nature human behaviour

Article

https://doi.org/10.1038/s41562-022-01455-y

Trait somatic anxiety is associated with reduced directed exploration and underestimation of uncertainty

Received: 22 September 2021

Accepted: 26 August 2022

Published online: 3 October 2022

Check for updates

Haxue Fan [1], Samuel J. Gershman [1,2,3] and Elizabeth A. Phelps [1,2]

Anxiety has been related to decreased physical exploration, but past findings on the interaction between anxiety and exploration during decision making were inconclusive. Here we examined how latent factors of trait anxiety relate to different exploration strategies when facing volatility-induced uncertainty. Across two studies (total $N = 985$), we demonstrated that people used a hybrid of directed, random and undirected exploration strategies, which were respectively sensitive to relative uncertainty, total uncertainty and value difference. Trait somatic anxiety, that is, the propensity to experience physical symptoms of anxiety, was inversely correlated with directed exploration and undirected exploration, manifesting as a lesser likelihood for choosing the uncertain option and reducing choice stochasticity regardless of uncertainty. Somatic anxiety is also associated with underestimation of relative uncertainty. Together, these results reveal the selective role of trait somatic anxiety in modulating both uncertainty-driven and value-driven exploration strategies.

From daily errands to life decisions, people constantly face the explore-exploit dilemma: should I stick with the current best option (exploit), or should I try something else that could potentially be better (explore)? Because exploration entails inherent risk, it may evoke a heightened level of anxiety and may be avoided altogether by individuals with chronically high levels of anxiety. However, the multifaceted nature of both exploration and anxiety complicates the study of their relationship. In this paper, we use computational models of exploration in tandem with dimensionality reduction of anxiety measures to cleave the exploration-anxiety relationship into multiple distinct components.

An anticorrelation between anxiety and exploration has long been assumed in animal models[1–3]. In the open-field test, a paradigm used to measure exploratory behaviour, researchers put the animal in a square box and compare the time it spends in the centre and the outer edge of the box. A decrease in time spent in the exposed centre area is regarded as a core feature of an anxiety phenotype in non-human animals. Likewise, anxiolytic drugs restore exploratory behaviour[4]. With paradigms mimicking the open-field test and its variants, a few human studies have also shown that anxiety inversely correlates with exploration[5,6].

Beyond the exploration of physical locations, exploration is a key component of human decision-making broadly. Specifically, it is a crucial element of tasks examining decision making under uncertainty[7] in which, to maximize the long-term reward, people need to collect information about a range of options, sometimes at the expense of foregoing the current most rewarding one. In other words, exploration is a reward-information trade-off: the action of directly experiencing an alternative to gather information about its value sometimes comes at the opportunity cost of not benefitting from the current best option[8]. Although studies have shown that anxious individuals are more averse to uncertainty[9], how anxiety influences exploration during decision making remains unclear. On one hand, people with anxiety symptoms tend to avoid uncertain options[10], which reduces exploration. On the other hand, anxiety is associated with the elevated valuation of information[11,12], which encourages exploration behaviour to reduce uncertainty. It is also possible that

[1]Department of Psychology, Harvard University, Cambridge, MA, USA. [2]Center for Brain Science, Harvard University, Cambridge, MA, USA. [3]Center for Brains, Minds and Machines, Cambridge, MA, USA. e-mail: phelps@fas.harvard.edu

https://doi.org/10.1038/s41562-022-01455-y

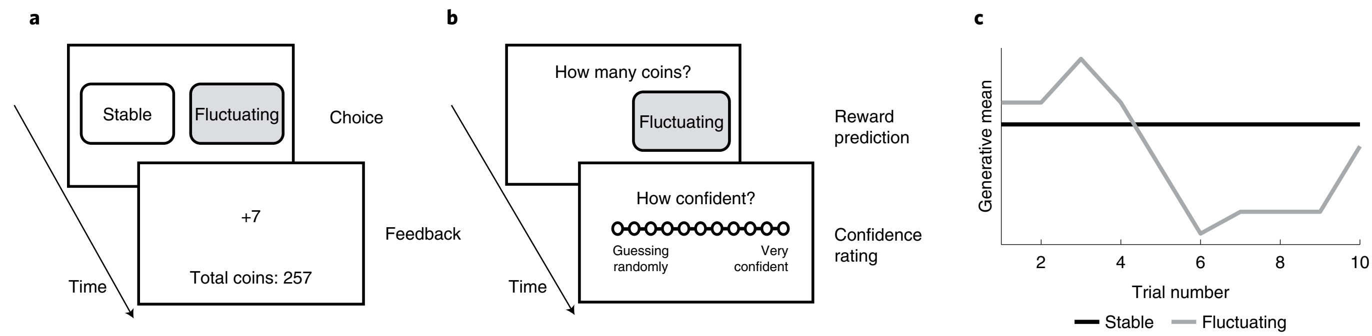

Fig. 1 | Study designs. a, On each trial, participants chose between two slot machines and received feedback on their choices. The participants could see the labels of the slot machines before making a choice and were explicitly instructed about the differences between a stable and a fluctuating machine before the task. **b**, The prediction task was completed at the end of each block in Experiment 2, where participants reported their reward predictions for the options they encountered in this block and rated their confidence. **c**, An example reward structure for a stable and a fluctuating arm. In the experiment, the means of both options were resampled from the zero-mean Gaussian at the beginning of each block.

the apparently increased exploration reflects an overall tendency to behave more randomly[13].

