which did not incorporate the process noise and observation noise into the learning rate and value updating, thus they cannot separate different sources of uncertainty.

Revisions in the paper:

Based on your suggestions, we expanded our **Supplement Method S6**, adding:

"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”

→Also, while parameter recovery was performed, we could envisage also to perform model recovery (simulate behavioral data from model, fit various models to simulated data and see whether we recover the same model)?

Responses:

To check for model recovery, four datasets with nSubjects = 50 and nTrials = 300 each were simulated, based on the Kalman filter model, the volatile Kalman filter model, the Rescorla-wagner model, and the Rescorla-wagner model with two learning rates. Each simulated subject was again fitted using all models and compared with BIC (Bayesian Information Criterion)-based goodness-of-fit.

The figure below shows the confusion matrices. The Kalman filter model showed good identifiability with a 96% successful recovery rate, there was very little confusion with VKF (0.04%). The volatile Kalman filter was recovered in 44% of cases, suggesting that while its unique features can be identified, there is some overlap with the standard Kalman filter model (56%). This is theoretically sensible given that VKF is an extension of KF.

The Rescorla-Wagner model demonstrated excellent identifiability with a 96% successful recovery rate. The two-learning-rate Rescorla-Wagner model showed poor recovery (0.04%), suggesting that its additional complexity might not provide sufficiently distinguishable behavioral patterns from the simpler RW model in our task context.

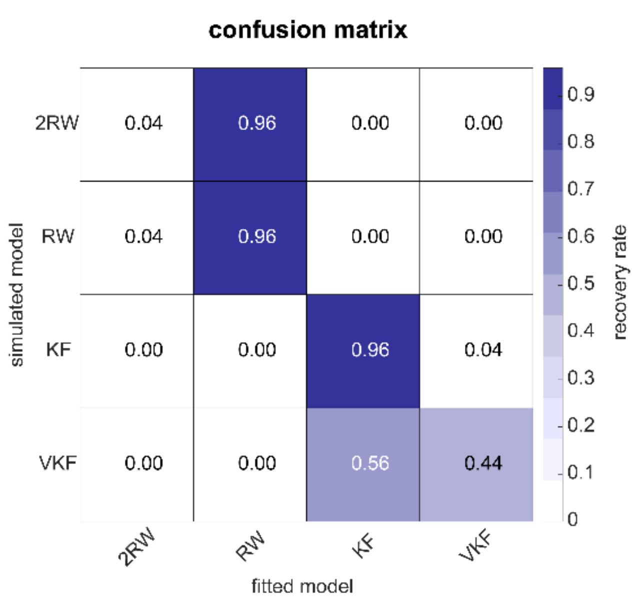

→A potentially significant problem in the analysis performed on page 12 is the artificial categorization of participants into "high" and "low" groups based on their scores on continuous scales (apathy and anxiety). This can have several drawbacks such as : a) Loss of information: By converting continuous data into categorical data, we lose the nuanced information contained in the original scores. This can reduce statistical power and the ability to detect true relationships. b) Arbitrary cutoffs: Using the top and bottom 25% as cutoffs is arbitrary. There's no inherent reason why these particular thresholds are meaningful, and different cutoffs could lead to different results. A careful justification of why the authors have chosen to dichotomise and why they have chosen these particular thresholds would be needed.

Responses:

We appreciate your concerns about dichotomizing continuous variables. **Our primary analyses appropriately treat anxiety and apathy as continuous measures, with all key findings based on correlational analyses using the full range of scores. The categorical analyses (top/bottom 25%) were included only to aid visualization and interpretation of computational differences**, provide concrete examples for clinical audiences, and demonstrate robustness of effects at different symptom levels.

To ensure our findings were not dependent on arbitrary cutoffs, we also validated our results using a standard deviation approach (±1SD from mean):

High anxiety (n=186) vs Low anxiety (n=176)

High apathy (n=172) vs Low apathy (n=142)

measurements	high vs. low apathy	high vs. low anxiety
volatility	t= -2.803, *p*=0.005	t=2.377, *p*=0.017
stochasticity	t= 2.785, *p*=0.005	t = -2.522, *p*=0.012

We agree that continuous analyses are more appropriate for our primary conclusions, and we have moved the categorical analyses to the Supplementary Information (SI, Text S3). We kept the violin plots just for visualization purposes.

Revisions in the paper:

1. In the revised manuscript, we maintain primary analyses using continuous measures and added the figure to show the results (Figure 3C, FigureS6), we kept the high/low apathy/anxiety plots just for visualization. And we removed categorical analyses from the main text, organize them into Table S6 (top/bottom 25%) and Text S3 (1SD from the mean value).
2. Accordingly, we revised figure caption for and added note at the end of figure "Note: Violin plots in panels D and E are provided for visualization purposes only. For details on the grouping methodology and statistical analyses, please refer to Table S6 and Supplementary Text S3."

→In Figure S2 it is not quite clear what the three different plots are showing.

Also, why are there a few quite substantial outliers in the third plot?

Response:

Thank you for raising this question about the apparent extreme values in Figure S3 (originally Figure S2). Figure S3 shows the parameter estimation (volatility, stochasticity, inverse temperature) from the split-half reliability analysis. The split-half reliability analysis involved fitting our model separately to each participant's first and last 150 trials. For parameter transformation, we employed the Hierarchical Bayesian Inference (HBI) framework, which typically assumes distributed priors for all free parameters. Following the approach of Piray & Daw (2020), we applied an exponential transformation (exp(x)) as the previous study used for the inverse temperature (Piray & Daw, 2020). This exponential transformation explains why the beta values appear extreme in the visualization. For example, beta estimates of 5 and 6 become 148.41 and 403.43 after transformation. However, it's important to note that the underlying parameter estimates (before transformation) for beta remain typically distributed.

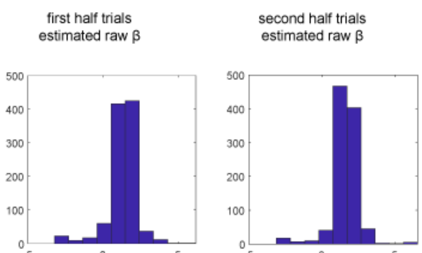

We have replaced **Figure S3** with the raw parameter estimates from the model to better illustrate the reliability of our measurements.

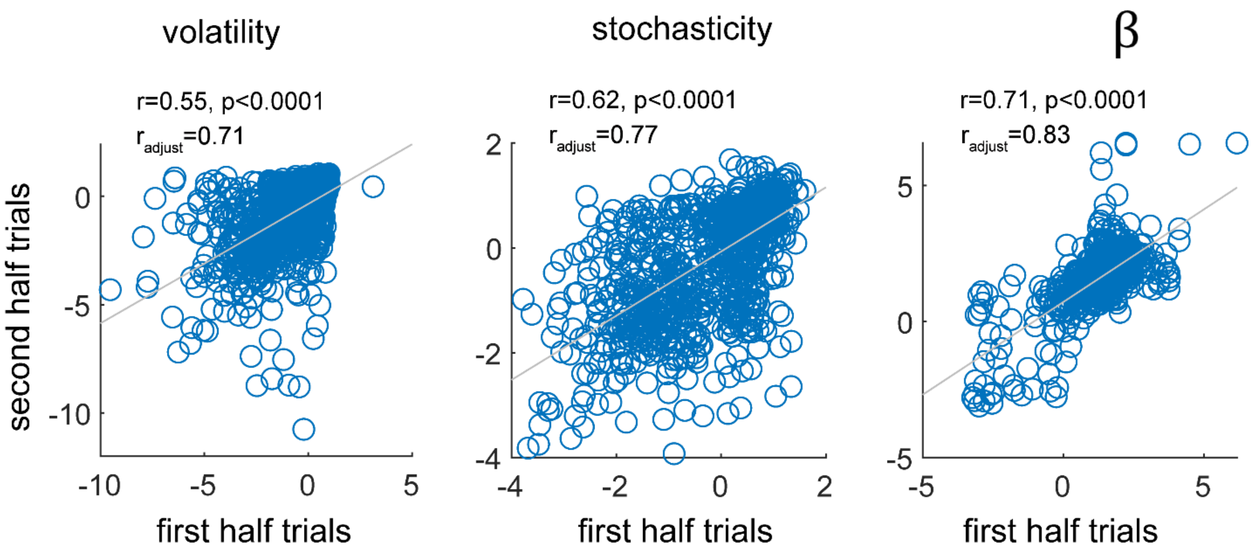

7. Minor comments on intro and discussion

→In the intro, the authors nicely mention two hypotheses. They could directly mention which one of those is supported by the data.

