Biological Psychiatry: CNNI

Archival Report

Impaired Goal-Directed Planning in Transdiagnostic Compulsivity Is Explained by Uncertainty About Learned Task Structure

Sirichat Sookud, Ingrid Martin, Claire M. Gillan, and Toby Wise

ABSTRACT

BACKGROUND: Diminished use of goal-directed (model-based) decision making is a hallmark of transdiagnostic compulsivity, promoting an overreliance on inflexible and habitual behaviors. However, the origin of this impairment remains unclear. Here, we tested the hypothesis that these impairments arise due to uncertainty within the internal world model that subserves goal-directed decision making.

METHODS: We adapted a validated gamified decision-making task to characterize how individuals build an internal model of an environment, pairing these data with computational modeling to uncover the exact mechanisms underpinning behavior and quantify individual differences. Two samples of participants (a discovery sample and a preregistered replication sample, n = 551 and 1322, respectively) performed the task, and we also acquired longitudinal data over 2-week and 3-month periods to assess task reliability and stability of behavior over time.

RESULTS: In the discovery and replication samples, we found that individuals higher in compulsivity and intrusive thought learned more slowly and formed a less certain representation of the task's structure. This uncertainty mediated the link between compulsive symptoms and use of goal-directed behavior. Behavior in the task was relatively stable over a 3-month (n = 385) and 1-year (n = 326) period and did not predict changes in symptoms.

CONCLUSIONS: Our results suggest that reliance on habitual behaviors seen in individuals with high levels of compulsive symptoms result from a tendency to form less certain internal models of the external world. Given the stability of this behavior and its links to symptoms, this may represent a trait-level vulnerability for this symptom dimension.

https://doi.org/10.1016/j.bpsc.2025.10.005

After a long day of work, you immediately head to the freezer and take out a frozen pizza, almost as if without thinking. If your goal is to lower your cholesterol, you might have chosen to make a salad instead, a decision guided by your ultimate goal of improving your health. This contrasts with your usual pizza routine, a simpler choice that does not account for longer-term health considerations. However, making this kind of goal-directed decision (accounting for distal outcomes) is only possible if you appreciate that eating healthily will reduce your cholesterol, ultimately improving your health; in essence, you must possess an internal model of how your actions lead to more distant consequences (1).

Impairments in goal-directed decision making are implicated in multiple psychiatric conditions, suggesting that symptoms in disorders such as obsessive-compulsive disorder (OCD) and addiction may arise from an overdependence on habitual (model-free) processes in place of goal-directed (model-based) strategies. For example, an individual with OCD may find themselves repeatedly checking whether they have locked the door due to this behavior having acquired model-free value rather than considering the long-term consequences that this behavior may cause. More recently, efforts to characterize transdiagnostic mechanisms have clarified this preponderance of results across disorders (2); rather than being tied to a diagnostic category, impairments in goal-directed decision making appear to underlie symptoms of compulsivity and intrusive thought patterns that are present across multiple diagnostic categories (3,4).

It has been suggested that this reliance on model-free decision making may arise due to uncertainty in an internal model of the world (5–7); if I am uncertain about the long-run consequences of my actions, I should be less likely to use a decision-making strategy that relies on this knowledge. While uncertainty is multifaceted, here we focus on reducible uncertainty (also referred to as estimation uncertainty), corresponding to our degree of knowledge when making a prediction. For example, I may currently believe that there is a 75% chance of rain today. Because this is not 100% or 0%, there is uncertainty. This kind of uncertainty is known as irreducible uncertainty; some events are inherently stochastic, precluding us from making entirely accurate predictions. However, perhaps I have only checked a single forecast and therefore am uncertain whether the true probability is 75%. This is reducible uncertainty; I could acquire additional information and become more confident about the true probability.

ISSN: 2451-9022 Biological Psychiatry: Cognitive Neuroscience and Neuroimaging March 2026; 11:357–365 www.sobp.org/BPCNNI

Human behavior follows this basic principle (8), and there is accumulating evidence that compulsivity symptoms are associated with deficits in learning an internal world model (9,10). While there is some evidence that compulsive individuals specifically overestimate uncertainty regarding task structure changeability (5), to our knowledge, it has not yet been established whether uncertainty more generally underlies habitual decision making.

Large-scale online citizen science studies now enable us to answer these questions with greater precision and robustness than was previously possible (2,11,12). Mental health symptoms are highly multifactorial in their etiology, meaning that any individual influence will likely have a subtle effect; thus, large samples are vital to achieve sufficient statistical power (2), rendering large-scale online data collection highly valuable. Concurrently, efforts to gamify decision-making tasks for mass online deployment encourages greater task engagement (3,13) and naturalism (14,15), improving our ability to detect subtle relationships and enabling research to move beyond highly selected samples and into the general population.

Here, we tested the hypothesis that a reliance on model-free decision making linked to transdiagnostic compulsivity arises due to uncertainty about learned task structure. We exploited a gamified task that indexes participants' ability to learn task structure combined with computational modeling that permits precise quantification of individuals' estimates of reducible uncertainty throughout the task. This is achieved by learning a distribution over transition probabilities, the variance of which serves as a measure of uncertainty. We demonstrate that transdiagnostic compulsivity is linked to heightened uncertainty about task structure and further suggest that this uncertainty mediates the link between compulsivity and reliance on habitual decision making.

METHODS AND MATERIALS

Ethical Approval

The study was approved by the King's College London Research Ethics Committee (LRS-21/22-31704).

Sample

Participants were recruited through Prolific (16) and were selected based on having a 95% approval rate, being based in the United Kingdom, and being age ≥18 years. Participants received compensation of £9/hour. We recruited 600 participants for the discovery sample, followed by 1400 for the replication sample (based on power calculations derived from the discovery sample). Of the discovery sample, 400 participants completed a second series of assessments approximately 3 months later, and 344 participants completed assessments for a third time approximately 1 year after that. Of the replication sample, 80 participants completed the task a second time approximately 2 weeks after their initial participation to allow us to assess test-retest reliability.