The inconsistency in these results is due to at least two reasons. First, people engage in distinct forms of exploration given different task structures (for example, the kind of uncertainty manipulated), making it hard to compare findings across heterogeneous experiments. Second, most existing studies view anxiety as a single-dimension construct and thus fail to investigate whether different anxiety dimensions (for example, cognitive and somatic anxiety) are linked to distinct exploration patterns. Therefore, the goal of the current study is to experimentally dissociate different exploration strategies to investigate their mapping onto distinct anxiety dimensions.

Recent work has revealed that people could use both the value and the uncertainty in the environment to guide exploration. A standard value-driven strategy is softmax exploration[14], which chooses an option with probability proportional to the exponential of its expected value. This strategy is closely related to other decision strategies such as probability matching and has received some empirical support (for a detailed review, see ref. [7]). The uncertainty-guided strategies fall into two distinct categories: directed and random exploration[15–18]. Directed exploration explicitly favours the uncertain option by adding an uncertainty bonus to each option's estimated value. Therefore, choices are directed towards particular options to obtain more information about their value. For instance, when choosing between a new cafe and your go-to place for lunch, you may prefer the novel option to assess whether it is better than your usual choice. This strategy is sensitive to the relative uncertainty between options[15,19] and can be captured by reinforcement learning algorithms such as Upper Confidence Bound[20,21]. In contrast, random exploration scales choice stochasticity in proportion to the total uncertainty aggregated across options. This kind of exploration can be captured by an algorithm called Thompson sampling[22]. For instance, imagine two cafes that offer different pastries every week (that is, volatile cafes). Since the pastry menus are constantly updated, the total uncertainty is higher than in a scenario where both cafes offer a fixed set of bakery items (that is, stable cafes). According to random exploration, people should explore more when facing volatile cafes.

On the basis of the link between uncertainty and exploration, past work has dissociated directed and random exploration by independently manipulating relative and total uncertainty[15,23–25]. This was achieved by changing riskiness, that is, the moment-to-moment outcome variability of an option (for example, the flavour difference of the same pastry between batches). This leaves open the question of whether exploration is truly driven by epistemic uncertainty (as proposed in computational models) or by aleatory uncertainty (the inherent variability of outcomes). To provide data on this question, here we experimentally manipulate epistemic uncertainty while holding aleatory uncertainty fixed. We do this by varying volatility[26], that is, the change rate of an option's true value (for example, the update speed of the pastry menus). Although the uncertainty-driven exploration framework predicts that the impact of uncertainty on exploration is independent of the way it is varied, this hypothesis is yet to be tested. Therefore, the current study examines whether the effects of volatility-induced uncertainty are comparable to those of riskiness-induced uncertainty. Focusing on volatility-related uncertainty also helps link the current study with the broader literature on the relationship between anxiety and decision making, where anxiety is associated with impaired performance as well as biased estimation of uncertainty in environments with changing volatility[27–33]. To pursue this connection more directly, we report a second study in which we measure uncertainty estimates directly and link them to anxiety dimensions.

Similar to exploration, anxiety is a multifactorial construct. There is considerable variation among individuals in the symptom profiles when experiencing anxiety[34–38], the majority of which can be sorted into cognitive and somatic categories[39,40]. The cognitive dimension includes symptoms associated with thought processes, including rumination, worry and intrusive thoughts. In contrast, the somatic dimension includes physical manifestations of anxiety such as sweating, trembling and palpitation. Recent work showed that cognitive and somatic anxiety have differential effects on aversive learning[41]. Therefore, it is possible that analogous individual differences could influence exploratory behaviour. Given the limited past research on the association between different anxiety components and learning, we do not have strong a priori hypotheses on the direction of the anxiety-exploration relationship.

Results

People use a hybrid of directed and random exploration strategies

In Experiment 1, participants ($N = 501$) performed a two-armed bandit task in which they were informed about the volatility of both options. During the task, participants repeatedly chose between two options ('arms') and received feedback (points delivered by the selected arm) (Fig. 1a). They were instructed to collect as many points as possible. Before making a decision, participants saw the labels of the arms, each of which could be either 'Stable' (S) or 'Fluctuating' (F). In other words, they knew whether an arm is stable or fluctuating before they interacted with it. Both types of arms delivered rewards drawn from a Gaussian distribution around its current generative mean. The generative mean of the fluctuating arm diffused in a Gaussian random walk within a block while the mean of the stable arm stayed the same (see Fig. 1c for an example). The trial type was denoted by the pair of option labels (for example, on SF trials, option 1 on the left is stable, and option 2 on the right is fluctuating), resulting in four trial types (SF, FS, SS and FF).