Responses and revisions in the paper:

Thank you for pointing out this opportunity to improve the clarity of our manuscript. Due to the word limits in the Main Text (4000 words), we added a very brief sentence at the end of the introduction:

"Our findings support the first hypothesis, revealing distinct behavioral patterns and computational mechanisms in apathetic and anxious individuals when faced with uncertainty."

→In the discussion, the authors could expand a bit on the group of participants in which both apathy and anxiety scores are high/low versus those in which they differ.

Response and Results:

Our analyses found no significant differences in volatility/stochasticity estimation or exploration behavior between individuals with both high anxiety and high apathy (N=63) and individuals with both low anxiety and low apathy (N=54) (we grouped participants by 1SD criterion) (all p>0.360)

measures	stats (H anxiety & apathy vs L anxiety & apathy)
volatility	t= -0.222, p=0.824
stochasticity	t= -0.026,p=0.978
P(explore)	t= -0.906, p=0.366

This pattern aligns with our main findings about the distinct computational mechanisms of anxiety and apathy.

Reviewer 3:

Questionnaire data:

→1. Were any data checks run on the questionnaire data (see e.g., Zorowitz et al.)?

Responses:

Thanks for your question!

1.1 We applied several data checks on the questionnaire and task data.

We used attention checks including checking the consistency of forward and reverse scored survey responses and the face validity of direct questioning, for instance the question "answer with the color of grass". Participants also had to meet a score threshold of 42% and an exploration threshold of 2 unique selections during the 25 practice bandit trials.

1.2 Thanks for bringing up this paper! Following the recommendations of Zorowitz et al.(Zorowitz et al., 2023), we have double checked for our questionnaire data, see details below:

Revised inter-item standard deviation (ISD) analysis

While we acknowledge the importance of data screening, as Zorowitz et al. (2023) highlighted, we have carefully adapted these methods for the GAD-7 and AMI, considering its unique characteristics as a brief clinical measure. For example, the split-half reliability is not quite useful for GAD-7 scale since it only has 7 items.

For the GAD-7 and AMI scores, we recognize that **consistent scores across items may reflect valid symptom presentations**. For example, if an individual indeed does not feel anxious at all over the last two weeks, they will have a 0 score for every single item. The 4-point response scale (0-3) limits the possible response patterns. The short length of the scale also makes some traditional quality metrics (like split-half reliability) less applicable.

Therefore, we modified our quality-checking approach to:

(1) Only flag response patterns that show implausible alternations (e.g., in GAD-7, extreme oscillations between "not at all" and "nearly every day" 0-3-0-3-0-3-0, or 3-0-3-0-3-0-3)

(2) Consider the overall pattern rather than just statistical variation

Our adapted screening procedure includes:

Step 1. Calculating the inter-item standard deviation (ISD)

Calculating the standard deviation of responses across all 7 items (GAD7), 18 items for AMI for each participant (Of note: low ISD is not considered problematic as it may represent consistent anxiety/apathy levels)

<u>Step 2. Detecting the extreme alternation pattern:</u>

Examining consecutive item responses for extreme jumps

Calculating the absolute differences between adjacent responses

Counting the proportion of extreme jumps (differences ≥ 3 points)

<u>Step 3. Flagging Criteria: a response pattern is flagged as suspicious only when BOTH conditions are met</u>

ISD > 2 (very high response variation) & Extreme alternations > 0.6 (60% of responses show extreme jumps)

This more fine-grained approach helps 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.**

<u>Revisions in the paper</u>

In SI, we added Method S14

<u>Method S14. Data quality check</u>

We have attention checks including checking the consistency of forward and reverse scored survey responses and the face validity of direct questioning including "answer with the color of grass". 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.

→2. The AMI scores are quite high - well into the clinical range - particularly for emotional and behavioral apathy. Symptom scores tend to be somewhat elevated in online (vs. normative) samples but this is particularly elevated, which raises concerns about the validity of this data.

Thank you for raising this concern, which results from a simple misunderstanding. For the Table S2 (originally Table S1), which is likely the source of the reviewer's observation, we presented *total scores* rather than subscale means, which we believe the reviewer likely assumed. In fact the AMI scores are well within the range found in prior studies and below the suggested clinical thresholds.

To clarify we added **another table** as follows:

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

Furthermore, we compared our AMI mean score with previous findings

1. Petitet, P., Scholl, J., Attaallah, B., Drew, D., Manohar, S., & Husain, M. (2021). The relationship between apathy and impulsivity in large population samples. *Scientific Reports*, *11*(1), 4830.(Petitet et al., 2021)

Table 1 Demographics characteristics.

From: The relationship between apathy and impulsivity in large population samples

Dataset	F/M	Age	AES	AMI	BIS-11	UPPS-P
1. Gillan et al. [42]	823/590	33.0 (10.8)	31.5 (8.2)	–	61.1 (11.2)	–
2. Rouault et al. [44]	257/240	35.6 (10.6)	32.3 (9.7)	–	58.3 (12.2)	–
3. Patzelt et al. [43]	377/461	36.1 (10.3)	32.5 (9.9)	–	57.1 (12.4)	111 (26)
4. Seow and Gillan [45]	228/209	35.3 (10.3)	30.4 (8.8)	–	56.9 (12.6)	–
5. Online dataset	195/194 (5)	28.0 (6.0)	–	1.72 (0.52)	60.4 (9.8)	–
6. Laboratory dataset	89/87	24.8 (4.3)	–	1.33 (0.41)	61.9 (10.8)	–

The table shows the number of females (F) and males (M) included in each study ("Prefer not to say" in parenthesis). Group mean total scores (standard deviation in parenthesis) are shown for each questionnaire included in a dataset.

Our sample's AMI score vs their online dataset' AMI

t = -0.888, p = 0.374

2. Norbury, A., Hauser, T. U., Fleming, S. M., Dolan, R. J., & Huys, Q. J. (2024). Different components of cognitive-behavioral therapy affect specific cognitive mechanisms. Science Advances, 10(13), eadk3222. (Norbury et al., 2024)*Table 1*

Our sample's AMI score vs their online datasets' AMI

Published AMI (Norbury et al., 2024)	**Our sample AMI vs. published AMI**
N=100, mean(SD) = 1.8(0.8)	**t = -1.317, p = 0.187**
N=208, mean(SD) = 1.6(0.8)	**t=1.59, p=0.110**

In summary, our AMI score fall within the normal range.

→3. A methodological strength is the model checking, but some of the results are concerning. In particular, the parameter recovery values (in table S5) for the volatility and stochasticity parameters are quite low.

Responses:

Thank you for highlighting this important concern about the parameter recovery values. We acknowledge that our previous recovery rates were low. In response, we have conducted a new parameter recovery analysis that strictly follows the standard procedure described in PLOS Computational Biology papers (page 21/26) and implements the official code for the original Kalman Filter model shared by Piray & Daw (Piray & Daw, 2020) on GitHub (https://github.com/payampiray/piray_daw_2020_ploscb). We modified the simulation process according to our specific task parameters and Kalman filter function.