Exclusion Criteria

Participants failing any attention checks or failing more than 1 infrequency item check (17) were excluded from analyses. We also excluded individuals reporting a gender other than male or female from analyses that included gender as a covariate as we had insufficient numbers to properly estimate these effects.

Task

Modified Cannon Blast Task. The primary task was a modified version of the Cannon Blast task (3), which is a gamified version of the 2-step decision-making task for assessing goal-directed versus habitual decision making (18). The original task asks participants to make decisions about first-stage options based on their likelihood of transitioning to second-stage states, which are associated rewards that vary over the course of the task according to a random walk and must be learned. A model-based decision maker will use the knowledge of transition probabilities to determine the best first-stage option based on its likelihood of transitioning to a rewarding second-stage state. A model-free decision maker will simply rely on the rewards that follow their first-stage choices, ignoring the structure of the task.

In our task, participants were tasked with firing cannon balls to hit moving aliens (Figure 1A). A hit provided 100 points. During a standard trial, participants operated a cannon to choose the firing direction and selected between 2 containers, each holding purple and pink balls. Balls had a predetermined probability of being unusable on each trial, meaning that they would explode before reaching their target, and 100 points would be lost. These probabilities varied throughout the task but were displayed on screen via 2 colored bars. Probabilities were high to prevent frustration and disengagement and were anticorrelated. We also used rapid switches between reward probabilities to make model-free strategies less effective (because the best option changed frequently), as we wanted to encourage model-based planning to facilitate exploration of uncertainty in the learning of task structure. To maximize success, participants should choose the color that they see has the best chance of being usable. Thus, the container represents the first-stage choice, and each container has a probability of transitioning to a second-stage state represented by the ball color. Each of these then has a probability of reward (i.e., not exploding).

The key manipulation was that the ball color was not shown to participants in advance of their choice. Instead, each container had a varying probability of firing pink or purple balls (Figure 1E), which participants had to learn. This contrasts with the original 2-step task, where participants are provided with this information. Participants were informed that the probabilities for the left and right containers were inversely related, such that if the left container had a 90% chance of pink, the right container had a 90% chance of purple. This constitutes the structure of the task; participants needed to learn an internal model of the likelihood of each color given their choice of container and combine this with their knowledge of explosion probabilities to determine the best option on each trial. Therefore, this constitutes an inverted variant of the original task; in the original task, participants are provided with the transition probabilities but must learn the rewards associated with second-stage states, whereas here they are

