Associations between aversive learning processes and transdiagnostic psychiatric symptoms revealed by large-scale phenotyping

Toby Wise[1,2,3] & Raymond J Dolan[1,2]

[1]Wellcome Centre for Human Neuroimaging, University College London, London, UK
[2]Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London, UK
[3]Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA, USA

Abstract

Background

Aversive learning processes are a candidate source of dysfunction in psychiatric disorders. Here symptom expression in a range of conditions is linked to altered threat perception, manifesting particularly in uncertain environments. How precise computational mechanisms that support aversive learning, and uncertainty estimation, relate to the presence of specific psychiatric symptoms remains undetermined.

Methods

400 subjects completed a novel online game-based aversive learning task, requiring avoidance of negative outcomes, in conjunction with completing measures of common psychiatric symptoms. We used a probabilistic computational model to measure distinct processes involved in learning, in addition to inferred estimates of safety likelihood and uncertainty. We tested for associations between learning processes and traditional psychiatric constructs alongside transdiagnostic factors using linear models. We used partial least squares regression to identify components of psychopathology grounded in both aversive learning behaviour and symptom self-report.

Results

State anxiety and a transdiagnostic compulsivity-related factor were associated with enhanced learning from safety. However, data-driven analysis using partial least squares regression indicated the presence of two separable components across our behavioural and questionnaire data: one linked enhanced safety learning and lower estimated uncertainty to physiological anxiety, compulsivity, and impulsivity; the other linked enhanced threat learning and heightened uncertainty estimation to symptoms of depression and social anxiety.

Conclusions

Our findings implicate aversive learning processes under uncertainty to the expression of psychiatric symptoms that cut across traditional diagnostic boundaries. These relationships are more complex than previously conceptualised. Future research should focus on understanding the neural mechanisms underlying alterations in aversive learning and how these lead to the development of symptoms and disorder.

Introduction

Many core symptoms of mental illness are linked to learning about unpleasant events in our environment. In particular, symptoms of mood and anxiety disorders, such as apprehension, worry, and low mood can intuitively be related to altered perception of the likelihood of aversive outcomes. Indeed, the importance of altered threat perception is a feature of many disorders that extend beyond disorders of mood to encompass conditions such as psychosis (1) and eating disorders (2). As a result, research into how individuals learn about aversive events holds great promise for enhancing our understanding across a diverse range of mental health problems.

Computational approaches are a powerful means to characterise the precise mechanisms underpinning learning, as well as uncovering how these relate to psychiatric symptom expression (3, 4). Recent studies have leveraged computational modelling to describe associations between learning processes and psychiatrically-relevant traits in non-clinical samples (5–8), as well as in clinical conditions ranging from anxiety and depression to psychosis (9–12). A common finding across studies is that of altered learning rates, where psychopathology is linked to inappropriate weighting of evidence when updating value estimates (13, 14). Notably, there is evidence suggesting that people with clinically significant symptoms of anxiety and depression show biased learning as a function of the valence of information, updating faster in response to negative than positive outcomes (12), a bias that might engender a negative view of the environment. In healthy individuals, this bias is associated with trait optimism (15). However, it is unclear to extent such biased learning relates to the spectrum of mental health problems.

One process implicated in the genesis of psychiatric disorders is that of uncertainty estimation. Uncertainty plays a fundamental role in learning, and computational formulations optimise learning in the face of non-stationary probabilistic outcomes based on uncertainty (11, 16–19). While psychiatric traits, including anxiety, are linked to an inability to adapt learning in response to environmental statistics such as volatility (5, 9), little research has investigated whether how individuals estimate, or respond to, uncertainty in aversive environments is associated with psychiatric traits. This is a critical question given that core features of anxiety revolve around a concept of uncertainty; individuals with anxiety disorders report feeling more uncertain about threat and being less comfortable in situations involving uncertainty (20–23).

Existing work on aversive learning has had a particular focus on symptoms of anxiety and depression (7, 12). However, these approaches have not been designed optimally for identifying mechanisms that span traditional diagnostic boundaries. This assumes importance in light of recent studies, using large samples, that show several aspects of learning and decision making relate more strongly to transdiagnostic factors than to any specific categorical psychiatric disorder (6, 8, 24, 25). Applying such an approach to aversive learning may yield better insights into the role of learning in psychiatric disorders.

Here, we aimed to identify specific aversive learning processes that relate to both traditional measures of anxiety, and transdiagnostic psychiatric traits, in a large sample collected online.

Specifically, we used a computational approach to test whether anxiety and transdiagnostic symptoms are associated with biased learning from safety and threat, whether these traits relate to altered estimates of threat likelihood, and whether they are associated with different levels of uncertainty during threat learning. Given difficulties in using traditional aversive stimuli in an online setting, we developed a novel game-based avoidance task designed to engage threat and avoidance processes without the need for administration of painful or noxious stimuli.

Methods and materials

Ethics

This research was approved by the University College London research ethics committee (reference 9929/003). All participants provided informed consent and were compensated financially for their time at a rate of at least £6 per hour.

Participants

We recruited 400 participants through Prolific (26). Subjects were selected based on being aged 18-65 and having at least a 90% approval rate across studies they had previously participated in.

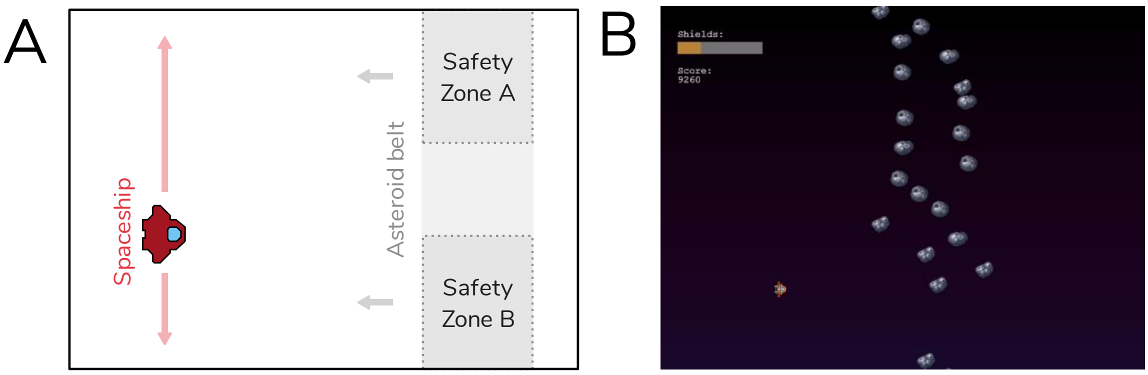

Figure 1. A) Task design. Subjects were tasked with playing a game that had a cover story involving flying a spaceship through asteroid belts. Each asteroid belt featured two locations that could potentially contain escape holes (safety zones), and subjects were instructed to aim to fly their spaceship through these to gain the highest number of points. Subjects were only able to move the spaceship in the Y-dimension, while asteroid belts moved towards the spaceship. The probability of each zone being safe varied over the course of the task but this could be learned, and learning this probability facilitated performance. B) Screenshot of the task, showing the spaceship, an asteroid belt with a hole in the lower safety zone (safety zone B), a representation of the spaceship's integrity (shown by the coloured bar in the top left corner) and the current score.