*"**Recovery analysis of parameters***

For this analysis, data were generated based on the binary VKF (Eqs 14–19). In particular, the observation on trial t, o_t, was randomly drawn based on the sigmoid-transformation of m_{t-1}.

The choice data were also generated randomly by applying the softmax as the response model with parameter β. Similar to experiment 1, for each artificial subject, we assumed 4 sequences of observations and actions (i.e. 4 cues) with 120 trials. These values were used as the group parameters: λ = 0.2, v0 = 5, ω = 1, and β = 1. For generating synthetic datasets for simulations, the parameters of the group of subjects (50 subjects) assigned to each model were drawn from a normal distribution with the standard deviation of 0.5."

Following this standard procedure, 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 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.

Here are the updated recovery results

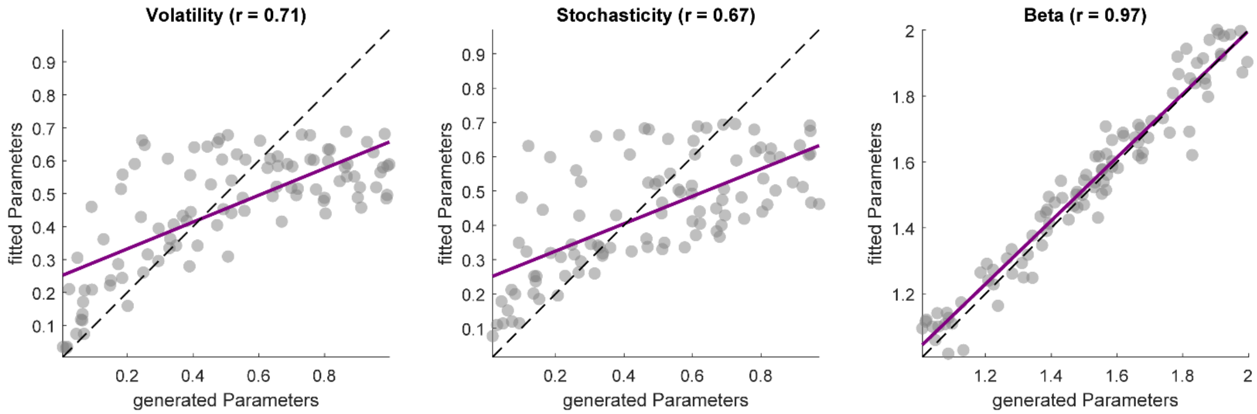

Revisions in the paper:

We updated the Method S8. Model validation and Figure S2

→4. It appears that the data were estimated hierarchically, with individual estimates then used for analyses - this may be reducing reliability and affecting the validity of other results.

Responses:

We value this opportunity to clarify our methodological approach and address your important concerns about reliability and validity.

However, based on both statistical theory and empirical evidence, we would like to respectfully explain why hierarchical methods enhance rather than reduce the reliability of parameter estimates and how Hierarchical Bayesian Inference (HBI) (Piray et al., 2019) is appropriate for both group and individual parameter estimates.

Theoretical foundation of hierarchical methods

First, hierarchical Bayesian estimation implements partial pooling across subjects, which helps constrain individual estimates when data are noisy (Bailey, 2005; Gelman, 2003). This approach acknowledges both individual variations and group-level patterns, providing a balanced framework for studying individual differences (Karvelis et al., 2023).

Statistical advantages of HBI

As the original paper introduced HBI (Piray et al., 2019) stated on Page 23

*"**Empirical Bayes methods play an increasing role in modern statistics. These methods essentially take a hierarchical approach, by assuming that individual data are generated based on the probabilistic properties of the population. This hierarchical approach has important consequences**. The most important consequence is that they provide a promising solution to the classical problem of priors in Bayesian statistics by providing informative, yet objective, priors at the individual level. Furthermore, by partly sharing parameters across subjects, they reduce overfitting relative to non-hierarchical models, which in turn allows them to confidently fit more complex models with a smaller penalty for overfitting. This is because non-hierarchical methods assume that the extra parameters of a complex model are independent. **For example, consider a model space in which the more complex model has one extra free parameter and there are 40 subjects in the dataset. Fitting the dataset with the complex model using non-hierarchical methods introduces 40 additional independent free parameters, driving the danger of overfitting**, and accordingly an excessive penalty to account for this possibility in assessing the evidence for the model. The hierarchical approach, however, assumes that the individual parameters are dependent, as they are all generated according to the same distribution, sharing a single mean parameter and smaller deviations from it. Modeling this hierarchical dependency enables those methods to avoid penalizing complex models as excessively."*

Empirical evidence supporting hierarchical methods

Multiple studies have rigorously validated hierarchical approaches in individual difference research, computational psychiatry, and neuroscience:

Ahn et al. (2011)(Ahn et al., 2011) use empirical research to argue that, compared to MLE, the hierarchical Bayesian estimation is the best method for obtaining accurate individual and group parameter estimates.

Karvelis, Paulus, and Diaconescu in their latest review (Karvelis et al., 2023), Section 4.1. Hierarchical model fitting methods can improve reliability, where they argued that hierarchical Bayesian methods improve parameter estimation by: (1) accounting for uncertainty at different levels; (2) allowing individual parameter estimates to be informed by group statistics and vice versa; (3) reducing the impact of noise in individual-level data

The reliability of hierarchical methods for studying individual differences has been demonstrated in several high-impact studies, for example: Swart et al. (Swart et al., 2017) employed hierarchical parameter estimation to study individual differences in dopaminergic drug effects on learning. Zhang and Gläscher (Zhang & Gläscher, 2020) applied hierarchical parameter estimation to investigate the individual differences in social observational learning. More related research see (Chakroun et al., 2020; Piray & Daw, 2020; Sapey-Triomphe et al., 2023).

Thus, the hierarchical approach actually strengthens rather than weakens individual difference analyses.

→5. Figure S2 also indicates other model estimation issues leading to extreme beta values (assuming beta is the last panel).

Responses:
Thank you for raising this question about the apparent extreme values in Figure S3 (originally Figure S2). This figure shows the parameter estimation (volatility, stochasticity, inverse temperature) from the split-half reliability analysis. The split-half reliability analysis involved fitting our model separately to each participant's first and last 150 trials. For parameter transformation, we employed the Hierarchical Bayesian Inference (HBI) framework, which typically assumes distributed priors for all free parameters. Following the approach of Piray & Daw (2020), we applied an exponential transformation (exp(x)) as the previous study used for inverse temperature (Piray & Daw, 2020). This exponential transformation explains why the beta values appear extreme in the visualization. For example, beta estimates of 5 and 6 become 148.41 and 403.43 after transformation. However, it's important to note that the underlying parameter estimates (before transformation) for beta remain typically distributed.

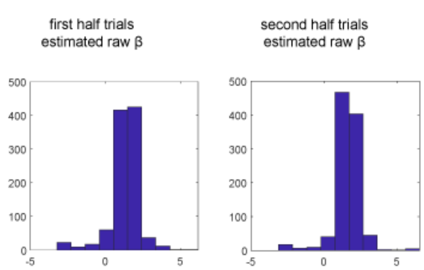

And here is the distribution for beta value (before transformation) with all trials.

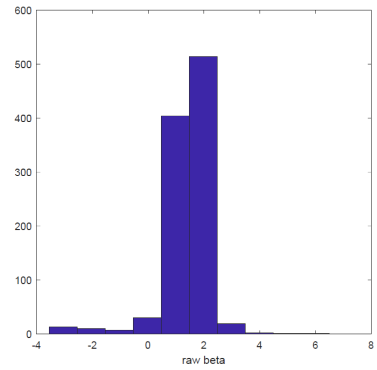

We have **Figure S3** with the raw parameter estimates from the model to better illustrate the reliability of our measurements.