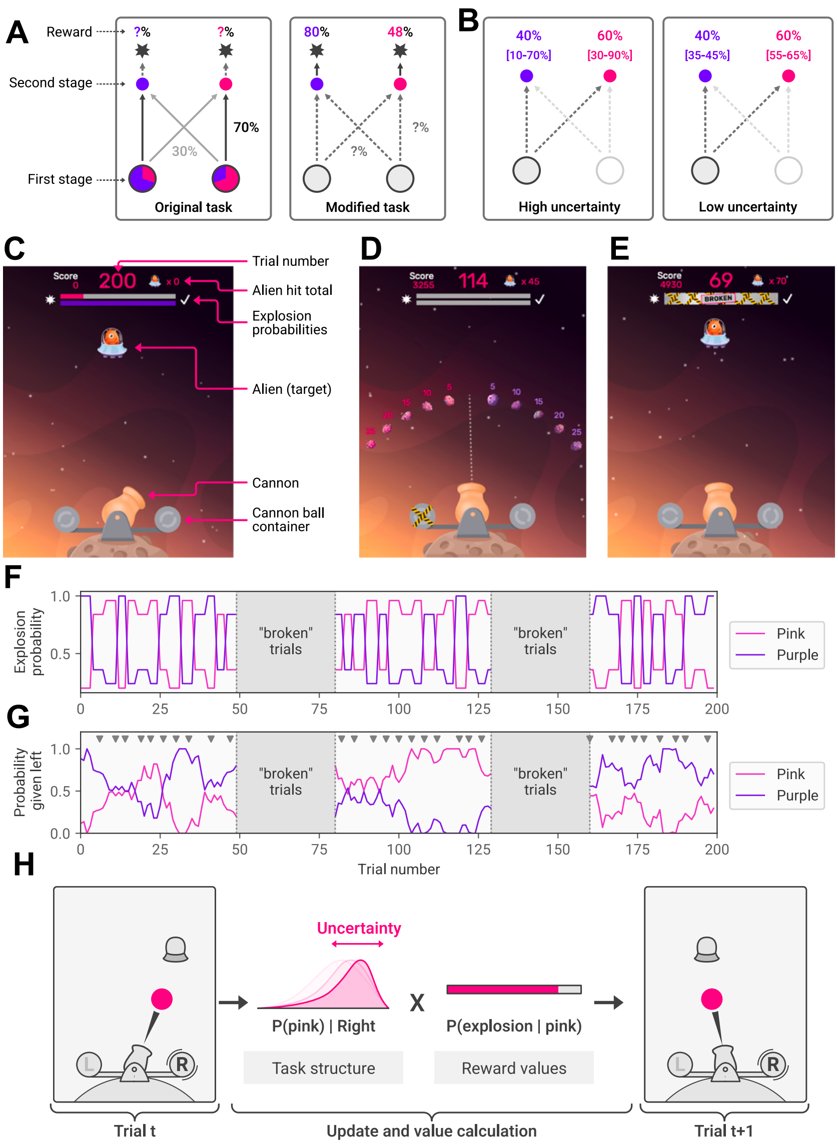

Figure 1. Task overview. **(A)** Schematic of the original and modified Cannon Blast task, a gamified implementation of the 2-step task. First-stage states (cannon ball containers) have a fixed and shown probability of transitioning to second-stage states (ball color), and participants must learn the probability of each color ball exploding. In the modified task, participants are shown the probabilities of explosion for each ball color but must learn the probability of first- to second-stage state transitions. **(B)** Estimated transition probabilities can be associated with different levels of reducible uncertainty, represented here through confidence intervals around an estimated probability. **(C)** Screenshot of a standard trial in the game. Participants were tasked with firing cannon balls from the cannon to hit an alien that moved around the top portion of the screen. They could select to fire a ball from 1 of 2 containers, each of which had a varying but hidden probability of producing pink or purple balls. The probability of each ball color exploding before reaching its target was shown by colored bars at the top of the screen. **(D)** A confidence trial. Participants were tasked with shooting at one of a range of pink or purple asteroids with varying values. If the color of the ball matched the color of the asteroid, they were awarded the number of points shown next to that asteroid, but they lost that number of points if the color did not match. **(E)** A broken trial, where all information regarding ball color was hidden. **(F)** Explosion probabilities for pink and purple balls changed rapidly across the task. All information regarding ball color was removed in broken trials. **(G)** Probabilities of pink and purple balls being produced by the left and right containers drifted over the course of the task. Probabilities for the left container are shown, and probabilities for the right container were the inverse of these probabilities. Pointers at the top indicate confidence trials. **(H)** Overview of the computational modeling approach. Observations of ball colors following left/right choices are used to update a probability distribution over task structure, the variance of which provides an index of uncertainty. This is then combined with the provided reward probabilities to inform the next choice.

given the rewards but must learn the transition probabilities. As a result, the task enables quantification of learning about transition probabilities, whereas the original version does not.

We also included trials to assess implicit confidence in task structure, where participants were effectively asked to bet on the likelihood of a given ball color (Figure 1D; see Supplemental Methods for further information).

Finally, we included a set of trials where information about ball color was concealed (described to participants as broken trials), prompting participants to guess which containers held good balls. The purpose of this was to heighten uncertainty about the task structure by removing information about ball color for a period.

Original Cannon Blast Task. We collected data in the replication sample using the original Cannon Blast task, described in detail elsewhere (3), to determine the correspondence of measures across related tasks. See Supplemental Methods for details.

Questionnaire Measures

Symptom Measures. We measured 3 transdiagnostic factors that had been identified in previous work (4): anxious-depression, compulsivity and intrusive thought, and social withdrawal. These factors have been linked to behaviors including model-based decision making (4), metacognition (19), and aversive learning (13). We chose to focus on transdiagnostic symptoms rather than condition-specific measures because these have been more strongly associated with model-based planning in previous work (2). Following our previous work, we used a subset of items from the original measures to predict the true scores (13,20). In addition, we measured state depression and anxiety using the Patient Health Questionnaire (21) and the 7-item Generalized Anxiety Disorder assessment (22); we refer to these measures as depression and anxiety.

Computational Modeling

Model Definitions. To provide a mechanistic account of participants' behavior in the task, we designed computational models of the processes underpinning learning about task structure. These models were built on our previous work using approximately Bayesian learning strategies (13,23) to learn a probability distribution over a given outcome. A significant strength of these models is their ability to estimate uncertainty during learning and thus to provide individual-level estimates of participants' subjective uncertainty while remaining very simple (3 free parameters including decision temperature). These models were fit to all learning trials (including broken trials), excluding confidence trials. Broken trials served to manipulate uncertainty, because the model gained no information regarding transition structure during these trials; instead, its estimates were allowed to decay but not update, increasing uncertainty.

These models, along with fitting procedures, are described in full in Supplemental Methods. Briefly, we used 4 variants: 1 used only a model-based strategy based on a learned task structure, 1 used a purely model-free strategy, and 2 used either weighted or fixed combinations of model-based and model-free strategies. Our modeling approach provides 3 key measures of interest, which are described in Table 1 and Supplemental Methods.

For comparison, we also tested 3 equivalent models that do not represent uncertainty and instead use a Rescorla-Wagner learning rule rather than the approximate Bayesian beta learning rule (24).

Statistical Analysis

Test-Retest Reliability Analyses. We quantified test-retest reliability and longer-term stability using intraclass correlation coefficients (ICCs) (25). We used consistency ICCs [also known as ICCs (1,3)], which estimate reliability based on the rank order of the observations.

Table 1. Key Model Parameters (Update Rate, Model-Basedness) and Latent Variables (Mean Uncertainty)

Parameter/Latent Variable	Description
Update Rate, τ	The update rate in the model, representing the extent to which evidence for transitions (pink or purple balls) is updated on each trial. Greater values generally result in faster learning.
Mean Uncertainty	The variance of the distribution over transition probabilities, averaged across trials within participant. This represents reducible uncertainty and generally reduces as the model acquires more information.
Model-Basedness	The difference in model fit between models representing purely model-based and purely model-free decision making, providing an index of the extent to which participants use model-based (goal-directed) over model-free (putatively habitual) strategies.

Other parameters are included in the model (decay rate λ and softmax temperature) but are not of interest for our primary hypotheses. These are described in Supplemental Methods.

Modeling of Confidence Ratings. We verified that the quantities derived from the model were meaningfully related to participant's behavior on the task by establishing that confidence ratings were predicted by different types of uncertainty as quantified by our model. Full details are provided in Supplemental Methods.

Symptom-Behavior Regression Models. Relationships between symptoms and behavior were tested using linear models with *p* values estimated using bootstrapping. Full details are provided in Supplemental Methods.

Preregistration

For our replication study, we preregistered the primary hypotheses and methods based on results from the discovery sample (https://osf.io/52awb). These pertained to analyses predicting 3 core behavioral metrics (model-basedness, mean uncertainty about task structure, and transition update rate) based on transdiagnostic symptoms and state depression/anxiety.

RESULTS

Samples

Our final samples after exclusions consisted of 551 participants for the discovery sample and 1322 participants for the replication sample (see Table 2 and Figure 2 for demographic information). Of the discovery sample, 385 participants completed a second series of assessments after 3 months, and 326 participants returned a third time after 1 year; 21.05% and 23.23% scored above established cutoffs for clinically relevant symptom severity based on anxiety and depression measures, respectively, a pattern that replicated in the replication sample (18.38% and 23.6%).