 https://doi.org/10.1038/s41562-022-01455-y

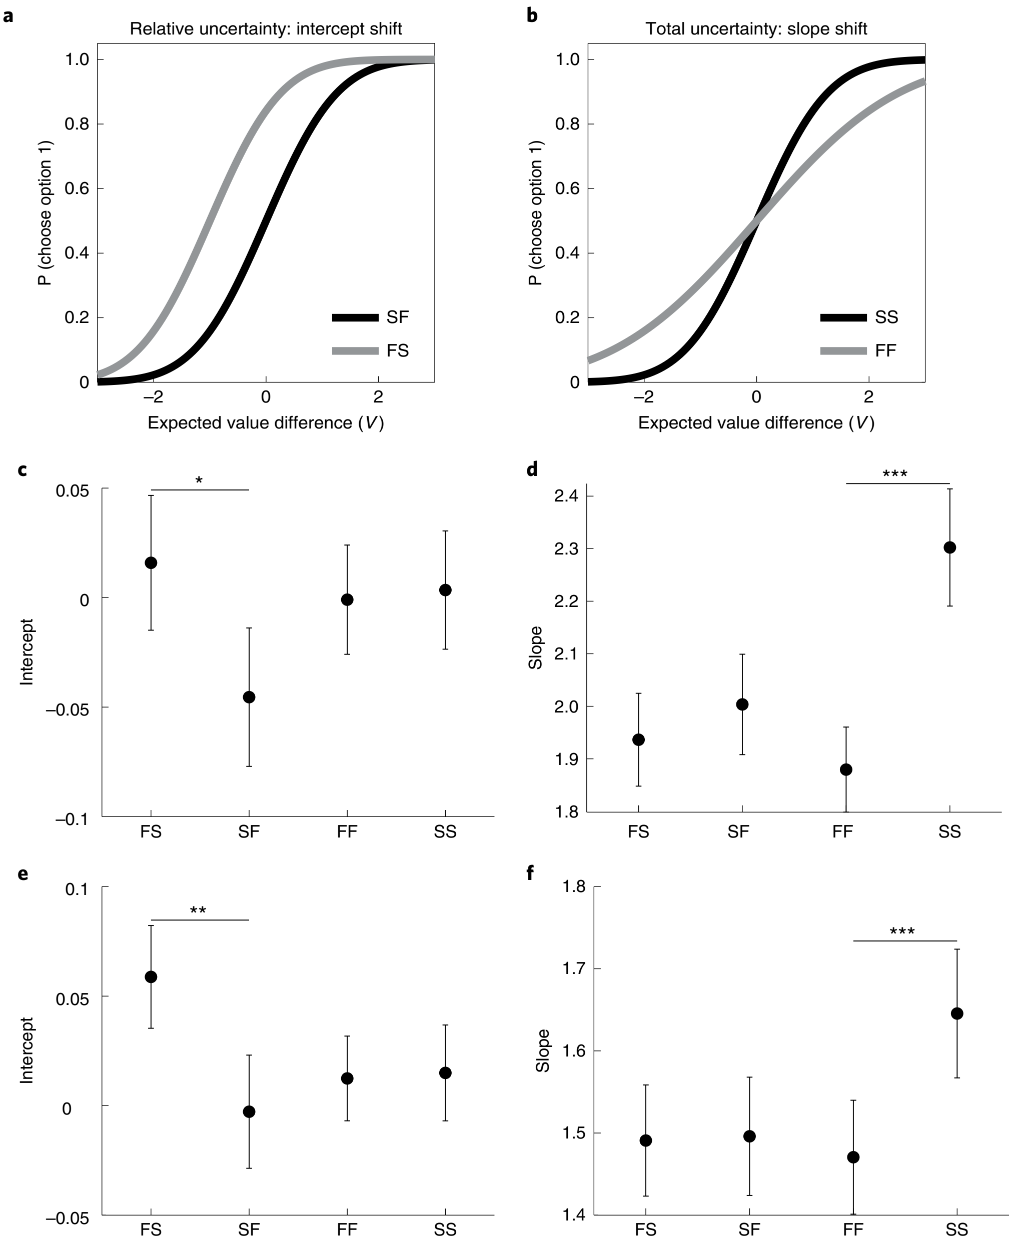

Fig. 2 | Predictions of choice probability function change across conditions and probit regression results. a, Directed exploration predicts a preference for the uncertain option, which manifests as a shift in intercept in opposite directions for SF and FS trials. P (choose option 1): probability of choosing the option on the left. **b**, Random exploration predicts more choice stochasticity when total uncertainty is high, equivalent to a steeper curve for FF trials than for SS trials. **c**–**f**, Across two experiments**:** Experiment 1, $N = 501$ (**c**,**d**) and Experiment 2, $N = 484$ (**e**,**f**), the intercept of FS trials was larger than that of SF trials (Experiment 1: $F(1, 150{,}292) = 6.20$, $P = 0.013$, $\Delta M = 0.061$; Experiment 2: $F(1, 144{,}892) = 10.17$, $P = 0.001$, $\Delta M = 0.062$), while the intercepts of FF and SS trials did not differ (Experiment 1: $F(1, 150{,}292) = 0.06$, $P = 0.814$, $\Delta M = 0.004$; Experiment 2: $F(1, 144{,}892) = 0.04$, $P = 0.834$, $\Delta M = 0.003$). The slope of FF trials was larger than that of SS trials (Experiment 1: $F(1, 150{,}292) = 81.07$, $P < 0.001$, $\Delta M = 0.423$; Experiment 2: $F(1, 144{,}892) = 31.23$, $P < 0.001$, $\Delta M = 0.175$), while the slopes of SF and FS trials did not differ (Experiment 1: $F(1, 150{,}292) = 2.77$, $P = 0.096$, $\Delta M = 0.067$; Experiment 2: $F(1, 144{,}892) = 0.03$, $P = 0.867$, $\Delta M = 0.005$). Points indicate the fixed effect coefficients of intercept and slope for each condition using maximum likelihood estimation. *F*-tests (two-sided) were used to compare intercepts and slopes across conditions. No multiple comparisons correction was applied. Error bars are 95% confidence intervals. *$P < 0.05$, **$P < 0.01$, ***$P < 0.001$.