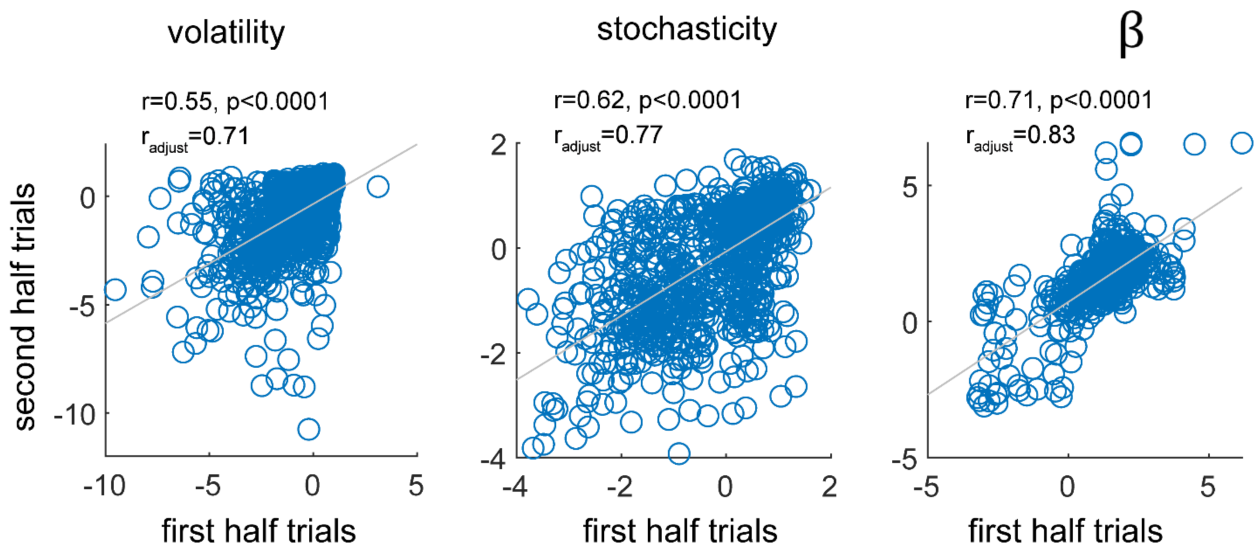

→6. What was the parameter recovery for the HMM transition probabilities?

Responses:

We conducted a parameter recovery analysis for the HMM transition probabilities in response to your query.

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 via *hmmgenerate* function in matlab, 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 by using the Pearson correlation. The recovery results were robust, we found the mean value of correlation between true and fitted parameters for P(explore | explore)= 0.644, and the averaged correlation for P(exploit | exploit)= 0.793.

Transition parameter	recovery metrics	
	mean r	*SD*
Explore → Explore	0.644	0.093
Exploit → Exploit	0.793	0.098

Note. mean r = Pearson correlation coefficient between true and recovered parameters, mean recovery rate from 50 simulations; *SD* = Standard deviation.

Revisions in the paper:

1, we added Method S7. Parameter recovery for HMM

2, we added one sentence to index Method S7 in Main Text, Method Section

Please note the last sentence underlying ***Hidden Markov Model*** **section**, "Model details are provided in Supplement Method S2. Model results of HMM can be found in Figure2 and Table S4. Parameter recovery for HMM see Method S7."

→7. HBI does not use empirical priors, as stated, but estimates hyperparameters and lower-level parameters simultaneously. The authors do use empirical priors based on a MAP analysis but this is a separate modeling step.

→8. What were the starting values based on the MAP analysis? This is not reported.

→9. No convergence statistics or details on the estimation (number of chains, samples per chain, how convergence was assessed) are reported for the HBI estimation.

We appreciate your detailed methodological questions regarding HBI implementation. Since questions 7-9 are closely related, we will address them together to provide a comprehensive response.

We acknowledge that our Main Text did not provide sufficient methodological details (but see SI, MethodS6 and original HBI methodological paper by Piray et al), as we relied on the well-documented HBI framework from Piray et al. (2019).

In HBI framework, priors are constructed based on data

HBI does estimate hyperparameters and individual parameters simultaneously, but HBI implements empirical Bayes principles where group-level distributions serve as priors for individual parameters. These priors are empirical in that they are informed by the data through an iterative process. Key evidence from Piray et al. (Piray et al., 2019), in the Discussion, Page 23, paragraph01:
"In this work, we took an empirical Bayes approach [31,32], in which priors are constructed based on data. In other words, parameters at the individual level are regularized by statistics across all individuals in the group"

HBI does not use MAP analysis.

The HBI does not rely on MAP analysis for analyses. Instead, it uses (1) direct initialization of individual parameters using **Laplace approximation** (2) full **variational Bayesian inference** for hierarchical parameter estimation and (3) iterative updates of both individual and group parameters

The initialization process is described in the Methods (Page 31): *"We initialize the parameters θkn and Akn by fitting all models separately to all participants (with an initial Gaussian prior), i.e., assuming as if zkn = 1."*

Based on HBI's feature and MAP's feature, here we listed the reasons that the HBI did not involve MAP: (1) HBI needs full posterior distributions, MAP's point estimates would lose uncertainty information but only provide point estimate; (2) HBI performs hierarchical inference on multiple levels, MAP is a single-step optimization

Regarding starting values and convergence

HBI uses variational inference, not MCMC, and thus does not require multiple chains. The convergence is monitored through changes in normalized parameters between iterations; the HBI algorithm terminates when dx < 0.01 or a maximum of 50 iterations is reached. Of note, HBI rarely hits the iteration limit of 50. As the paper Page 31 noted: "*In our analyses, we terminated the algorithm if the change in the normalized value of parameters between two consecutive iterations, j − 1 and j, was smaller than 0.01*"
where the dx = sqrt(mean((x-xpre).^2) (in matlab code)

Figures:

→10. The binning of scatterplot data (in e.g., Figure 1C) is confusing. Why not plot all data points as done in 1B?

Responses:

Thank you for this suggestion for data visualization. **We added scatter plots showing all individual data points in the Supplementary Information**

Figure S4 (corresponding to main Figure 1C-1D)

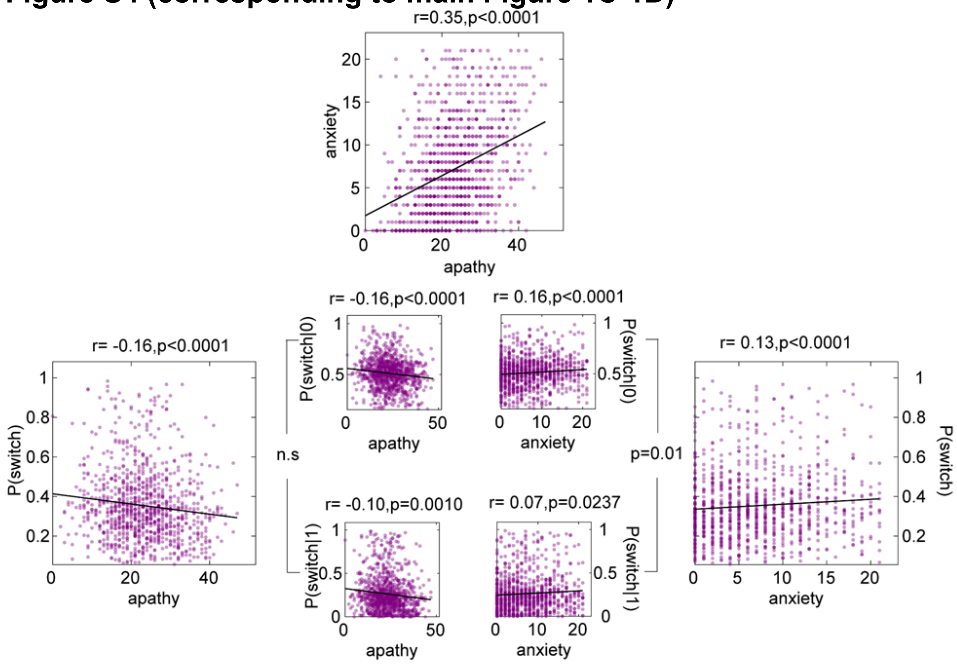

Figure S5 (corresponding to main Figure 2C-2D)