Task Behavior

Participants were able to perform the task successfully, as demonstrated by increasing scores across the entire task (Figure 2F). We observed a pattern of switching between options indicative of model-based decision making in trials with a clear transition pattern ($\geq$80% chance of 1 color); as in the original 2-step task, participants switched after observing a rewarded rare transition and stayed when observing a rewarded common transition (discovery $\beta_{\text{rewarded} \times \text{common}}$ = 0.23; 95% highest density interval [HDI], 0.20–0.26; replication $\beta_{\text{rewarded} \times \text{common}}$ = 0.25; 95% HDI, 0.23–0.28).

Computational modeling revealed that participants' behavior was predominantly goal directed, being best described by a model that did not feature model-free value learning, in both samples (Figure 3A–C). Test-retest reliability analyses indicated that key model parameters exhibited test-retest reliability in the moderate range, with values of 0.4 to 0.6 (26) (Table 3). Reliability was higher over 3 months and 1 year (Table 3). The model-basedness estimate from the modified task

Table 2. Demographic Characteristics of the Samples

Sample	Sample Size	Age, Years	Gender, Female/ Male
Discovery	551	40.75 (12.32)	276/275
3-Month follow-up	385	41.80 (12.39)	187/198
1-Year follow-up	326	43.78 (12.22)	154/172
Replication	1322	41.15 (12.69)	667/655
2-Week follow-up	80	–	–

Values are presented as *n* or mean (SD). Demographic data were not available for the 2-week follow-up sample due to a technical error.

was also correlated with model-basedness from the original Cannon Blast task ($\sigma = 0.17$, $p < .001$).

To provide evidence that uncertainty plays a role in learning, we compared these models with equivalent models that learned transition probabilities using a Rescorla-Wagner learning rule, maintaining only point estimates of probabilities. Uncertainty-driven learning models outperformed these comparison models (Figure 3A–C), suggesting that participants learned transition structure in a manner that tracked uncertainty.

Confidence ratings were strongly predicted by model-derived transition expectations (probability of pink vs. purple), alongside both reducible and irreducible uncertainty (Figure 3D; see Supplemental Results for details). Model-derived reducible uncertainty was also strongly positively correlated with age, indicating that this metric is sensitive to individual differences (discovery $r = 0.20$, $p < .001$; replication $r = 0.25$, $p < .001$) (Figure 3E). Notably, an optimal model demonstrated low uncertainty relative to human participants, with 97.1% of participants having higher uncertainty than was optimal.

Relationships Between Behavior and Transdiagnostic Symptom Dimensions

In the discovery sample, we found that model-basedness in the task was predicted positively by anxious-depression ($\beta = 0.19$, $p_{corrected} < .001$) (Figure 4A) and negatively by compulsivity and intrusive thought ($\beta = -0.19$, $p_{corrected} < .001$) (Figure 4A). These effects were replicated in the replication sample (compulsivity and intrusive thought: $\beta = -0.10$, $p_{corrected} = .001$; anxious-depression: $\beta = 0.07$, $p_{corrected} = .04$). Separating this metric into its constituent parts (given that it represents the difference between model-based and model-free fit) revealed that this was driven specifically by the use of model-based decision making (see Tables S14, S15, S22, and S23). We also replicated previous findings using the original Cannon Blast task (4) in the replication sample, finding a reduction in model-basedness associated with compulsivity and intrusive thought together with an increase in model-basedness associated with anxious-depression and social withdrawal (compulsivity and intrusive thought $\beta = -0.11$, $p < .001$; anxious-depression $\beta = 0.06$, $p = .039$; social withdrawal $\beta = 0.08$, $p = .008$) (Figure 4E).

In both discovery and replication samples, compulsivity and intrusive thought was further associated with heightened uncertainty about task structure (discovery $\beta = 0.16$, $p_{corrected} = .003$; replication $\beta = 0.09$, $p_{corrected} = .009$) (Figure 4A) and more gradual updating of task transition probability estimates, as indexed by the update parameter in our model (discovery $\beta = -0.15$, $p_{corrected} = .015$; replication $\beta = -0.10$, $p_{corrected} = .002$) (Figure 4A). Anxious-depression consistently demonstrated the opposite pattern with respect to uncertainty estimates (discovery $\beta = -0.16$, $p_{corrected} = .003$, replication $\beta = -0.11$, $p_{corrected} = .001$) (Figure 4A) and update rates (discovery $\beta = 0.15$, $p_{corrected} = .012$; replication $\beta = 0.10$, $p_{corrected} = .001$) (Figure 4A).

We sought to clarify these relationships further by determining whether the relationship between compulsivity and intrusive thought and model-basedness was mediated by uncertainty. In both samples, mediation analysis revealed a significant indirect effect (path values are shown in Figure 4C for the replication sample and in the Supplement for the discovery sample) and a nonsignificant direct effect (Figure S4C), suggesting that the observed association was fully mediated by uncertainty. Notably, we also found that the association between compulsivity and intrusive thought and model-basedness