Avoidance learning task

Traditional lab-based threat learning tasks typically use aversive stimuli such as electric shocks as outcomes to be avoided. As it is not possible to use these stimuli online, we developed a game-based task in which subjects' goal was to avoid negative outcomes. In this game, participants were tasked with flying a spaceship through asteroid belts. Subjects were able to move the spaceship in the Y-axis alone, and this resulted in a one dimensional behavioural output. Crashing into asteroids diminished the spaceship's integrity, and if enough asteroids were hit the game finished. In this eventuality subjects were able to restart and continue where they left off. The overarching goal was to maximise the number of points scored, where the latter accumulated continuously for as long as the game was ongoing, and reset if the spaceship was destroyed. Subjects were shown the current integrity of the spaceship by a bar displayed in the corner of the screen, along with by a display of their current score. No actual monetary reward was given to the subjects for performance on the task.

Crucially, the location of safe spaces in the asteroid belts could be learned, and learning facilitated performance as it allowed correct positioning of the spaceship prior to observing the safe location. The task was designed such that without such pre-emptive positioning it was near impossible to successfully avoid the asteroids, thus encouraging subjects to learn the safest positions. Holes in the asteroids could appear either at the top or bottom of the screen (Figure 1A), and the probability of safety associated with either location varied independently over the course of the task. Thus, it was possible to learn the safety probability associated with each safety zone and adapt one's behaviour accordingly. Participants also completed a control task that required avoidance that was not dependent on learning, enabling us to control for general motor-related avoidance ability in further analyses (described in supplementary material). We elected a priori to exclude subjects with limited response variability (indicated by a standard deviation of their positions below 0.05) so as to remove subjects who did not move the spaceship. However, no subject met this exclusion criterion.

After completing the task, subjects were asked to provide ratings indicating how anxious the task made them feel and how motivated they were to avoid the asteroids, using visual analogue scales ranging from 0 to 100.

Behavioural data extraction

For analysis, we treated each pass through an asteroid belt as a trial. Overall there were 269 trials in total. As a measure of behaviour, we extracted the mean Y position across the 1 second prior to observing the asteroid belt, representing where subjects were positioning themselves in preparation for the upcoming asteroid belt.

Computational modelling of behaviour

Our modelling approach focused on models that allowed the quantification of subjective uncertainty. To this end, we modelled behaviour using approximate Bayesian models that assume subjects estimate safety probability using a beta distribution. This approach is naturally suited to probability estimation tasks, as the beta distribution is bounded between zero and one, and provides a measure of uncertainty through the variance of the distribution. While certain reinforcement learning formulations can achieve similar uncertainty-dependent learning and quantification of uncertainty, we chose beta models as they have an advantage of being computationally simple. Empirically, these models have been used successfully in previous studies to capture value-based learning (27), where they explain behaviour in aversive learning tasks better than commonly used reinforcement learning models (28, 29), a pertinent characteristic in the current task.

The basic premise underlying these models is that evidence for a given outcome is dependent on the number of times this outcome has occurred previously. For example, evidence for safety in a given location should then be highest when safety has been encountered many times in this location. This count can be represented by a parameter *A*, which is then incremented by a given amount every time safety is encountered. Danger is represented by a complementary parameter *B*. The balance between these parameters provides an indication of which outcome is most likely. Meanwhile, the overall number of outcomes counted influences the variance of the distribution and hence the uncertainty about this estimate. Thus, uncertainty is highest when few outcomes have been observed. The exact amount by which *A* and *B* are updated after every observed outcome can be estimated as a free parameter (here termed τ), and we can build asymmetry in learning into the model, so that learning about safety and danger have different rates, allowing updates for A and B to take on different values (here termed τ^+ and τ^-).

Such a model is appropriate in stationary environments, when the probability of a given outcome is assumed to be constant throughout the experiment. However, in our task the probability of safety varied, and so it was necessary to build a forgetting process into the model. This is achieved by incorporating a decay (represented by parameter λ) which diminishes the current values of A and B on every trial. The result of this process is akin to reducing the number of times they have been observed, and maintains the model's ability to update in response to incoming evidence. Estimates for A and B are therefore updated on each trial (*t*) according to the following equation for both shock zones independently (termed X and Y here). Both zones are updated on every trial, as subjects saw the outcome associated with both simultaneously. This formed the basis of all the probabilistic models tested:

$$A_{t+1}^X = (1-\lambda)\cdot A_t^X + outcome_t^X \cdot \tau^+ \cdot W \tag{1}$$

$$B_{t+1}^X = (1-\lambda)\cdot B_t^X + (1-\ outcome_t^X)\cdot \tau^- \ \cdot W \tag{2}$$

We also observed in pilot data that subjects tended to be influenced more by outcomes occurring in the zone they had previously chosen, an effect likely due to attention. On this basis, we incorporated a weighting parameter that allowed the outcome of the unchosen option to be down-weighted by an amount shown in the above equation (W) determined by an additional free parameter, ω.

$$W_{t+1}^X = \begin{matrix} 1\ if\ chosen \\ \omega\ if\ \ unchosen \end{matrix} \tag{3}$$

We can calculate the estimated safety probability for each zone (*P*) by taking the mean of this distribution:

$$P_{t+1}^X = \frac{A_{t+1}^X}{\left(A_{t+1}^X + B_{t+1}^X\right)} \tag{4}$$

Similarly, we can derive a measure of uncertainty on each trial by taking the variance of this distribution.

$$\sigma_{t+1}^{X} = \frac{A_{t+1}^{X} \cdot B_{t+1}^{X}}{\left(A_{t+1}^{X} + B_{t+1}^{X}\right)^{2} \cdot \left(A_{t+1}^{X} + B_{t+1}^{X} + 1\right)} \tag{5}$$

Our best fitting model included a “stickiness” parameter, which caused the chosen safety zone to be more likely to be chosen on the following trial (other models tested and compared with the winning model are described fully in supplementary material). This was achieved by boosting the value of the chosen option (determined by whether the position was above or below the centre) by raising the estimated safety probability of the option to the power of S, a free parameter bounded between zero and one governing the degree of stickiness. Here an exponential effect was chosen over an additive one as it ensures the probability does not take a value above 1.

$$P_{t+1}^{X} = \begin{matrix} {P_{t+1}^{X}}^{S} \; if \;\; chosen \\ P_{t+1}^{X} \; if \; unchosen \end{matrix} \tag{6}$$