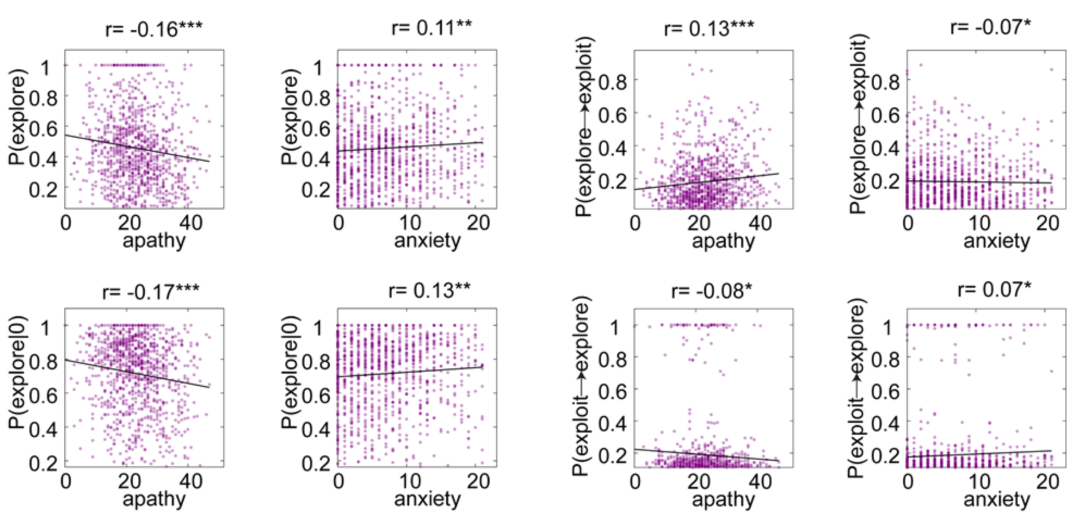

Figure S6 (corresponding to main Figure 3C)

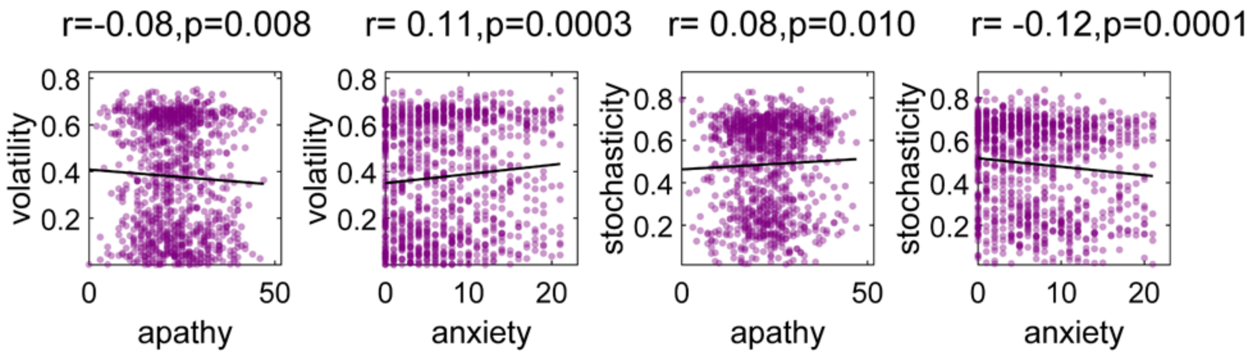

Figure S7 (corresponding to main Figure 4A)

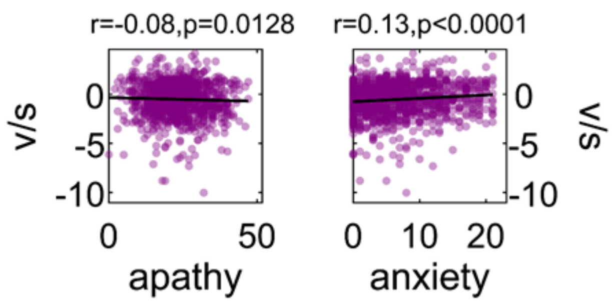

We maintained the binned correlation plots (25 quantile bins based on the x-axis) with error bars showing the standard error in **Main Text** for the following reasons:

1. Specific reason for Figure 1B with all individual points

The aim of Figure 1B is to show most of our individuals had good performance in this task, thus the individual points are essential here to clearly show the distribution of performance across participants and identify outliers

2. Data density:

Plotting all 1001 participants can create significant overplotting. The overplotting can obscure underlying patterns, especially in dense regions of the plot. The binned approach helps reveal the central tendency and spread of the data.

3. Visual clarity:

The binning helps highlight the underlying relationship between variables, and it reduces visual noise while maintaining statistical accuracy (All statistical analyses were performed on the raw (unbinned) data, as noted in the figure caption)

→11. Panels in S2 are not labeled.

<u>Revisions in the paper:</u>

We revised Figure S3 (originally FigureS2) with labels for each panel.

Framing, analysis, and interpretation:

→12. More explanation and contextualization of what P(explore) from the HMM means, and where this value comes from, is needed in the main text.

<u>Responses:</u>

We agree that the manuscript would benefit from a better explanation of P(explore). P(explore) is defined as the proportion of trials classified as **exploratory states** in a participant's choice sequence derived from HMM.

<u>Revisions in the paper:</u>

Hidden Markov Model

"We fit a Hidden Markov Model (HMM) to the behavior, to decode the hidden state of each trial for each participant. We fit HMM via expectation-maximization using the Baum-Welch algorithm and decode hidden states from observed choice sequences by the Viterbi algorithm(32). From this analysis, we extracted two types of measures: First, p(explore), which quantifies the overall proportion of trials classified as exploratory states for each participant. Second, the transition probabilities, which characterize the temporal dynamics of state switching: p(explore→explore) indicates the probability of maintaining an exploratory state between consecutive trials, with higher values reflecting more sustained exploration periods; p(exploit→exploit) represents the probability of maintaining an exploitative state, with higher values indicating more persistent exploitation of chosen options (32). Model details are provided in Supplement Method S2. Model results of HMM can be found in Figure2 and Table S4. Parameter recovery for HMM see Method S7."

→13. How correlated are p(switch) and p(explore) from the model-free and HMM analyses? Do they measure separable constructs?

Responses:

Thank you for this question.

→How correlated are p(switch) and p(explore) from the model-free and HMM analyses

The correlation between P(explore) and P(switch) is: $r=0.916$, $p<0.0001$

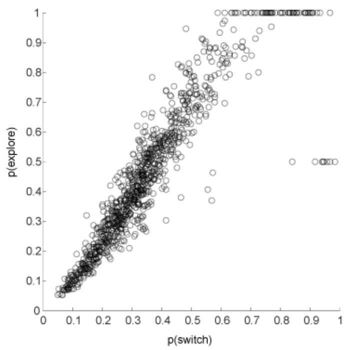

P(switch) is calculated as the proportion of trials where participants selected a different option from their previous choice. P(explore) quantifies the overall proportion of trials classified as exploratory states from HMM for each participant.

→Do they measure separable constructs?

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.

Revisions in the paper:

We added Figure S10.

→14.1 How correlated are p(explore) from the HMM and the beta parameter from the Kalman filter model?