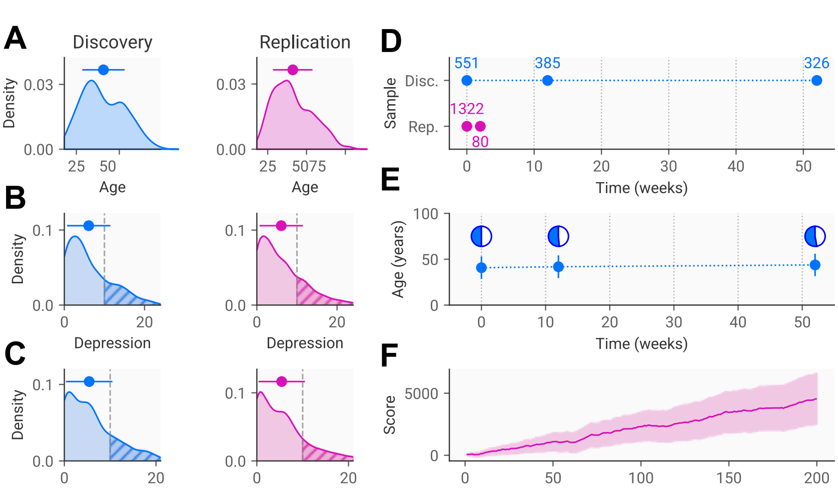

Figure 2. **(A)** Age distribution across both samples. The point above represents the mean, with the error bars showing the SD. **(B)** Depression scores across both samples, measured using the Patient Health Questionnaire-8. The dashed line represents a cutoff for identifying clinically significant depression (36), while the shaded area represents the distribution of scores above this cutoff. **(C)** Anxiety scores across both samples, measured using the 7-item Generalized Anxiety Disorder scale. The dashed line represents a cutoff for identifying clinically significant generalized anxiety disorder (22). **(D)** Illustration of sample sizes for discovery (Disc.) and replication (Rep.) samples at different time points. **(E)** Age and gender in the discovery sample across multiple time points. The points and error bars represent the mean and standard error of participants' ages. The pie charts represent the proportion of male (shaded) and female (unshaded) participants. **(F)** Participants' scores on the task in the replication sample as a function of trial. The shaded area represents the standard error.

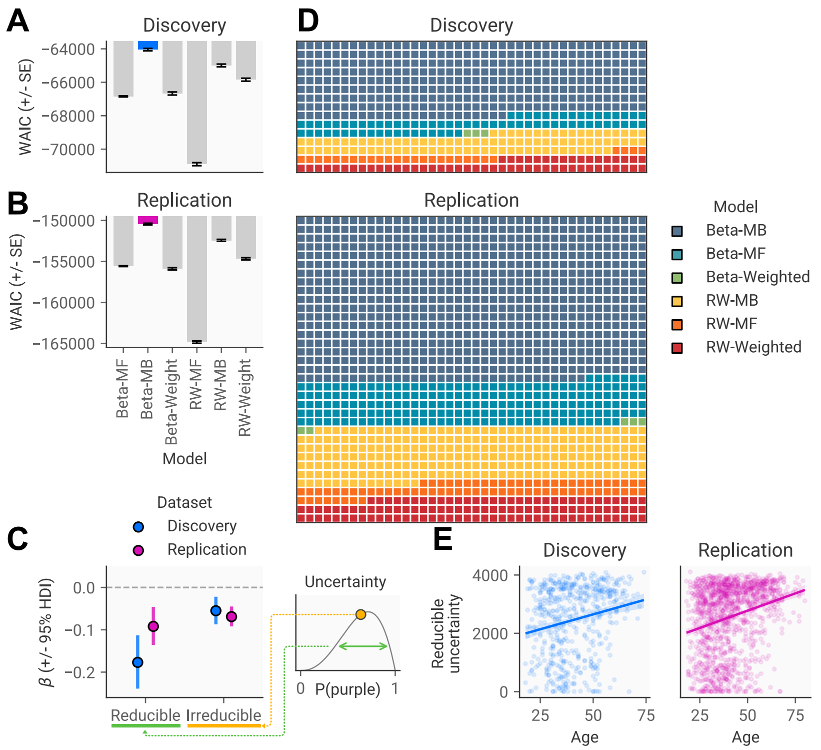

Figure 3. Computational modeling. **(A)** Model comparison results for the discovery sample, showing Watanabe-Akaike information criterion (WAIC) scores for each model (higher scores = better fit). Fit is shown for 2 families of models: beta models that represent uncertainty and Rescorla-Wagner (RW) models that do not. The winning model is highlighted in blue. **(B)** Model comparison results for the replication sample. The winning model is highlighted in pink. **(C)** The best-fitting model for every participant in both samples. **(D)** Regression coefficients from a Bayesian hierarchical model predicting implicit confidence ratings from model-derived irreducible and reducible certainty metrics. Note that the direction of the reducible uncertainty (derived from the variance) effect has been reversed for illustrative purposes such that greater uncertainty predicts lower confidence. **(E)** Associations between age and reducible uncertainty across both samples. Weight indicates parameterized weighting of model-based (MB) and model-free (MF) structures. a.u., arbitrary unit; HDI, highest density interval.

in the original Cannon Blast task was partially mediated by uncertainty estimated from the novel task (Table S30).

Effects of Anxiety and Depression Symptoms on Task Behavior

We observed a positive relationship between depressive symptoms and relative model-basedness in both the discovery ($\beta = 0.29$, $p_{corrected} < .001$) (Figure 4B) and replication ($\beta = 0.14$, $p_{corrected} = .003$) (Figure 4B) samples, with no replicable results related to anxiety. Breaking the relative measure into its component parts, depressive symptoms were specifically

Table 3. Test-Retest Reliability and Stability Estimates for Model Parameters (Update Rate, Model-Basedness) and Latent Variables (Mean Uncertainty) Over 2-Week, 3-Month, and 1-Year Intervals

Interval	Sample Size	Parameter/Latent Variable	ICC
2 Weeks	80	Update rate, τ	0.50
		Mean uncertainty	0.50
		Model-basedness	0.42
3 Months	385	Update rate, τ	0.57
		Mean uncertainty	0.57
		Model-basedness	0.60
1 Year	326	Update rate, τ	0.62
		Mean uncertainty	0.60
		Model-basedness	0.60

ICC, intraclass correlation coefficient.

associated with greater use of model-based strategies in both samples (see Supplemental Results).