In order to fit our model to the observed behaviour, we require an output that represents the position of the spaceship on the screen. This position (pos) was calculated based on the safety probability of the two safety zones, such that the position was biased towards the safest location and was nearer the centre of the screen when it was unclear which position was safest.

$$pos_{t+1} = \frac{\left(P_{t+1}^{X} - P_{t+1}^{Y}\right) + 1}{2} \tag{7}$$

Additionally, based on pilot data we had noted that subjects tended to position themselves nearer to either the top or bottom of the screen, and tended to avoid the centre. To encourage similar performance from our model, a softmax function with inverse temperature (β) as a free parameter was applied to the estimated position:

$$pos_{t+1} = \frac{e^{(\beta \cdot pos_{t+1})}}{e^{(\beta \cdot pos_{t+1})} + e^{(\beta \cdot (1 - pos_{t+1}))}} \tag{8}$$

Although softmax functions are typically used for converting continuous value estimates to choice probabilities in decision making tasks, we emphasise that here this method is used simply to produce position estimates that avoid the centre of the screen rather than the conventional usage of computing choice probability.

For completeness, we also tested two reinforcement learning models, a Rescorla-Wagner model and a variant of this model with different learning rates for better and worse than expected outcomes (15), both of which are described in supplementary material.

Models were fit with a hierarchical Bayesian approach using variational inference implemented in PyMC3, through maximising the likelihood of the data given a reparametrised beta distribution with a mean provided by the model and a single free variance parameter. Model fit was assessed using the Watanabe-Akaike Information Criterion (WAIC)(30), an index of model fit designed for Bayesian models that accounts for model complexity.

Measures of psychiatric traits

Our primary hypotheses focused on state/trait anxiety and intolerance of uncertainty. These were measured using the State Trait Inventory of Cognitive and Somatic Anxiety (STICSA)(31) and the Intolerance of Uncertainty Scale (IUS)(32) respectively. We also wished to examine how behaviour in our task related to the three transdiagnostic factors identified by Gillan et al. (2016), based on factor analysis of a range of psychiatric measures. To measure these factors more efficiently, we developed a reduced set of questions that provided an accurate approximation of the true factor scores, details of which are provided in supplementary material.

Regression models

Bayesian regression models were used to investigate relationships between behaviour and psychiatric measures, predicting each behavioural measure of interest from the psychiatric measures. Our dependent variables were parameters and quantities derived from our model, which represented the way in which an individual learns about safety probability and how they estimate uncertainty. Specifically, we used the two update parameters from our model (τ^+ and τ^-, referring to the extent to which subjects update in response to safety and danger respectively) and the mean safety probability and uncertainty estimates across the task (generated by simulating data from the model with each subject's estimated parameter values). Crucially, the fact that task outcomes were identical for every subject ensured these values were dependent only on the manner by which subjects learned about safety, not the task itself.

These models were constructed using Bambi (33) and fit using Markov chain Monte Carlo (MCMC) sampling, each with 8000 samples, 2000 of which were used for tuning. All models included age and sex as covariates, along with performance on our control task to account for non-learning related avoidance ability. For analyses predicting state and trait anxiety and intolerance of uncertainty, we constructed a separate model for each variable due to the high collinearity between these measures. For analyses including the three transdiagnostic factors, these were entered into a single model. When reporting regression coefficients, we report the mean of the posterior distribution along with the 95% highest posterior density interval (HPDI), representing the points between which 95% of the posterior distribution's density lies.

Partial least squares regression

To provide a data-driven characterisation of the relationship between task behaviour and psychiatric traits, and take advantage of item-level responses in our questionnaire measures, we used partial least squares (PLS) regression to identify dimensions of covariance between individual questions and the measures derived from our modelling. To ensure robustness of these results, we split our data into training and testing sets, made up of 75% and 25% of the data respectively. To identify the appropriate number of components within the training set, we used a 10-fold cross-validation procedure. The mean squared error of the model's predictions was then averaged across test folds to provide an index of the model's predictive accuracy with different numbers of components.

Once the number of components was determined, we validated the model's predictions by testing its predictive accuracy on the held-out 25% of the data. To provide a measure of statistical significance we used permutation testing, repeating this process 1000 times with shuffled outcome variables to assess its predictive accuracy when fitted on data where no relationship exists between the predictors and outcomes.

Preregistration

The main hypotheses and methods of this study were preregistered on the Open Science Framework (https://osf.io/jp5gn). The data-driven PLS regression analysis was exploratory.

Results

Task performance

Subjects engaged and performed well at the task, on average hitting enough 'asteroids' to destroy the spaceship 1.92 times (SD= 1.95) over the course of the task. To provide a simple behavioural measure of subjects' avoidance tendencies, we calculated the degree to which they moved their position following safe and dangerous outcomes at their chosen location. As expected, this revealed subjects tended to shift their position more following danger than following safety outcomes, indicating they adapted their behaviour in the task, and learnt about the safest locations (Figure 2A).

Subjects reported high motivation to perform the task, providing a mean rating of 85.70 (SD = 18.44) when asked to rate how motivated they were to avoid asteroids on a scale from 0-100 (Figure 3B). When asked to rate how anxious the task made them feel on a similar scale, subjects provided an average rating of 40.89 (SD = 30.45), although visual inspection of their responses indicated a strongly bimodal distribution (Figure 3A).

Computational modelling of behaviour

To quantitatively describe behaviour, we fit a series of computational models to the behavioural data. Probabilistic models tended to outperform reinforcement learning models according to WAIC scores that penalise based on model complexity (Figure 2B), with the winning model being an asymmetric leaky beta model incorporating choice stickiness. Simulating responses using the model, using each subject's estimated parameter values, produced profiles that demonstrated a high concordance with the true data, reproducing broad behavioural patterns seen in the true data (Figure S1).

Due to an interest in asymmetric learning about safety and punishment, the two parameters of the preferred model with greatest relevance to our hypotheses were update rates for safety and danger. To verify that the estimated parameters truly reflected behavioural tendencies, we examined correlations between their estimated values and the tendency to stay following safety and move following danger across subjects (Figure 2D). These were robustly correlated, providing confidence that these model-derived parameters related to relevant behavioural measures. Interestingly, we observed a negative bias in the values of these parameters whereby subjects tended to update to a greater extent in response to danger than safety ($t(400) = 26.76$, $p < .001$, Figure 2E).