Responses:

The correlation between P(explore) from HMM and beta from KF is: $r = -0.667$, $p<0.0001$ indicating a strong negative relationship. This result makes theoretical sense because the beta indicates choice consistency; high beta (more consistent choices) should correspond to low P(explore), while low beta (more random choices) should correspond to high P(explore). So, this result provided convergent validity between our two modeling approaches

→14.2 How much of exploration is directed vs. random?

Responses:

While our task was not designed to distinguish between directed and random exploration, we conducted analyses following the classification approach used in previous work to address your question (Chakroun et al., 2020; Daw et al., 2006; Wiehler et al., 2021). **We were able to compute directed/random exploration ratios (analysis details see below, point #5) and show the results below:**

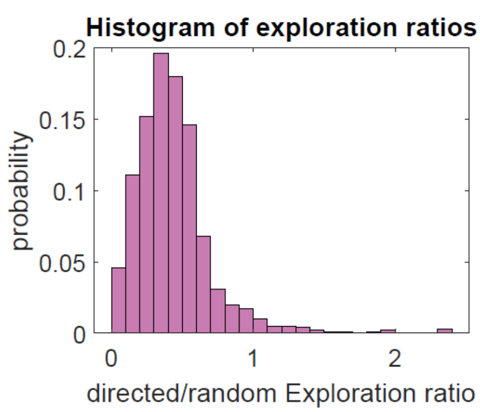

While these results suggest more random than directed exploration, we prefer not to include them in the manuscript for the following reasons: (1) they are not directly comparable to previous findings due to fundamental differences in task design; (2) they would not substantially contribute to our paper's main focus on the relationship between uncertainty processing and affective states. This question would be more appropriately addressed in future work specifically designed to dissociate these forms of exploration.

To be more specific, we provided details to support our statements as follows:

1. In Daw's 2006 paper, and Jan Peter's team's 2021 paper "Attenuated Directed Exploration during Reinforcement Learning in Gambling Disorder" and 2020 paper "Dopaminergic modulation of the exploration/exploitation trade-off in human decision-making", they used the task from Daw's 2006 paper. Although both are multi-armed restless bandit tasks, there are fundamental differences in the random walk settings between our task and Daw's 2006 design.
   From Daw et al., 2006, SI, Page1
   "The payoff for choosing the ith slot machine on trial t was between 1 and 100 points, drawn from a Gaussian distribution (standard deviation $\sigma_o = 4$) around a mean $\mu_{i,t}$ and rounded to the nearest integer. At each timestep, the means diffused in a decaying Gaussian random walk, with $\mu_{i,t}+1 = \lambda\mu_{i,t} + (1 - \lambda)\theta + v$ for each i. The decay parameter λ was 0.9836, the decay center θ was 50, and the diffusion noise v was zero-mean Gaussian (standard deviation $\sigma_d = 2.8$). "

   Our setting:
   "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. "

   Task parameter settings directly influence model selection and construction (Wilson & Collins, 2019). In papers based on Daw 2006 and Jan Peter's team's work (using the same task and model), the model was constructed as:

$$\widehat{\mu_{i,t+1}} = \lambda\widehat{\mu_{i,t}} + (1-\lambda)\theta$$

$$\widehat{\sigma^2_{i,t+1}} = \lambda^2\widehat{\sigma^2_{i,t}} + \sigma^2_d$$

μ is the mean expected value, σ is the SD of the expected value, λ is the decay rate (fixed to 0.9836), u is the decay center (fixed to 50), and σ_d^2 is the SD of the diffusion noise (fixed to 2.8)., as these parameters are from the task settings.

Update rules:

$$\widehat{\mu_{c_t,t+1}} = \widehat{\mu_{c_t,t}} + k_t\delta_t$$

Where the δ_t is the prediction error

$$\delta_t = outcome_t - \widehat{\mu_{c_t,t}}$$

And k_t is the learning rate (kalman gain)

$$k_t = \frac{\widehat{\sigma^2_{c_t,t}}}{\widehat{\sigma^2_{c_t,t}} + \sigma_o^2}$$

Specifically, $\widehat{\sigma^2_{c_t,t}}$ refers to the estimated uncertainty of the expected value of the chosen bandit, and σ_o is the observation SD, that is, the variance of the normal distribution from which payouts are drawn (**fixed to 4**).

The uncertainty of the expected value of the chosen bandit is then updated according to:

$$\widehat{\sigma^2_{c_t,t+1}} = (1 - k_t)\widehat{\sigma^2_{c_t,t}}$$

The modified softmax function is:

$$p(c_t = i) = \frac{exp\ (\beta[\widehat{\mu_{i,t}} + \varphi\widehat{\sigma_{i,t}}])}{\sum \mathrm{j}\ exp\ (\beta[\widehat{\mu_{j,t}} + \varphi\widehat{\sigma_{j,t}}])}$$

where the φ implements directed exploration.

However, in our paper, as described in **Main Text and Method S3**, we use the Kalman filter approach of Piray and Daw (2020)(Piray & Daw, 2020), a model better suited for broader reinforcement learning tasks. We cannot define observational noise, diffusion noise, decay rate, and decay center parameters like these three papers.

2. Our paper's purpose is to explore how people's weighting of two types of uncertainty relates to their mental states and associated exploratory behavior
3. Currently, no direct evidence indicates whether higher weighting of volatility corresponds to more directed exploration. To verify such a relationship would require (1) using the same 4-armed-restless bandit task as Daw et al., 2006 (2) implementing both their original model and their 2019 PCB model (3) analyzing parameters' correlations

   This is indeed an interesting question: while directed exploration is defined as uncertainty-guided exploration, we don't know whether it's guided by process noise (volatility) or observational noise (stochasticity). The same question applies to random exploration. To our knowledge, there's no current evidence showing direct relationships between volatility, stochasticity, and exploration.

4. Another approach would be using paradigms that introduce novel options (Hogeveen et al., 2022) to more objectively quantify directed exploration.
5. Following your suggestion and Jan Peter's team's definition of directed/random exploration (page 4)

   "Based on the best-fitting computational model, trials were classified. Exploitation trials are trials with choices of the bandit with the highest sum of expected value, uncertainty bonus and perseveration bonus (i.e., the highest softmax probability). Exploration trials are all other trials. These were further subdivided into trials on which participants selected the bandit with the highest exploration bonus (directed exploration trials) and all other trials (random exploration trials)"

   We attempted similar analyses by adding a *beta1* parameter to the softmax to capture directed exploration. However, this model was not optimal, with a BIC of 399056. Here, we plotted the distribution of directed/random exploration ratios across participants.

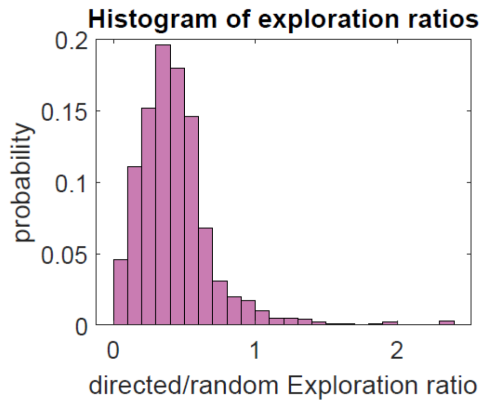

In summary, we prefer not to incorporate these results into our current paper's framework. Instead, we appreciate this thoughtful suggestion as it points to important questions for future research investigating how different forms of exploration relate to uncertainty processing and affective states.

→15. The group categorization section, where the top and bottom 25% of scores are binned, isn't justified. It appears to just be reporting the same results from the analyses with continuous questionnaire scores, but with an artificial dichotomization of the data.

Responses:

We agree with your opinion about the limitations of dichotomizing continuous variables. Our primary analyses appropriately use continuous measures of anxiety and apathy, with the correlational results reported throughout the paper. Given this feedback and a similar concern from Reviewer 2, we have moved the categorical analyses to the **Table S6 and Text S3**.