The trial type remained the same within a block and varied randomly across blocks (that is, the types and locations of two arms were fixed within a block and participants encountered a new pair of arms in each block). Participants learned the statistics from instructions (that is, they were explicitly instructed on the difference between a stable and a fluctuating arm and did not need to learn it during the experiment) and had a chance to practice before the formal experiment.

This experimental design allows us to independently manipulate relative and total uncertainty to investigate their separate influences on different exploration strategies. Because the volatility of the fluctuating arm is higher than the stable arm, the fluctuating arm is more uncertain (Supplementary Fig. 8). If directed exploration is, as hypothesized, sensitive to relative uncertainty, people should prefer option 1 (the left arm) in condition FS and option 2 (the right arm) in condition SF, even the value difference suggests otherwise. This is equivalent to an increase in the intercept of the psychometric curve for condition FS vs SF, that is, the probability of choosing option 1 is higher in condition FS (Fig. 2a). It is worth noting that the comparison between condition FS and SF holds the total uncertainty constant because there is always one fluctuating arm and one stable arm (Supplementary Fig. 8).

https://doi.org/10.1038/s41562-022-01455-y

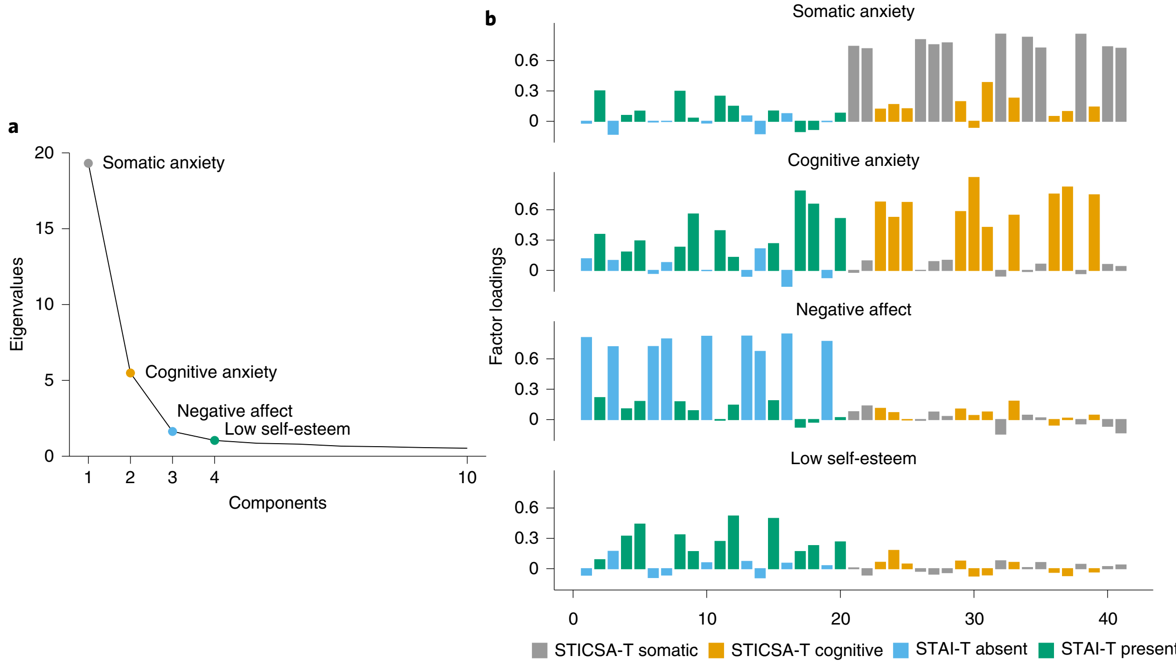

Fig. 3 | Exploratory factor analysis results (N = 501). **a**, Scree plot of eigenvalues. **b**, Factor loadings of items on trait anxiety factors. Items from different subscales are distinguished by their colours. STICSA-T somatic: somatic subscale of STICSA-T; STICSA-T cognitive: cognitive subscale of STICSA-T; STAI-T absent: STAI-T anxiety absent items; STAI-T present: STAI-T anxiety present items.