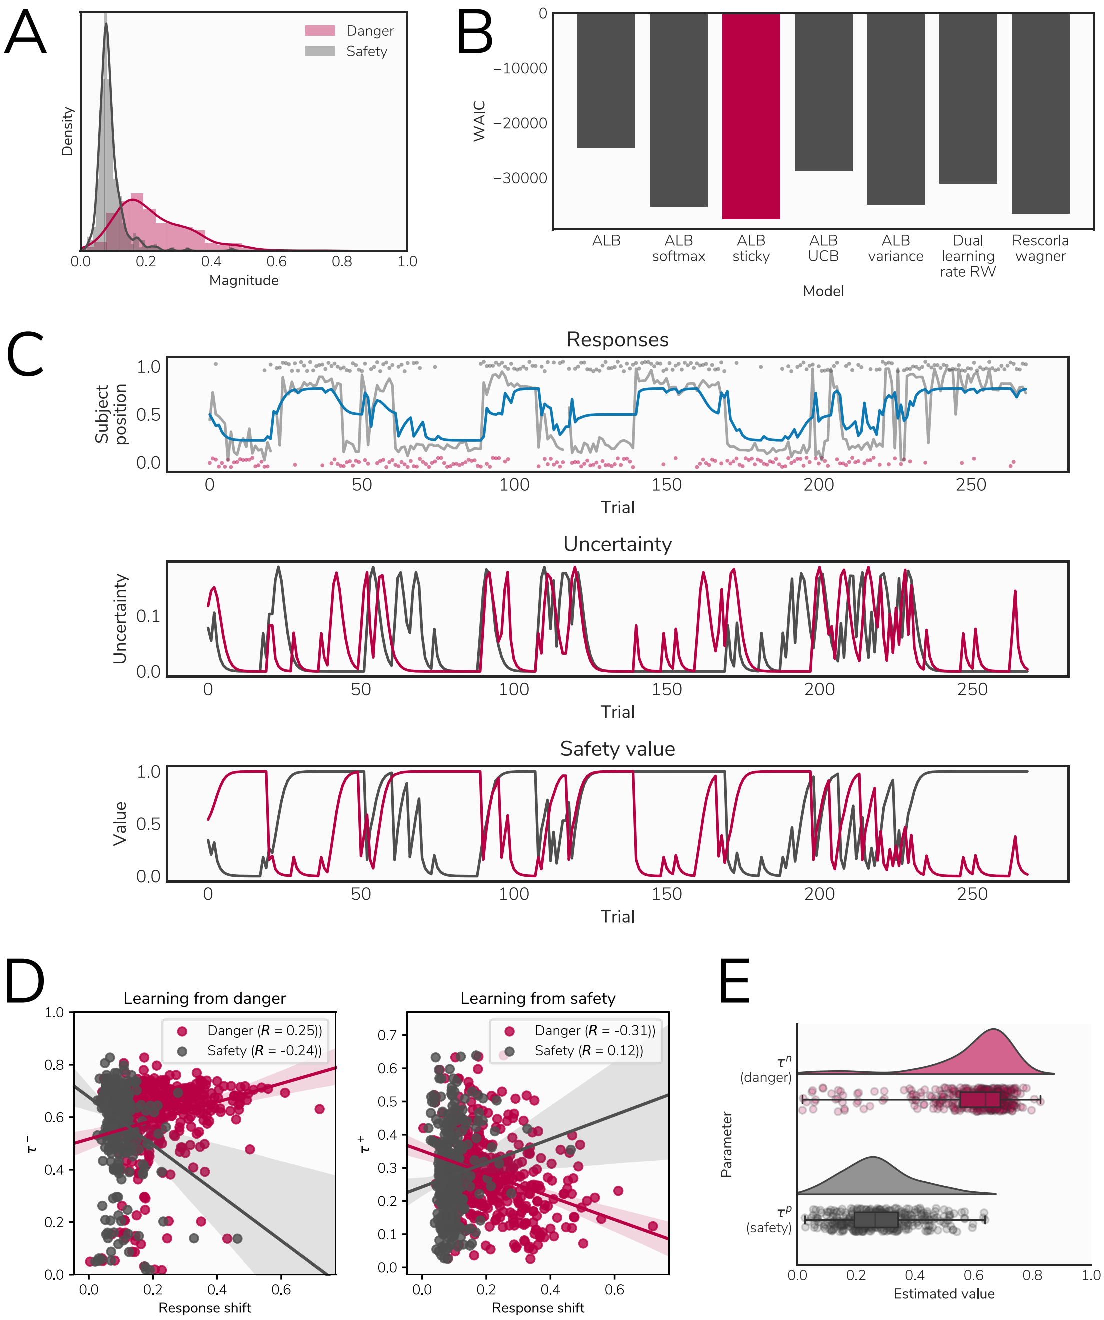

Figure 2. A) Degree of location switching after encountering danger and safety across subjects. The switch magnitude is the average absolute change in position between trial *n* and trial *n*+1. As expected, subjects showed more switching behaviour after encountering danger and were more likely to stay in the same position following a safe outcome. B) Model comparison results, showing the WAIC score for each model with the winning model highlighted. ALB = asymmetric leaky beta, RW = Rescorla-Wagner. C) The top panel shows responses and model fit for an example subject. In the top panel, the grey line represents the subject's position throughout the task, with the grey and red dots representing safe locations on each trial. The blue line represents simulated data from the model for this subject. The lower two panels show estimated uncertainty and safety probability for each stimulus (represented by the grey and red lines) across the duration of the task, generated by simulating data from the model. D) Correlations between estimated update parameters for danger (left) and safety (right) and our behavioural measure of position switching after these outcomes across subjects, demonstrating that parameters from our model reflect purely behavioural characteristics. E) Distributions of estimated parameter values for τ^+ and τ^-, representing

update rates following danger and safety outcomes respectively, showing a bias in updating whereby subjects update to a greater extent in response to danger than safety.

Relationships with psychiatric measures

First, as a validation of our task, we tested whether task-induced anxiety was predicted by state and trait anxiety. Pearson correlations revealed significant relationships with both these measures (state R=0.25, p<.001; trait R=0.21, p<.001), indicating that individuals reporting more general anxiety also found the task more anxiety-inducing.

Target variable	Predictor	Estimate (+/- 95% HPDI)
Threat update (τ^-)	IUS	0.07 (-0.03, 0.16)
	STICSA S	-0.09 (-0.19, 0.01)
	STICSA T	-0.03 (-0.12, 0.07)
Safety update (τ^+)	IUS	0.0 (-0.1, 0.09)
	STICSA S	**0.12 (0.02, 0.21)**
	STICSA T	0.09 (-0.01, 0.18)
Mean safety uncertainty	IUS	-0.05 (-0.14, 0.05)
	STICSA S	-0.09 (-0.19, 0.01)
	STICSA T	-0.08 (-0.18, 0.01)
Mean safety probability	IUS	-0.02 (-0.12, 0.07)
	STICSA S	**0.12 (0.02, 0.21)**
	STICSA T	0.06 (-0.03, 0.15)

Table 1. Estimates from regression model predicting learning-related variables derived from our computational model from measures of intolerance of uncertainty, state anxiety, and trait anxiety. Effects with HPDIs excluding zero are highlighted. IUS: Intolerance of uncertainty scale; STICSA S: State-trait inventory of cognitive and somatic anxiety, state measure; STICSA T: State-trait inventory of cognitive and somatic anxiety, trait measure; HPDI: Highest posterior density estimate

We next asked whether our four behavioural variables of interest (threat and danger update parameters, mean estimated safety probability, and mean estimated uncertainty) were associated with anxiety (both state and trait) and intolerance of uncertainty respectively. Contrary to our hypotheses, the only relationships with HPD intervals that did not include zero were positive effects of state anxiety on safety update rates and mean estimated safety probability (Figure 3E, Table 1), indicating more anxious individuals learned faster about safety and perceived safety as more likely overall.