We kept the violin plots strictly for visualization.

→16. Cross-sectional mediation analyses make causal claims that cannot be justified by the present data.

We have revised the manuscript to reflect this limitation:

Revisions in the paper:

"To examine potential associations between individual differences in the perception of uncertainty, exploratory behavior, and affect, we conducted a mediation analysis (see Supplement Method S10) using anxiety, switching after reward omission (P(switch | 0)), and v/s. The results suggest that v/s may partially account for the relationship between anxiety and the tendency to switch after receiving no reward (Figure 4B). Similar patterns were observed for the analogous HMM model-based measures (see Figure S9). No significant mediation effect was found for apathy."

Also we revised the tone in Discussion, Pargraph01.

"Our mediation analysis suggests that the perception of volatility relative to stochasticity may be associated with the relationship between anxiety and exploratory behavior after reward omission. The apparent relationship between anxiety and a higher weighting of volatility relative to stochasticity may be linked to increased information-seeking behavior. This could potentially reflect a strategy aimed at reducing uncertainty and managing perceived risks more effectively"

→17.1 The two dimensions "found" by the UMAP seem to reflect differences in choice, not learning, behavior that are not captured by the models here

Responses:

We believe this comment is based on a misinterpretation and we appreciate the opportunity to clarify it. The UMAP analysis is indeed sensitive to learning, as explained below.

1. Temporal learning dynamics

The behavioral sequences used in our UMAP analysis ({choice_{t-1}, outcome_{t-1}, choice_{t}}) inherently capture learning dynamics by incorporating both previous outcomes and subsequent choices. These sequences reflect how participants learn from feedback and adjust their decisions accordingly, not just their raw choice patterns.

2. Integration with computational parameters

The strong correlation between dimension 2 and the v/s ratio ($r = -0.72$, $p<10^{-185}$) demonstrates that this dimension captures computational aspects of learning. The v/s ratio represents how participants learn from and integrate feedback over time. This is not merely a choice metric but rather reflects the underlying learning process that guides those choices.

3. Dissociable relationships

Importantly, we found that dimension 1 correlates strongly with exploratory behavior (**r = -0.90**) but not with v/s (r = 0.03), while dimension 2 shows the opposite pattern (r = -0.19 with exploration, **r = -0.72 with v/s**). This double dissociation suggests that these dimensions capture distinct aspects of learning and decision-making, rather than just choice behavior.

→17.2 specifically, whether the ratio of volatility to stochasticity is related to increased versus decreased exploration.

Responses:

Our analysis reveals that the relationship between the ratio of volatility to stochasticity (*v/s*) and exploration is more complex than a simple linear correlation. Specifically, we found:

1. Non-linear relationship

As shown in **Figure S13**, we found both significant linear and quadratic relationships between *v/s* and exploration (linear term, coefficient = 0.03, SE = 0.005, $t(996)=5.69$, $p<10^{-8}$; quadratic term, coefficient = 0.009, SE = 13×10^{-4}, $t(996)=6.948$, $p<10^{-11}$). This non-linear relationship helps explain why the correlation appears different in different parts of the manifold.

2. Context-dependent effects

When we divided the manifold into two groups based on dimension 1 scores (as shown in Figure 5H), we found that the relationship between v/s and exploration differs between groups:

2.1 In the monotonically decreasing group (individuals with high anxiety & low apathy), higher v/s correlates with decreased exploration

2.2 In the monotonically increasing group (individuals with low anxiety & high apathy), higher v/s correlates with increased exploration

→a. The claim that the UMAP analysis has found the "latent structure of adaptive behavior" is an overly strong claim when based on a few measures from one behavioral task.

Revision in the paper

We have revised the manuscript to limit the scope of our interpretation of our UMAP results.
original version: "Our results showed that exploration and uncertainty estimation related closely to the two axes of a parabolic latent structure of adaptive behavior "
revised version: "Our results showed that exploration and uncertainty estimation related closely to the two axes of a parabolic latent structure of the explore-exploit trade-off in our task."

→18. It would be a useful validation of the UMAP results (especially given that UMAP finds structure regardless of whether structure actually exists) to examine these results in other ways.

Responses:
We agree that additional validation is warranted, and we have already taken several steps to confirm our findings:

PCA and t-SNE lead to the same results

We applied multiple methods including t-SNE and PCA to the same behavioral data (see **Figure S11 and Table S7**). All methods revealed similar manifold structures and correlational patterns, suggesting the structure is robust to the specific dimensionality reduction technique used.

Simulation validation: different decision strategies associate with distinct low-dimensional shape

We simulated 1001 agents' decision strategies in the three-armed restless bandit task including

(1) Ideal agent: Agent always selects the option with the highest reward probability (optimal choice based on perfect knowledge of the reward probabilities)

(2) Random choice: Agent randomly selects among the three options with equal probability

(3) Pure staying: Agent consistently chooses the same option

(4) Pure switching: Agent switches to a different option on every trial

(5) Win-switch-lose-stay: Agent maintains the same choice after no reward and switches after reward

(6) Win-stay-lose-switch: Agent maintains the same choice after reward and switches after no reward

Then we followed the same method to obtain all possible sequences of choices and rewards. Applying UMAP, we found each strategy produced distinctly different manifold structures (Figure S14), confirming that UMAP does not simply impose artificial structure in this case but rather reveals meaningful patterns when they exist in the data. The win-stay-lose-switch manifold resembled the real human data manifold most closely, which is consistent with the observation that a win-stay-lose-switch strategy is a good first-order approximation of the behavior.

Revision in the paper

We added **Figure S14** to describe the simulation results.

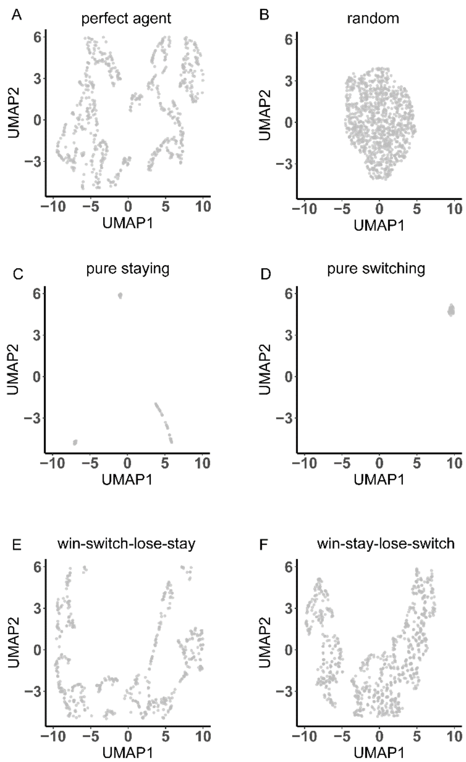

Figure S14. UMAP structures for simulated decision strategies. Low-dimensional representations of different simulated decision strategies reveal distinct patterns. (A) Perfect agent. (B) Random choice. (C) Pure staying. (D) Pure switching shows systematic alternation between options. (E) Win-switch lose-stay (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.

→For example, if anxiety/apathy are correlated with the UMAP dimensions, do the relationships between v/s and p(explore) differ in high vs. low apathy or anxiety groups?

Responses:
If we used the traditional quartile splits (or group mean±1SD) to split high/low anxiety and high/low apathy groups separately, the higher v/s consistently predicts increased exploration regardless of group membership.

We believe this result demonstrates that the traditional analysis examining separate anxiety and apathy groups may not provide a full picture of the results. In contrast, our manifold analysis integrates traits, behaviors, and model parameters to uncover two distinct subgroups with opposing patterns between v/s and P(explore).