A comparison between conditions SS and FF reveals how people respond to a shift in total uncertainty (FF > SS) while holding the relative uncertainty. If random exploration is sensitive to total uncertainty, it should predict that people behave more randomly in condition FF. This can be reflected as an increase in slope in the psychometric curve for condition FF vs SS, that is, the probability of choosing option 1 is closer to random ($P = 0.5$) in condition FF (Fig. 2b).

Our data suggest that people performed well in the task, choosing the most rewarding option 79.76% (s.e.m. = 5.22%) of the time on average (Supplementary Fig. 1a). In line with our predictions, there was a significant intercept shift between the psychometric curves for condition FS and SF ($F(1, 150{,}292) = 6.20$, $P = 0.013$, $\Delta M = 0.061$, two-sided), showing that people directed their exploration towards the uncertain option (Fig. 2c and Supplementary Fig. 2a). Additionally, the intercept of SF was significantly negative ($t(150{,}292) = -2.82$, $P = 0.005$, $\beta = -0.045$, 95% CI (−0.077, −0.014), two-sided; see Supplementary Table 9 for the full regression result table), indicating that people formed a bias towards choosing option 2 when its relative uncertainty was high. The intercept of FS was not significantly different from zero ($t(150{,}292) = 1.01$, $P = 0.312$, $\beta = 0.016$, 95% CI (−0.015, 0.047)). There was no significant difference in intercept between conditions FF and SS ($F(1, 150{,}292) = 0.06$, $P = 0.814$, $\Delta M = 0.004$), suggesting that people did not adjust directed exploration according to total uncertainty.

We also found evidence for random exploration, manifesting as a steeper curve for FF than for SS ($F(1, 150292) = 81.07$, $P < 0.001$, $\Delta M = 0.423$; Fig. 2d and Supplementary Fig. 3b). This suggests that people chose more randomly when facing high total uncertainty, that is, random exploration is sensitive to total uncertainty. There was no significant difference in slope between SF and FS ($F(1, 150{,}292) = 2.77$, $P = 0.096$, $\Delta M = 0.067$), corroborating that random exploration is not sensitive to relative uncertainty. These behavioural signatures of directed and random exploration are present irrespective of whether all trials were used, or only trials in the late stage of each block (that is, trials 4–10; Supplementary Fig. 14).

To examine the exploration strategies people used on a trial-by-trial basis, we pooled participants' data across conditions and used the following probit regression model:

$$P(a_t = 1|\mathbf{w}) = \Phi\left(w_1 V_t + w_2 \mathrm{RU}_t + w_3 V_t/\mathrm{TU}_t\right), \qquad (1)$$

where $V_t = Q_t(1) - Q_t(2)$ is the value difference between two options, $\mathrm{RU}_t = \sigma_t(1) - \sigma_t(2)$ denotes the relative uncertainty, $\mathrm{TU}_t = \sqrt{(\sigma t^2(1) + \sigma t^2(2))}$ denotes the total uncertainty, and $\Phi(\cdot)$ is the standard Gaussian cumulative distribution function. $Q_t(k)$ and $\sigma_t(k)$ respectively denote the value estimate and estimate uncertainty of option k obtained using Kalman filtering (equations (2)–(5)). Past work[15] demonstrated that this is the exact analytic form of a hybrid of Thompson Sampling and Upper Confidence Bound algorithms. A positive w_2 means that people add an uncertainty bonus to an option's value proportional to its relative uncertainty, which directs exploration towards the option they are more uncertain about. In contrast, a positive w_3 indicates that as the total uncertainty in the environment goes up, people increase choice randomness accordingly. Specifically, if $w_2 = 0$, the model is insensitive to relative uncertainty and is reduced to pure random exploration. If $w_3 = 0$, the model is insensitive to total uncertainty and is reduced to pure directed exploration. Finally, if $w_2 = w_3 = 0$, the model is only influenced by the value of options. This value-driven strategy is similar to softmax exploration and we term it undirected exploration because it is not influenced by uncertainty in the environment.