We then tested whether behaviour in the task was associated with three transdiagnostic factors identified in previous research (6). Here, we observed unexpected relationships with a factor labelled compulsivity and intrusive thought (Figure 3F, Table 2), reflecting the fact that subjects scoring higher on this factor learned faster about safety and had higher safety probability estimates. No other effects had HPDIs that excluded zero.

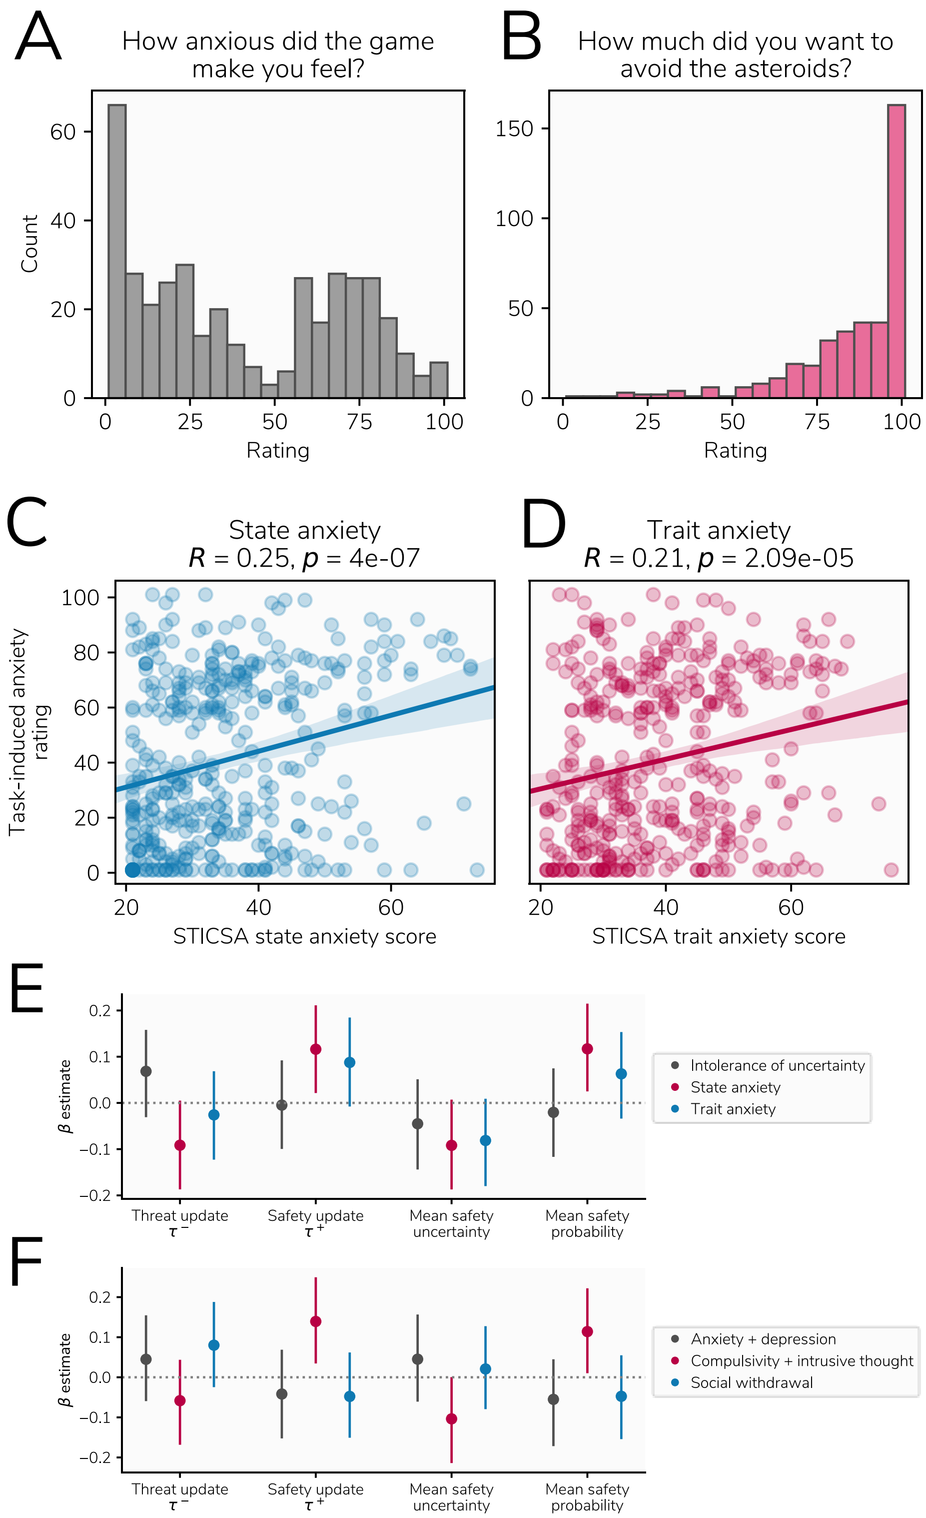

Figure 3. A) Distribution of task-induced anxiety ratings recorded after the task. B) Distribution of task motivation ratings. C and D) Relationships between task-induced anxiety ratings and state and trait anxiety scores. E) Results of state/trait anxiety and intolerance of uncertainty models, showing relationships between these psychiatric variables and behavioural variables. Points indicate the mean of the posterior distribution for the regression coefficient parameter, while error bars represent the 95% highest posterior density interval. The β estimate here refers to the regression coefficient for each predictor. F) Results of three factor model, showing relationships between behaviour and factors labelled anxiety and depression, compulsivity and intrusive thought, and social withdrawal.