→19. "Within the decreasing group, higher perceived volatility correlates with reduced exploration. Conversely, in the increasing group, an increased perception of volatility tends to stimulate more exploratory actions" (p. 17) - couldn't the opposite also be true, that more exploration could increase the perception of volatility?'

Responses:
The reviewer is correct, we cannot know the direction of causality. We changed the tone and revised the related text in Results, Figure caption (Figure 5H) and Discussion (paragraph03) (see ***Revision in the paper*** below)

→specifically, whether the ratio of volatility to stochasticity is related to increased versus decreased exploration.(from comments#17)

Responses:
v/s has a quadratic relationship with exploration, such that for the majority of participants, higher v/s predicts increased exploration, for a smaller subset of individuals, higher v/s is associated with decreased exploration (see SI **FigureS12**).

Revision in the paper

1. Revised the results interpretation (to address Comments#19)
revised version:
"To better understand the structure of the manifold, we examined its relationship with affective

states and exploration patterns. Using a critical dimension 1 score of -0.671 as the dividing point (see Method S12 and Figure S12), we identified two distinct groups: a monotonically decreasing group (N=390) and a monotonically increasing group (N=611). These groups showed markedly different characteristics. The decreasing group exhibited higher overall exploration rates and was characterized by slightly higher anxiety levels ($t(999)=2.08$, $p=0.037$) and lower apathy levels ($t(999)=-3.56$, $p=0.0003$), with higher v/s ratios associated with decreased exploration. In contrast, the increasing group showed lower overall exploration rates, lower anxiety levels, and higher apathy levels, with higher v/s ratios correlating with increased exploration.

These patterns reveal complexity that is not captured by traditional analyses. When using simple quartile splits of anxiety and apathy groups, higher v/s consistently predicts increased exploration regardless of group membership. The manifold approach, however, integrates multiple behavioral aspects, including affective states (anxiety and apathy), exploration behaviors, and uncertainty processing (v/s ratio), allowing us to identify interaction patterns that would be missed when examining each factor in isolation. This integration also captures the non-linear relationship between v/s and exploration (Figure S13)"

And we moved the paragraph:

"It is worth noting that we only found linear relationships between apathy, anxiety, and exploration, as well as between these affective states and the ratio of volatility to stochasticity (our analysis using higher order effects among these variables did not yield significant results, more details can be found in Table S8)." After the revised paragraph, to further clarify, only the v/s and exploration have a quadratic relationship.

<u>2. Revised the figure caption for Figure 5H to be more clear</u>

<u>revised version</u>

"(H) We divided the manifold into the monotonically decreasing group (the most left panel) and monotonically increasing group (the most right panel). The decreasing group exhibited higher overall exploration rates and was characterized by slightly higher anxiety levels and lower apathy levels. Within this group, higher v/s ratios were associated with decreased exploration. In contrast, the increasing group showed lower overall exploration rates, lower anxiety levels, and higher apathy levels, with higher v/s ratios correlating with increased exploration."

<u>3. Revised the Discussion, paragraph 03</u>

<u>revised version</u>

"Segmenting the data on the manifold further illuminated the fine-grained interplay between affective states and exploratory behavior. The monotonically decreasing group (N=390, Figure 5H), characterized by relatively higher anxiety and lower apathy, showed higher overall exploration rates compared to the monotonically increasing group (N=611) who had lower anxiety and higher apathy. Intriguingly, these groups exhibited opposing relationships between uncertainty estimation and exploration: within the decreasing group (left), higher v/s ratios were associated with decreased exploration, while within the increasing group (right), higher v/s ratios predicted increased exploration. This pattern suggests that the relationship between uncertainty estimation and exploratory behavior depends on an individual's mixed mental states. The shape from UMAP captures the non-linear relationship between the ratio of volatility to stochasticity and exploration (Figure S13), raising important questions about how environmental volatility and stochasticity might affect exploration, and its implications for mental health."

→20. The interpretation of Figure 5H ("This exploration serves as a coping strategy to relieve anxious feelings in the environment") goes beyond what can be supported by the data.

Revision in the paper

We've deleted this sentence in the figure caption for Figure 5H

Reference

Ahn, W.-Y., Krawitz, A., Kim, W., Busmeyer, J. R., & Brown, J. W. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics*, *4*(2), 95–110.

Bailey, R. A. (2005). Designing experiments and analyzing data: A model comparison perspective, 2nd edn. *Journal of the Royal Statistical Society. Series A, (Statistics in Society)*, *168*(3), 634–635.

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.

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.

Chakroun, K., Mathar, D., Wiehler, A., Ganzer, F., & Peters, J. (2020). Dopaminergic modulation of the exploration/exploitation trade-off in human decision-making. *ELife*, *9*. https://doi.org/10.7554/eLife.51260

Charpentier, C. J., Cogliati Dezza, I., Vellani, V., Globig, L. K., Gädeke, M., & Sharot, T. (2022). Anxiety increases information-seeking in response to large changes. *Scientific Reports*, *12*(1), 7385.

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, P. Y., & Popovich, P. M. (2002). *Correlation: Parametric and nonparametric measures*. https://scholar.google.com/citations?user=wGqeIeQAAAAJ&hl=en&oi=sra

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876–879.

Ebitz, R. B., Albarran, E., & Moore, T. (01/2018). 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.

Fahed, M., & Steffens, D. C. (2021). Apathy: Neurobiology, assessment and treatment. *Clinical Psychopharmacology and Neuroscience: The Official Scientific Journal of the Korean College of Neuropsychopharmacology*, *19*(2), 181–189.

Fan, H., Gershman, S. J., & Phelps, E. A. (2022). Trait somatic anxiety is associated with reduced directed exploration and underestimation of uncertainty. *Nature Human Behaviour*, 1–12.

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.

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.

Gelman, A. (2003). A Bayesian formulation of exploratory data analysis and goodness-of-fit testing. *Revue Internationale de Statistique [International Statistical Review]*, *71*(2), 369–382.

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.

Hogeveen, J., Mullins, T. S., Romero, J. D., Eversole, E., Rogge-Obando, K., Mayer, A. R., & Costa, V. D. (2022). The neurocomputational bases of explore-exploit decision-making. *Neuron*, *110*(11), 1869-1879.e5.

Karvelis, P., Paulus, M. P., & Diaconescu, A. O. (2023). Individual differences in computational psychiatry: A review of current challenges. *Neuroscience and Biobehavioral Reviews*, *148*(105137), 105137.

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

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.

Nakagawa, S., & Cuthill, I. C. (2007). Effect size, confidence interval and statistical significance: a practical guide for biologists. *Biological Reviews of the Cambridge Philosophical Society*, *82*(4), 591–605.

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.

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=DHIsuVnduwnFI3SeAY-zL2_Jw78

Sapey-Triomphe, L.-A., Pattyn, L., Weilnhammer, V., Sterzer, P., & Wagemans, J. (2023). Neural correlates of hierarchical predictive processes in autistic adults. *Nature Communications*, *14*(1), 3640.

Scholl, J., Trier, H. A., Rushworth, M. F. S., & Kolling, N. (2022). The effect of apathy and compulsivity on planning and stopping in sequential decision-making. *PLoS Biology*, *20*(3), e3001566.

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

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.

Swart, J. C., Froböse, M. I., Cook, J. L., Geurts, D. E., Frank, M. J., Cools, R., & den Ouden, H. E. (2017). Catecholaminergic challenge uncovers distinct Pavlovian and instrumental mechanisms of motivated (in)action. *ELife*, *6*. https://doi.org/10.7554/eLife.22169

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.

Wiehler, A., Chakroun, K., & Peters, J. (2021). Attenuated directed exploration during reinforcement learning in gambling disorder. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *41*(11), 2512–2522.

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

Zhang, L., & Gläscher, J. (2020). A brain network supporting social influences in human decision-making. *Science Advances*, *6*(34), eabb4159.

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.