As expected, we found that people were sensitive to RU, V/TU and V, with their fixed effect coefficients all significantly larger than zero (RU: $t(150{,}297) = 5.89$, $P < 0.001$, $\beta = 0.137$, 95% CI (0.092, 0.183); V/TU: $t(150{,}297) = 15.98$, $P < 0.001$, $\beta = 1.060$, 95% CI (0.930, 1.190); V: $t(150{,}297) = 34.43$, $P < 0.001$, $\beta = 1.576$, 95% CI (1.486, 1.666)). The model with three regressors ('hybrid model') outperformed model candidates that only allowed for one exploration strategy, other nested model candidates, and reinforcement learning models with constant learning rates (Supplementary Table 2). The hybrid model demonstrated good parameter recovery capability and did not introduce spurious

https://doi.org/10.1038/s41562-022-01455-y

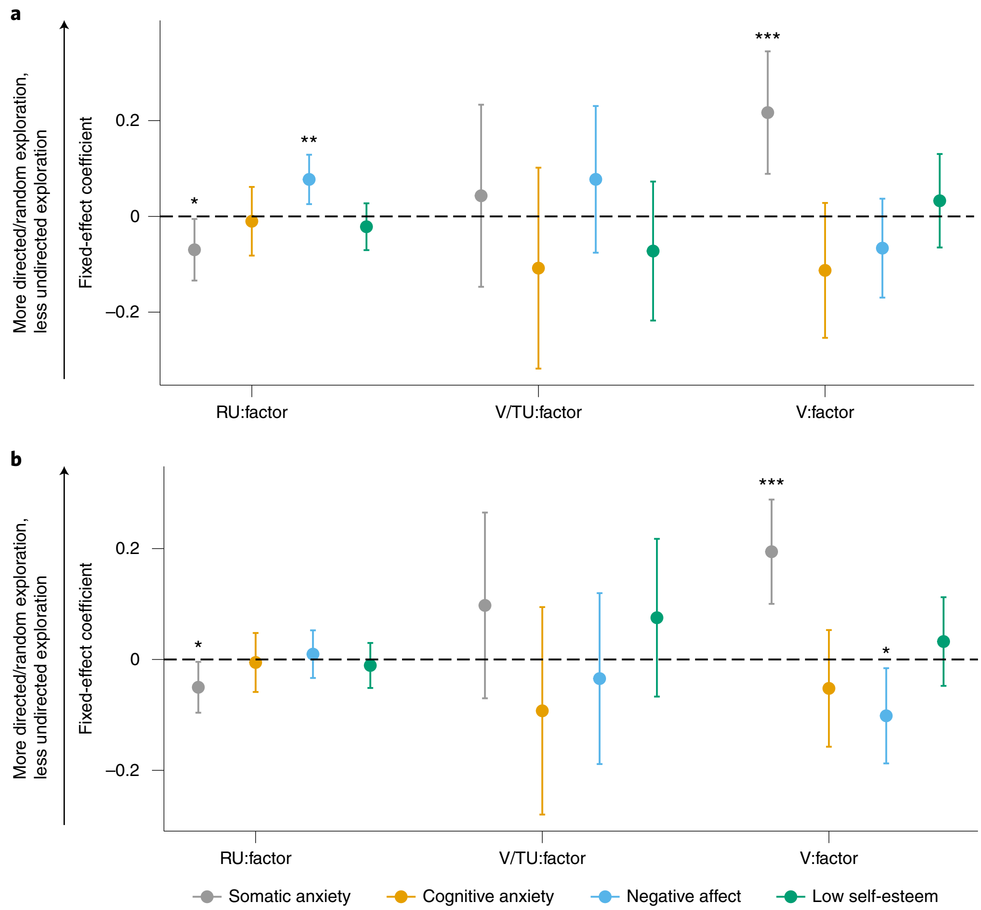

Fig. 4 | Effects of trait anxiety factors on exploration strategies. a, Data from Experiment 1 ($N = 501$). **b**, Data from Experiment 2 ($N = 484$). Factor scores were obtained using EFA results in Experiment 1. All factors were standardized and entered into the same model together with age and gender. Points indicate the fixed effect coefficients fit for each predictor using maximum likelihood estimation. Error bars represent 95% confidence intervals. A positive coefficient for RU:Factor (V/TU:Factor) indicates increased directed (random) exploration. A positive coefficient for V:Factor indicates decreased undirected exploration. Significance of the coefficients was assessed using t-test (two-sided). Multiple comparisons correction was not applied. Across the two experiments, Somatic anxiety factor negatively correlated with RU (Experiment 1: $t(150{,}273) = -2.12$, $P = 0.034$, $\beta = -0.070$, 95% CI (−0.134, −0.006); Experiment 2: $t(145{,}173) = -2.14$, $P = 0.032$, $\beta = -0.050$, 95% CI (−0.096, −0.004)) and V (Experiment 1: $t(150{,}273) = 3.32$, $P < 0.001$, $\beta = 0.217$, 95% CI (0.089, 0.345); Experiment 2: $t(145{,}173) = 4.05$, $P < 0.001$, $\beta = 0.194$, 95% CI (0.100, 0.289)). $^{*}P < 0.05$, $^{**}P < 0.01$, $^{***}P < 0.001$.

correlations between fitted parameters (correlation between fitted and simulated parameters: all $r_s > 0.99$, all $P_s < 0.001$; correlation between fitted parameters: all $P_s > 0.05$; Supplementary Fig. 6). It could also produce behavioural signatures of uncertainty-driven exploration, that is, intercept and slope differences in condition-based analysis comparable to human data (Supplementary Figs. 2 and 3). Together, our data support the hypothesis that people use a hybrid of directed and random exploration strategies, which are respectively sensitive to relative and total uncertainty induced by volatility.