We also found that higher depressive symptoms were linked to more rapid updating (discovery: $\beta = 0.35$, $p_{corrected} < .001$; replication: $\beta = 0.17$, $p_{corrected} < .001$) (Figure 4B) and lower uncertainty (discovery: $\beta = -0.32$, $p_{corrected} < .001$; replication: $\beta = -0.17$, $p_{corrected} = .002$) (Figure 4B). This association was fully mediated by uncertainty, as indicated by a significant indirect effect (Figure 4D) and a nonsignificant direct effect (Figure 4D).

Changes in Behavior Over Time

As described above, behavior was reasonably stable over time (Table 3 and Figure 5B), with subtle learning effects as demonstrated by increasing scores (Figure 5D). Symptoms exhibited even greater stability, with ICCs >0.77 (Table S1 and Figure 5A). We found no significant associations between changes in symptoms and changes in behavior across either time point when we controlled for baseline scores (Figure S5C).

DISCUSSION

Impairments in goal-directed decision making are an established correlate of compulsive symptoms (3,4,27,28). However, it has remained unclear why such impairments emerge. Here, we used a computational factor modeling approach (2) to suggest that greater uncertainty in a learned internal model of the external environment may underpin the association

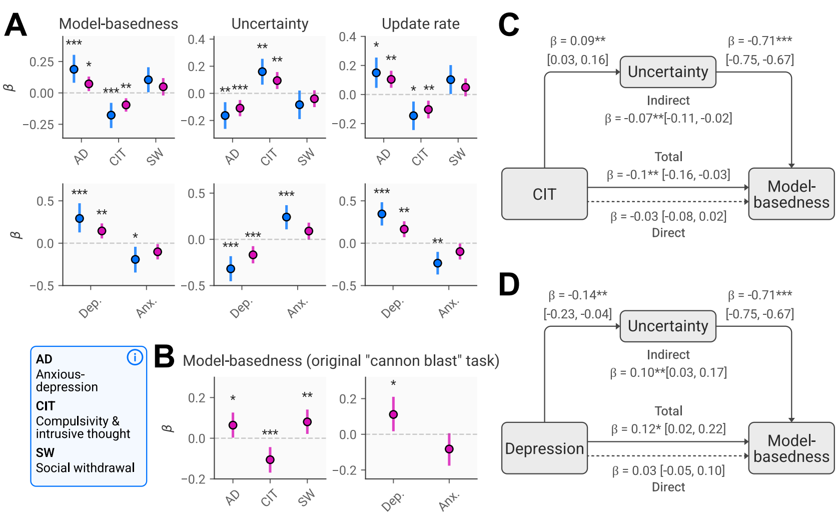

Figure 4. Associations between task behavior and symptoms. **(A)** Results from regression models predicting behavioral metrics (indicated by the title above each panel) from symptoms (shown on the x-axis) across both samples. The points represent the regression coefficient, while the error bars represent 95% CIs (estimated using bootstrapping). Stars above the points represent statistical significance, corrected for multiple comparisons. The top row represents analyses of transdiagnostic symptom dimensions, while the bottom row represents analyses of depression (Dep.) and anxiety (Anx.) symptoms (measured using the Patient Health Questionnaire-8 and 7-item Generalized Anxiety Disorder scale, respectively). **(B)** Results from regression models predicting model-basedness (MB) as assessed by the original Cannon Blast task (3) from both transdiagnostic dimensions (left) and state anxiety and depression (right). **(C)** Results of a mediation model testing mediation of the CIT > MB association by uncertainty about task structure. **(D)** Results of a mediation model testing mediation of the depression > MB association by uncertainty about task structure. Values represent regression weights (β) and confidence intervals (estimated using bootstrapping). *$p < .05$, **$p < .01$, ***$p < .001$.

between transdiagnostic compulsivity and intrusive thought and model-based planning.

We developed a novel extension of a validated, gamified 2-step decision-making task (3) that required participants to learn the structure of the task in order to perform well. Using computational modeling, we extracted a measure of model-basedness for each individual, finding in 2 samples that individuals higher in compulsivity were less model based in their decision making, as has been observed in previous studies using different tasks (3,4). Critically, our modeling also demonstrated that people higher in this symptom dimension learned the task structure more slowly and were more uncertain as a result. Mediation analyses demonstrated that the link between compulsivity and model-basedness was mediated by this uncertainty metric.

Taken together, these results suggest that impaired goal-directed decision making may emerge as a result of an uncertain internal model of the world; simply, an uncertain model is less likely to be relied on to guide behavior. As a result, behavior is likely to become rigid and dependent on habits. This is consistent with previous theoretical work (6,7) and empirical findings regarding perceptions of changeability (5) alongside evidence that successful cognitive map formation predicts use of model-based control (29). These results also echo previous findings suggesting alterations in task structure learning associated with compulsivity (9,10) but provide an important extension by providing empirical evidence regarding the specific role of uncertainty. Notably, previous investigations into value learning processes in transdiagnostic compulsivity have observed either null results (30) or associations in an opposite direction (13), suggesting that these results are specific to task structure learning. Future work could consider the role of additional factors that are known to influence cognitive map formation, such as motivation (29), in this relationship. Future work in clinical samples would also be valuable to establish the extent to which these findings translate to clinically relevant settings, and the integration of transdiagnostic approaches into clinical practice represents a broader challenge within the field.