Target variable	Predictor	Estimate (+/- 95% HPDI)
Threat update (τ^-)	AD	0.04 (-0.06, 0.15)
	CBIT	-0.06 (-0.17, 0.04)
	SW	0.08 (-0.02, 0.19)
Safety update (τ^+)	AD	-0.04 (-0.15, 0.07)
	CBIT	**0.14 (0.03, 0.25)**
	SW	-0.05 (-0.15, 0.06)
Mean safety uncertainty	AD	0.05 (-0.06, 0.16)
	CBIT	-0.1 (-0.21, 0.0)
	SW	0.02 (-0.08, 0.13)
Mean safety probability	AD	-0.06 (-0.17, 0.04)
	CBIT	**0.11 (0.01, 0.22)**
	SW	-0.05 (-0.15, 0.05)

Table 2. Estimates from regression model predicting learning-related variables derived from our computational model from the three transdiagnostic factors identified by Gillan et al. (2016). Effects with HPDIs excluding zero are highlighted. AD: Anxious-depression; CBIT: Compulsive behaviour and intrusive thought; SW: Social withdrawal; HPDI: Highest posterior density estimate

Psychiatric constructs derived from behaviour and self-report

Finally, we sought to identify factors represented shared covariance between our behavioural measures and individual questionnaire items using PLS regression. We first identified the number of components that best described our data by evaluating the performance of a predictive PLS model using cross-validation. Two components gave the best predictive performance (Figure 4A). We then evaluated the performance of this model on held out data using permutation testing, and this indicated our model achieved a statistically significant level of predictive accuracy (permutation $p = 0.025$, Figure 4B).

To aid interpretation of these two components we examined how our behavioural variables loaded on each. The first component had positive weights on update rates in response to safety and estimated safety likelihood, and negative weights on update rates in response to threat, decay, stickiness, and mean uncertainty estimates (Figure 4C), while the reverse was true of the second component. Loadings on questionnaire items were varied, but the first component tended to load most strongly on items describing physical symptoms of anxiety, compulsive behaviour, and impulsivity. In contrast, the second factor loaded primary on items describing social anxiety and depressed mood. For illustrative purposes, a selection of items showing strong differences in loadings between components are shown in Figure 4D, with full details available in supplementary material.

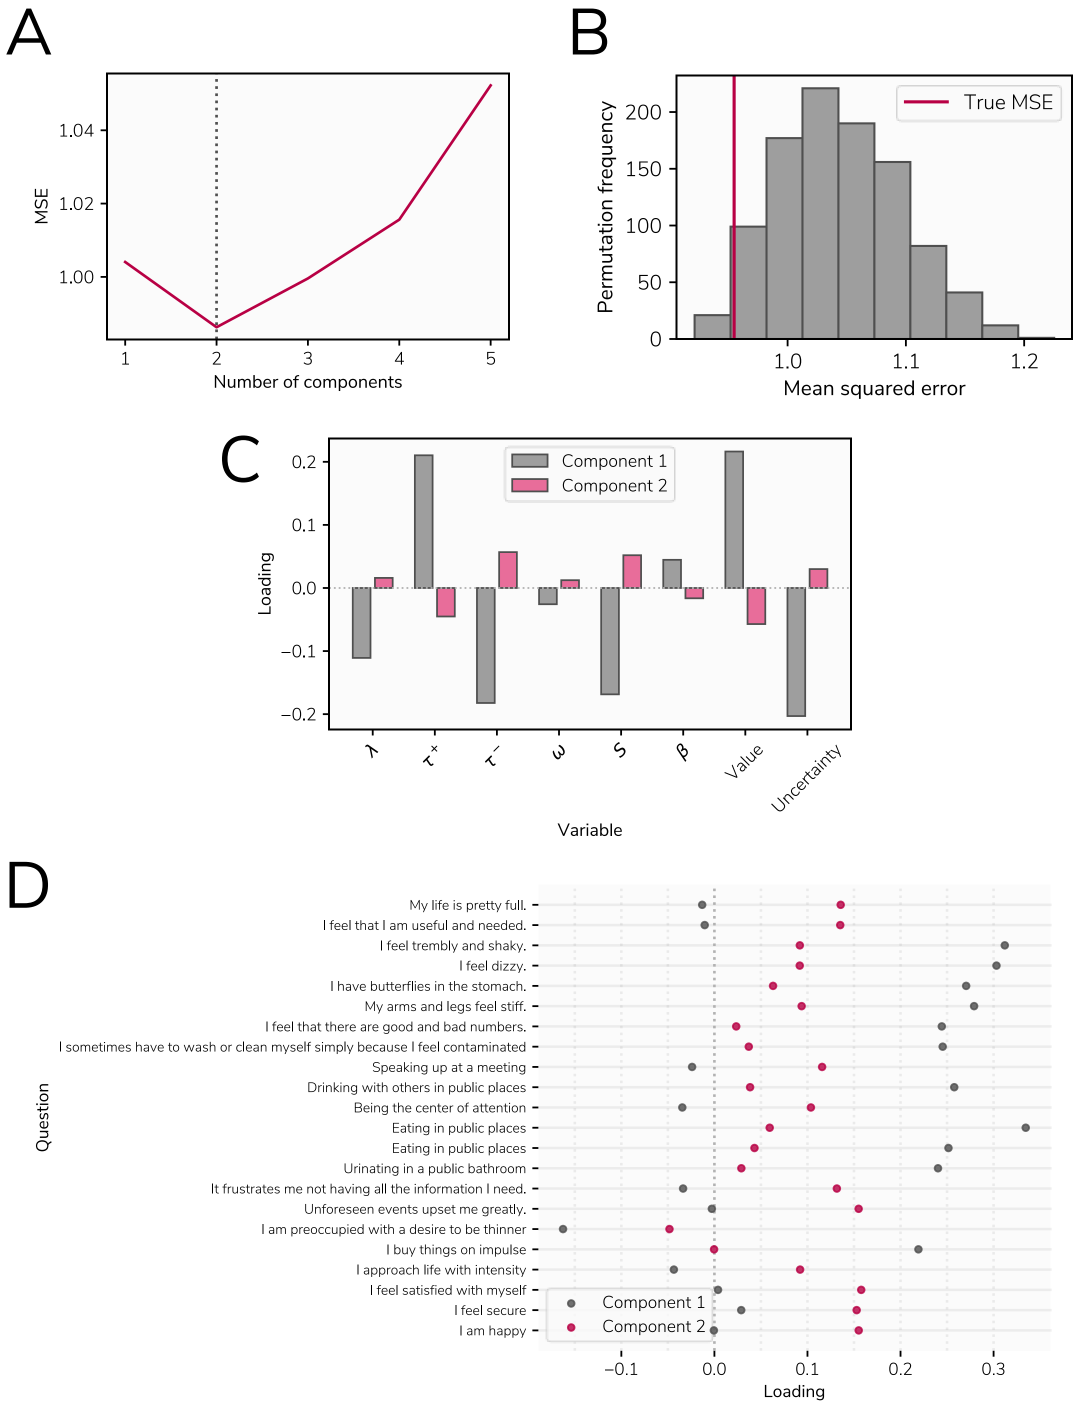

Figure 4. Results of PLS regression analysis. A) The optimal number of components was determined based on cross-validated predictive accuracy within 75% of the data used for training the model. This figure represents the negative mean squared error of these predictions across models with between one and five factors, showing best performance with 2 components. B) Null distribution of predictive accuracy scores generated by retraining our PLS regression model on 1000 permuted datasets and testing on the held out 25% of the data set, with the MSE achieved by the model trained on the true data shown by the red line. C) Loadings for the two components on behavioural variables, including all parameters in the model and mean safety probably and uncertainty estimates. D) Loadings on a selection of questionnaire items showing the largest dissociations in loadings between the two components.