Trait somatic anxiety is associated with reduced directed exploration and choice stochasticity

We administered State-Trait Anxiety Inventory trait scale (STAI-T)[42] and State-Trait Inventory for Cognitive and Somatic Anxiety trait scale (STICSA-T)[40] to capture the multidimensionality of trait anxiety. The questionnaires were filled out after the behavioural task (other scales administered are reported in Supplementary Methods). To extract latent factors of trait anxiety and reduce dimensionality, we conducted an exploratory factor analysis (EFA) on all items from STAI-T and STICSA-T ($N = 82$). The EFA resulted in a four-factor structure, which was validated with the confirmatory factor analysis in an independent sample ($N = 797$, see Supplementary Methods and Results for details). On the basis of the item loadings (Fig. 3 and Supplementary Table 3), we labelled the four factors as 'Somatic anxiety' (subjective experience of physiological symptoms), 'Cognitive anxiety' (worrying thoughts and rumination), 'Negative affect' (lack of positive affective experience) and 'Low self-esteem' (negative self-image).

We included trait anxiety factor scores and their interactions with V, RU and V/TU in equation (1) to model the effects of trait anxiety factors on exploration. We found that the Somatic anxiety factor had two distinct effects on exploration strategies. First, its interaction with relative uncertainty was significantly negative ($t(150{,}273) = -2.12$, $P = 0.034$, $\beta = -0.070$, 95% CI (−0.134, −0.006); Fig. 4a and Supplementary Fig. 10b). In other words, given the same level of relative uncertainty, people high on trait somatic anxiety were less likely to engage in directed exploration. Second, the Somatic anxiety factor was associated with increased sensitivity to value difference between options ($t(150{,}273) = 3.32$, $P < 0.001$, $\beta = 0.217$, 95% CI (0.089, 0.345); Supplementary Fig. 10a), indicating reduced undirected exploration. These effects were robust to different ways of trait somatic anxiety measurement, remaining significant when trait somatic anxiety was quantified by raw STICSA-T somatic subscale score (Supplementary Table 10). We also found a positive relationship between the Negative affect factor and directed exploration ($t(150{,}273) = 2.93$, $P = 0.003$,

 https://doi.org/10.1038/s41562-022-01455-y

Table 1 | Interactions between trait anxiety factors and exploration strategies

Trait anxiety factor	RU:Factor			V/TU:Factor			V:Factor		
	β(s.e.m.)	t	P	β(s.e.m.)	t	P	β(s.e.m.)	t	P
Experiment 1 (N=*501*)									
Somatic anxiety	−0.070(0.033)	−2.12	0.034*	0.043(0.097)	0.45	0.657	0.217(0.065)	3.32	< 0.001***
Cognitive anxiety	−0.010(0.037)	−0.28	0.782	−0.108(0.107)	−1.01	0.313	−0.113(0.072)	−1.57	0.117
Negative affect	0.077(0.026)	2.93	0.003**	0.077(0.078)	0.99	0.322	−0.066(0.053)	−1.26	0.208
Low self-esteem	−0.022(0.025)	−0.86	0.387	−0.072(0.074)	−0.98	0.328	0.033(0.05)	0.66	0.512
Experiment 2 (N=*484*)									
Somatic anxiety	−0.050(0.023)	−2.14	0.032*	0.098(0.086)	1.14	0.254	0.194(0.048)	4.05	<0.001***
Cognitive anxiety	−0.005(0.027)	−0.20	0.842	−0.093(0.096)	−0.97	0.332	−0.052(0.054)	−0.97	0.331
Negative affect	0.010(0.022)	0.44	0.662	−0.035(0.079)	−0.44	0.661	−0.102(0.044)	−2.32	0.020*
Low self-esteem	−0.011(0.021)	−0.52	0.602	0.075(0.073)	1.04	0.300	0.032(0.041)	0.79	0.428

Factor scores were obtained using EFA results in Experiment 1. All trait anxiety factors were standardized and entered into the same model together with age and gender. A positive coefficient for RU:Factor (V/TU:Factor) indicates increased directed (random) exploration. A positive coefficient for V:Factor indicates decreased undirected exploration. We used *t*-tests (two-sided) to assess the significance of the coefficients. Multiple comparisons correction was not applied. *$P<0.05$, **$P<0.01$, ***$P<0.001$

$\beta = 0.077$, 95% CI (0.026, 0.129); Fig. 4a and Supplementary Fig. 10e). However, this effect was not replicated in Experiment 2 (see below). No trait anxiety factor significantly influenced random exploration, and neither the Cognitive anxiety nor the Low self-esteem factor interacted with any exploration strategies (all $P_s > 0.05$, Table 1; see Supplementary Results and Table 11 for a similar analysis conducted using a regression model that does not include the random exploration component).