More generally, our task provides a novel measure of goal-directed decision making and the task structure learning that underpin this process. Combined with a modeling approach adapted from previous research (13,23) that captures subjective uncertainty during learning, this method provides an

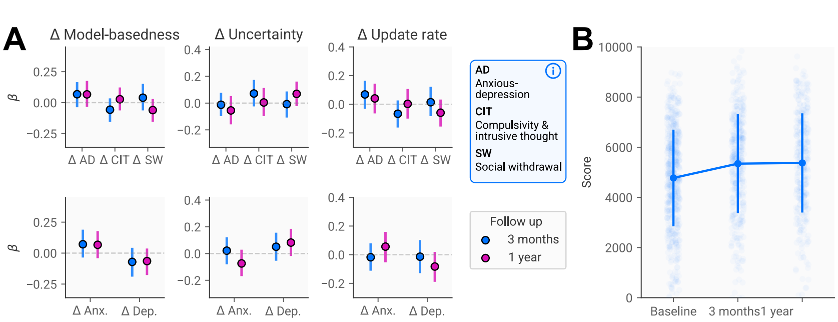

Figure 5. Longitudinal results. **(A)** Results from regression models predicting change (Δ) in behavioral metrics across the follow-up period from change in symptom scores (top) and state anxiety (Anx.) and depression (Dep.) (bottom), controlling for baseline scores. The points represent the regression coefficient, while the error bars represent 95% CIs (estimated using bootstrapping). **(B)** Participants' scores on the task across the 3 time points.

effective way to assess both decision making and the learning processes that facilitate it. We also demonstrated that implicit confidence in our task was predicted by different types of model-derived uncertainty, indicating that these quantities are not solely abstract values but have subjective meaning for participants. We also found that alternative models that did not represent uncertainty provided a poorer fit to the data, suggesting that uncertainty plays a role in learning task structure. An advantage of these models is that they allow learning to adapt to current uncertainty, such that unexpected outcomes are less influential during periods of certainty. However, we note that our models focus only on reducible (or estimation) uncertainty, because it represents the precision of a world model, and therefore should drive its use. Future work could investigate how these findings extend to our sources of uncertainty, such as environmental volatility (31).

Notably, our best-fitting model was one that relied purely on a model-based strategy (although there was variation in its fit across participants, with a substantial minority being best fit by a model-free variant). This is likely because model-free strategies performed poorly in the task due to very rapid changes in reward probabilities; the task was designed to promote model-based planning by using reward probabilities with low autocorrelation. However, we note that our model assumes that participants possess a correct, but uncertain, model of the task. An alternative explanation is that compulsivity is associated with an imperfect (but perhaps no less certain) model of the world. This was not possible to test with a simple task as was used here but would be an intriguing avenue for future work with more sensitive task designs. Furthermore, there are alternative models that may capture the pattern of dynamic learning seen here without explicitly representing uncertainty (for example, models that adjust learning rates according to prediction error magnitude). Therefore, future research should seek to confirm that the pattern of transition learning that we observed here is truly uncertainty dependent and explore alternative explanations for imperfect learning or use of world models.

Our longitudinal analyses demonstrated that both symptoms and behavior were relatively stable over time and that changes in symptoms were not predictive of changes in behavior. This may suggest that learning profiles are relatively stable markers of trait compulsivity, potentially reflecting a general vulnerability factor for more severe symptoms. This would be consistent with suggestions that deficits in model-based planning are not ameliorated following successful treatment in OCD (32). However, it is also possible that this null result arose as a result of limited variability in symptoms available in a community sample. Longer-term investigations, including in samples with more extreme symptoms, may provide greater insight into these relationships. This may also suggest that these processes are not a promising target for later intervention, but we might speculate that they could serve as a target for earlier preventive interventions. This work also has limitations. While the use of an unselected, general population sample recruited online allows for well-powered investigations (2), the symptoms and underlying mechanisms present in such a sample may be qualitatively different from those seen in clinical populations. However, 20% of the participants scored above established cutoffs for depression and anxiety. Furthermore, these is evidence that relationships between compulsive symptoms and goal-directed decision making are present in both general population and clinical samples (27), suggesting that this may not necessarily represent a significant problem. Furthermore, while the sample was not selected based on clinical status, such individuals were not excluded. We also relied on self-report symptom measures and used a brief questionnaire (20,33) rather than the full set of items originally used to derive these transdiagnostic dimensions (4), which precludes us from assessing effects related to individual questionnaire measures. Lastly, we also note that our approach relies on a somewhat simplistic definition of habitual, and the precise definition of this term is debated (34); it is possible that participants using a model-free strategy may have chosen this as a result of a goal-directed decision where the goal is to obtain points, potentially considering cognitive cost, although we note that a recent study found that compulsivity and intrusive thought were also associated with habitual responding in a task that offers a purer measure of habitual behavior (35).

ACKNOWLEDGMENTS AND DISCLOSURES

This work was supported by a Wellcome Trust Career Development Award (Grant No. 225945/Z/22/Z [to TW]).

A previous version of this article was published as a preprint on PsyArXiv: https://doi.org/10.31234/osf.io/zp6vk_v2.

Code and data to reproduce the analyses described here are available at https://github.com/the-wise-lab/transition-learning-public.

The authors report no biomedical financial interests or potential conflicts of interest.

ARTICLE INFORMATION

From the Department of Neuroimaging, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, United Kingdom (SS, IM, TW); Institute of Cognitive Neuroscience, University College London, London, United Kingdom (IM); Department of Psychology, Trinity College Dublin, Dublin, Ireland (CMG); Trinity College Institute for Neuroscience, Trinity College Dublin, Dublin, Ireland (CMG); and ADAPT Centre for Digital Technology, Trinity College Dublin, Dublin, Ireland (CMG).

Address correspondence to Toby Wise, Ph.D., at toby.wise@kcl.ac.uk.

Received Jun 13, 2025; revised Oct 1, 2025; accepted Oct 2, 2025.

Supplementary material cited in this article is available online at https://doi.org/10.1016/j.bpsc.2025.10.005.

REFERENCES

1. Drummond N, Niv Y (2020): Model-based decision making and model-free learning. Curr Biol 30:R860–R865.
2. Wise T, Robinson OJ, Gillan CM (2023): Identifying transdiagnostic mechanisms in mental health using computational factor modeling. Biol Psychiatry 93:690–703.
3. Donegan KR, Brown VM, Price RB, Gallagher E, Pringle A, Hanlon AK, Gillan CM (2023): Using smartphones to optimise and scale-up the assessment of model-based planning. Commun Psychol 1:1–15.
4. Gillan CM, Kosinski M, Whelan R, Phelps EA, Daw ND (2016): Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. Elife 5:e11305.
5. Fradkin I, Ludwig C, Eldar E, Huppert JD (2020): Doubting what you already know: Uncertainty regarding state transitions is associated with obsessive compulsive symptoms. PLoS Comput Biol 16:e1007634.
6. Fradkin I, Adams RA, Parr T, Roiser JP, Huppert JD (2020): Searching for an anchor in an unpredictable world: A computational model of obsessive compulsive disorder. Psychol Rev 127:672–699.