Discussion

Perceptions of danger and safety have been linked to key symptoms of psychiatric disorders. Here, we use computational modelling to show that when subjects learn to avoiding threat, transdiagnostic components of psychopathology relate to how they learn about both safety likelihood, and uncertainty.

We found an unexpected relationship between biases in learning and anxiety. Contrary to our a priori hypotheses, subjects scoring higher on state anxiety tended to update their predictions to a greater extent in response to safety, as well as perceiving safety to be more likely overall, than those scoring low on this measure. These results diverge from previous findings that reported individuals diagnosed with clinical anxiety and depression learn faster from punishment (12). One potential explanation for this discrepancy is that the aforementioned work included subjects with clinical diagnoses, involving a mix of anxiety and depressive disorders. When we examined transdiagnostic factors, we found that symptoms of depression were associated with elevated learning from threat, suggesting that such a bias in learning is associated with depressive symptoms. It is also possible that the nature of our game-based online task engaged processes distinct from that of standard lab-based tasks, though it is not clear what these might be. Nevertheless, we observed a similar pattern of results, where more anxious subjects learn faster about safety, in a study using a more traditional lab-based aversive learning task (Wise et al, 2019).

We found a similar pattern of enhanced learning from threat when examining a transdiagnostic factor representing compulsivity and intrusive thought. Although this factor has been shown to be associated with more model-free behaviour (6, 24), altered confidence judgements (25), and action-confidence coupling (8) in large-scale samples, to date it has not been investigated with regard to threat learning. Notably, we found a weak relationship between this factor and uncertainty, whereby more compulsive individuals had higher certainty in their safety estimates, echoing previous work in perceptual decision making that showed this factor is associated with higher confidence estimates (8, 25). We only found weak relationships with the other two factors, representing anxious-depression and social withdrawal, in a direction indicative of lower safety probability estimates and higher uncertainty.

Perhaps our most striking results come from our data-driven approach where we derived components of psychopathology grounded in computational analyses of aversive learning behaviour. Using PLS regression we identified two components, one associating greater learning from safety with physiological symptoms of anxiety and compulsivity, while the other associated greater learning from threat with depressive symptoms and social anxiety. These results suggest that symptoms which are typically grouped together, either within diagnostic categories or in factors identified using self-report measures alone, are associated with distinct patterns of learning behaviour. This data-driven analysis also revealed relationships between aversive learning and impulsive behaviour, encompassing a symptom dimension that is typically studied in the context of reward processing (34). Individuals scoring higher on these symptoms exhibited higher safety learning, which may explain previously observed relationships between impulsivity and risk tolerance (35).

A feature of this study is our development of a novel online task for measuring aversive learning. A number of studies examining other aspects of learning and decision making in the context of psychiatric disorders have also availed of large samples recruited through online services (6, 8, 24, 25). However, it has been difficult to examine aversive learning in online environments, as aversive lab stimuli such as shock cannot be easily administered online. A game-based design allowed us to design a task that required avoidance behaviour as well as evoke feelings of anxiety. Although qualitatively different from standard lab-based tasks, we observed similar patterns of biased learning to that seen in lab-based work (36). An added benefit of our task is that it is highly engaging, and subjects reported feeling motivated to perform well.

One potential limitation of this study is a focus on a general population sample. While this might limit applicability to clinical anxiety, other research indicates that findings from clinical samples replicate in samples recruited online (6, 8). Furthermore, it seems reasonable to assume that clinical anxiety is on a spectrum with less severe expressions of this condition. Although we did not deliberately set out to recruit individuals with clinically significant anxiety, 36% of our sample scored at or above a threshold designed for the detection of anxiety disorders on our measure of trait anxiety (see supplementary material). In light of this, and given limitations with research in clinical samples that includes medication load (37) and recruitment challenges (38), online samples provide an effective method for studying clinically-relevant phenomena. Additionally, it is important to note that the effects we observed were small, as in previous studies using large-scale online testing (6, 24, 24). However, large samples provide accurate effect size estimates in contrast to the exaggerated effects that are common in studies using small samples (39). Such small effects are unsurprising given the multifactorial nature of psychiatric disorders (40). While we have shown aversive learning to be important, we acknowledge this is likely to be one of a multitude of processes involved in the development of these conditions.

In conclusion, our results demonstrate novel links between transdiagnostic symptoms of psychiatric disorders and mechanisms of threat learning and uncertainty estimation in aversive environments. The findings emphasise the importance of these processes not only in anxiety but indicate a likely relevance across a spectrum of psychopathology.

Acknowledgements

T.W. is supported by a Wellcome Trust Sir Henry Wellcome Fellowship (206460/17/Z). R.J.D. holds a Wellcome Trust Investigator award (098362/Z/12/Z). The Max Planck UCL Centre is a joint initiative supported by UCL and the Max Planck Society. The Wellcome Centre for Human Neuroimaging is supported by core funding from the Wellcome Trust (203147/Z/16/Z). We thank Evan Russek for comments on an earlier version of this manuscript.

References

1. Reininghaus U, Kempton MJ, Valmaggia L, Craig TKJ, Garety P, Onyejiaka A, *et al.* (2016): Stress Sensitivity, Aberrant Salience, and Threat Anticipation in Early Psychosis: An Experience Sampling Study. *Schizophr Bull*. 42: 712–722.

2. Harrison A, Tchanturia K, Treasure J (2010): Attentional Bias, Emotion Recognition, and Emotion Regulation in Anorexia: State or Trait? *Biological Psychiatry*, Positron Emission Tomography Imaging of Dopamine and Opiate Receptors in Addiction. 68: 755–761.

3. Friston KJ, Stephan KE, Montague R, Dolan RJ (2014): Computational psychiatry: the brain as a phantastic organ. *The Lancet Psychiatry*. 1: 148–158.

4. Montague PR, Dolan RJ, Friston KJ, Dayan P (2012): Computational psychiatry. *Trends in Cognitive Sciences*, Special Issue: Cognition in Neuropsychiatric Disorders. 16: 72–80.

5. Browning M, Behrens TE, Jocham G, O'Reilly JX, Bishop SJ (2015): Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nat Neurosci*. 18: 590–596.