7. Rigoux L, Stephan KE, Petzschner FH (2024): Beliefs, compulsive behavior and reduced confidence in control. PLoS Comput Biol 20: e1012207.
8. Lee SW, Shimojo S, O'Doherty JP (2014): Neural computations underlying arbitration between model-based and model-free learning. Neuron 81:687–699.
9. Sharp PB, Dolan RJ, Eldar E (2023): Disrupted state transition learning as a computational marker of compulsivity. Psychol Med 53:2095–2105.
10. Seow TXF, Benoit E, Dempsey C, Jennings M, Maxwell A, O'Connell R, Gillan CM (2021): Model-based planning deficits in compulsivity are linked to faulty neural representations of task structure. J Neurosci 41:6539–6550.
11. Gillan CM, Rutledge RB (2021): Smartphones and the neuroscience of mental health. Annu Rev Neurosci 44:129–151.
12. Gillan CM, Daw ND (2016): Taking psychiatry research online. Neuron 91:19–23.
13. Wise T, Dolan RJ (2020): Associations between aversive learning processes and transdiagnostic psychiatric symptoms in a general population sample. Nat Commun 11:4179.
14. Mobbs D, Wise T, Suthana N, Guzmán N, Kriegeskorte N, Leibo JZ (2021): Promises and challenges of human computational ethology. Neuron 109:2224–2238.
15. Wise T, Emery K, Radulescu A (2024): Naturalistic reinforcement learning. Trends Cogn Sci 28:144–158.
16. Palan S, Schitter C (2018): Prolific.ac—A subject pool for online experiments. J Behav Exp Fin 17:22–27.
17. Zorowitz S, Solis J, Niv Y, Bennett D (2023): Inattentive responding can induce spurious associations between task behaviour and symptom measures. Nat Hum Behav 7:1667–1681.
18. Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ (2011): Model-based influences on humans' choices and striatal prediction errors. Neuron 69:1204–1215.
19. Rouault M, Seow T, Gillan CM, Fleming SM (2018): Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance. Biol Psychiatry 84:443–451.
20. Hopkins AK, Gillan C, Roiser J, Wise T, Sidarus N (2022): Optimising the measurement of anxious-depressive, compulsivity and intrusive thought and social withdrawal transdiagnostic symptom dimensions. PsyArXiv https://doi.org/10.31234/osf.io/q83sh.
21. Kroenke K, Spitzer RL, Williams JBW (2001): The PHQ-9: Validity of a brief depression severity measure. J Gen Intern Med 16:606–613.
22. Spitzer RL, Kroenke K, Williams JBW, Löwe B (2006): A brief measure for assessing generalized anxiety disorder: The GAD-7. Arch Intern Med 166:1092–1097.
23. Wise T, Michely J, Dayan P, Dolan RJ (2019): A computational account of threat-related attentional bias. PLoS Comput Biol 15: e1007341.
24. Rescorla RA, Wagner AR (1972): A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black AH, Prokasy WF, editors. Classical conditioning II: Current research and theory. New York: Appleton- Century-Crofts, 64–99.
25. Weir JP (2005): Quantifying test-retest reliability using the intraclass correlation coefficient and the SEM. J Strength Cond Res 19:231–240.
26. Koo TK, Li MY (2016): A guideline of selecting and reporting intraclass correlation coefficients for reliability research. J Chiropr Med 15:155–163.
27. Gillan CM, Kalanthroff E, Evans M, Weingarden HM, Jacoby RJ, Gershkovich M, *et al.* (2020): Comparison of the association between goal-directed planning and self-reported compulsivity vs obsessive-compulsive disorder diagnosis. JAMA Psychiatry 77:77–85.
28. Brown VM, Chen J, Gillan CM, Price RB (2020): Improving the reliability of computational analyses: Model-based planning and its relationship with compulsivity. Biol Psychiatry Cogn Neurosci Neuroimaging 5:601–609.
29. Karagoz AB, Reagh ZM, Kool W (2024): The construction and use of cognitive maps in model-based control. J Exp Psychol Gen 153:372–385.
30. Lee JK, Rouault M, Wyart V (2023): Compulsivity is linked to maladaptive choice variability but unaltered reinforcement learning under uncertainty. bioRxiv https://doi.org/10.1101/2023.01.05.522867.
31. Piray P, Daw ND (2024): Computational processes of simultaneous learning of stochasticity and volatility in humans. Nat Commun 15:9073.
32. Wheaton MG, Gillan CM, Simpson HB (2019): Does cognitive-behavioral therapy affect goal-directed planning in obsessive-compulsive disorder? Psychiatry Res 273:94–99.
33. Wise T, Sidarus N (2025): Reducing the burden of psychological questionnaire measures through selective item re-weighting. R Soc Open Sci 12:241857.
34. Miller KJ, Shenhav A, Ludvig EA (2019): Habits without values. Psychol Rev 126:292–311.
35. Donegan KR, Suddell S, De Forest S, Sun J, Gallagher E, Hanlon A, *et al.* (2025): Compulsivity is associated with accelerated stimulus-response habit learning. PsyArXiv https://doi.org/10.31234/osf.io/xwtsu_v1.
36. Wu Y, Levis B, Riehm KE, Saadat N, Levis AW, Azar M, *et al.* (2020): Equivalency of the diagnostic accuracy of the PHQ-8 and PHQ-9: A systematic review and individual participant data meta-analysis. Psychol Med 50:1368–1380.