6. 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.

7. Huang H, Thompson W, Paulus MP (2017): Computational Dysfunctions in Anxiety: Failure to Differentiate Signal From Noise. *Biological Psychiatry*, Computational Psychiatry. 82: 440–446.

8. Seow TX, Gillan C (2019): Transdiagnostic phenotyping reveals a range of metacognitive deficits associated with compulsivity. *bioRxiv*. 664003.

9. Lawson RP, Mathys C, Rees G (2017): Adults with autism overestimate the volatility of the sensory environment. *Nat Neurosci*. advance online publication. doi: 10.1038/nn.4615.

10. Mkrtchian A, Aylward J, Dayan P, Roiser JP, Robinson OJ (2017): Modeling Avoidance in Mood and Anxiety Disorders Using Reinforcement Learning. *Biological Psychiatry*, Stress: Mechanisms in Gut and Brain. 82: 532–539.

11. Vinckier F, Gaillard R, Palminteri S, Rigoux L, Salvador A, Fornito A, *et al.* (2016): Confidence and psychosis: a neuro-computational account of contingency learning disruption by NMDA blockade. *Mol Psychiatry*. 21: 946–955.

12. Aylward J, Valton V, Ahn W-Y, Bond RL, Dayan P, Roiser JP, Robinson OJ (2019): Altered learning under uncertainty in unmedicated mood and anxiety disorders. *Nature Human Behaviour*. 1.

13. Robinson OJ, Chase HW (2017): Learning and Choice in Mood Disorders: Searching for the Computational Parameters of Anhedonia. *Computational Psychiatry*. 1–26.

14. Sharp PB, Eldar E (2019): Computational Models of Anxiety: Nascent Efforts and Future Directions. *Curr Dir Psychol Sci*. 28: 170–176.

15. Lefebvre G, Lebreton M, Meyniel F, Bourgeois-Gironde S, Palminteri S (2017): Behavioural and neural characterization of optimistic reinforcement learning. *Nature Human Behaviour*. 1: 0067.

16. Bach DR, Dolan RJ (2012): Knowing how much you don't know: a neural organization of uncertainty estimates. *Nat Rev Neurosci*. 13: 572–586.

17. de Berker AO, Rutledge RB, Mathys C, Marshall L, Cross GF, Dolan RJ, Bestmann S (2016): Computations of uncertainty mediate acute stress responses in humans. *Nat Commun*. 7: 10996.

18. Mathys C, Daunizeau J, Friston KJ, Stephan KE (2011): A Bayesian foundation for individual learning under uncertainty. *Front Hum Neurosci*. 5: 39.

19. Mathys C, Lomakina EI, Daunizeau J, Iglesias S, Brodersen KH, Friston KJ, Stephan KE (2014): Uncertainty in perception and the Hierarchical Gaussian Filter. *Front Hum Neurosci*. 8: 825.

20. Grupe DW, Nitschke JB (2013): Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nat Rev Neurosci*. 14: 488–501.

21. Dugas MJ, Gagnon F, Ladouceur R, Freeston MH (1998): Generalized anxiety disorder: a preliminary test of a conceptual model. *Behaviour Research and Therapy*. 36: 215–226.

22. Barlow DH (2000): Unraveling the mysteries of anxiety and its disorders from the perspective of emotion theory. *Am Psychol*. 55: 1247–1263.

23. Grillon C (2002): Associative learning deficits increase symptoms of anxiety in humans. *Biological Psychiatry*. 51: 851–858.

24. Patzelt EH, Kool W, Millner AJ, Gershman SJ (2018): Incentives Boost Model-Based Control Across a Range of Severity on Several Psychiatric Constructs. *Biological Psychiatry*. doi: 10.1016/j.biopsych.2018.06.018.

25. Rouault M, Seow T, Gillan CM, Fleming SM (2018): Psychiatric Symptom Dimensions Are Associated With Dissociable Shifts in Metacognition but Not Task Performance. *Biological Psychiatry*, Translating Biology to Treatment in Schizophrenia. 84: 443–451.

26. Palan S, Schitter C (2018): Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*. 17: 22–27.

27. de Boer L, Axelsson J, Riklund K, Nyberg L, Dayan P, Bäckman L, Guitart-Masip M (2017): Attenuation of dopamine-modulated prefrontal value signals underlies probabilistic reward learning deficits in old age. (W. Schultz, editor) *eLife*. 6: e26424.

28. Tzovara A, Korn CW, Bach DR (2018): Human Pavlovian fear conditioning conforms to probabilistic learning. *PLOS Computational Biology*. 14: e1006243.

29. Wise T, Michely J, Dayan P, Dolan RJ (2019): A computational account of threat-related attentional bias. *PLOS Computational Biology*. 15: e1007341.

30. Watanabe S (2010): Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory. *J Mach Learn Res*. 11: 3571–3594.

31. Grös DF, Antony MM, Simms LJ, McCabe RE (2007): Psychometric properties of the State-Trait Inventory for Cognitive and Somatic Anxiety (STICSA): comparison to the State-Trait Anxiety Inventory (STAI). *Psychol Assess*. 19: 369–381.

32. Carleton RN, Norton MAPJ, Asmundson GJG (2007): Fearing the unknown: A short version of the Intolerance of Uncertainty Scale. *Journal of Anxiety Disorders*. 21: 105–117.

33. Yarkoni T, Westfall J (2016): Bambi: A simple interface for fitting Bayesian mixed effects models. doi: 10.31219/osf.io/rv7sn.

34. Plichta MM, Scheres A (2014): Ventral–striatal responsiveness during reward anticipation in ADHD and its relation to trait impulsivity in the healthy population: A meta-analytic review of the fMRI literature. *Neuroscience & Biobehavioral Reviews*. 38: 125–134.

35. Deakin J, Aitken M, Robbins T, Sahakian BJ (2004): Risk taking during decision-making in normal volunteers changes with age. *Journal of the International Neuropsychological Society*. 10: 590–598.

36. Wise T, Michely J, Dayan P, Dolan RJ (2019): A computational account of threat-related attentional bias.

37. Phillips ML, Travis MJ, Fagiolini A, Kupfer DJ (2008): Medication Effects in Neuroimaging Studies of Bipolar Disorder. *American Journal of Psychiatry*. 165: 313–320.

38. Wise T, Arnone D, Marwood L, Zahn R, Lythe KE, Young AH (2016): Recruiting for research studies using online public advertisements: examples from research in affective disorders. *Neuropsychiatr Dis Treat*. 12: 279–285.

39. Gelman A, Carlin J (2014): Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors. *Perspect Psychol Sci*. 9: 641–651.

40. Kendler KS (2019): From Many to One to Many—the Search for Causes of Psychiatric Illness. *JAMA Psychiatry*. 76: 1085–1091.