Replicating main results of Experiment 1 in Experiment 2

In Experiment 1, we documented an inverse relationship between trait somatic anxiety and directed and undirected exploration. However, it remained unclear whether the reduced exploration was attributed to a change in the process of learning (underestimated RU and/or overestimated V) or decision making (accurate value and uncertainty estimates but reduced sensitivity towards RU and/or increased sensitivity towards V), or both. In Experiment 2 ($N = 484$), we directly investigated this question by adding a reward prediction task at the end of each two-armed bandit task block (Fig. 1b). During the prediction task, participants guessed how many points one machine will generate in the next trial. They then rated their confidence in their reward prediction on a scale from 0 (guess randomly) to 10 (very confident).

We first repeated the analysis in Experiment 1 to make sure that adding the prediction task did not interfere with the exploration strategies people adopted in the two-armed bandit task. Indeed, the manipulation of relative and total uncertainty causally influenced uncertainty-driven exploration: people directed their exploration towards the relatively uncertain option ($F(1, 144{,}892) = 10.17$, $P = 0.001$, $\Delta M = 0.062$; Fig. 2e and Supplementary Fig. 2c). In addition, people chose more randomly when the total uncertainty was high ($F(1, 144{,}892) = 31.23$, $P < 0.001$, $\Delta M = 0.175$; Fig. 2f and Supplementary Fig. 2d). Examining participants' behaviour on a trial-by-trial basis, the model including V, RU and V/TU fit participants' behaviour better than models that nested within it, as well as reinforcement learning models with constant learning rates (Supplementary Table 2), again showing that people used a hybrid of undirected, directed and random exploration (V: $t(144{,}897) = 31.39$, $\beta = 1.153$, $P < 0.001$, 95% CI (1.081, 1.225); RU: $t(144{,}897) = 11.36$, $P < 0.001$, $\beta = 0.209$, 95% CI (0.175, 0.247); V/TU: $t(144{,}897) = 16.85$, $P < 0.001$, $\beta = 1.053$, 95% CI (0.938,1.185)).

Second, we replicated the inverse relationship between trait somatic anxiety and exploration. On the one hand, people high on the Somatic anxiety factor were less likely to direct exploration towards the option with high relative uncertainty ($t(145{,}173) = -2.14$, $P = 0.032$, $\beta = -0.050$, 95% CI (−0.096, −0.004); Fig. 4b, Table 1 and Supplementary Fig. 10h). On the other hand, they demonstrated a lower level of choice stochasticity, choosing the option with a higher expected value more often ($t(145{,}173) = 4.05$, $P < 0.001$, $\beta = 0.194$, 95% CI (0.100, 0.289); Supplementary Fig. 10g). The effects of Negative affect factor on exploration are inconsistent between Experiment 1 and Experiment 2: Experiment 2 data did not support a relationship between Negative affect and directed exploration ($t(145{,}173) = 0.44$, $P = 0.662$, $\beta = 0.010$, 95% CI (−0.033, 0.053)). Instead, there was a positive association between Negative affect and undirected exploration ($t(145{,}173) = -2.32$, $P = 0.020$, $\beta = -0.102$, 95% CI (−0.188, −0.016)).

Subjective value and uncertainty track the normative estimates

We treated participants' reward prediction as their subjective value estimation and their confidence rating as the inverse of their uncertainty estimation (that is, a high confidence level implies low estimation uncertainty). As expected, self-report reward predictions and confidence ratings tracked the normative value and uncertainty estimates well (reward prediction vs posterior mean: $t(24{,}935) = 113.023$, $P < 0.001$, $r_s\,(12{,}088) = 0.72$, 95% CI (0.71, 0.73); confidence vs posterior s.d.: $t(24{,}935) = -39.45$, $P < 0.001$, $r_s\,(12{,}088) = -0.24$, 95% CI (−0.26, −0.24)). The inverse relationship between normative uncertainty estimates and confidence is in line with past work showing that given Bayesian confidence hypothesis[43,44], confidence for the Gaussian estimation problem is monotonically related to the inverse posterior s.d.[45].

Because directed and undirected exploration strategies are sensitive to RU and V, we focused on the subjective estimates of RU and V (see Supplementary Methods and Results for a discussion on TU). The subjective counterpart of V and RU were defined as:

$$\text{Subjective } V = \text{option 1's reward prediction}$$
$$-\text{option 2's reward prediction,}$$

$$\text{Subjective RU} = \text{option 2's confidence} - \text{option 1's confidence.}$$

The subjective V and RU also significantly correlated with their normative counterparts (V: $t(24{,}935) = 123.297$, $P < 0.001$, $r_s\,(12{,}088) = 0.78$, 95% CI (0.77, 0.79); RU: $t(24{,}935) = 65.976$, $P < 0.001$, $r_s\,(12{,}088) = 0.42$, 95% CI (0.40, 0.43)).

Table 2 | The relationship between trait anxiety factors and subjective RU and V estimates

Trait anxiety factor	RU:Factor			V:Factor		
	β(s.e.m.)	t	P	β(s.e.m.)	t	P
Somatic anxiety	−0.293(0.06)	−4.85	<0.001***	−1.141(0.456)	−2.50	0.012*
Cognitive anxiety	0.087(0.076)	1.15	0.249	0.697(0.564)	1.24	0.217
Negative affect	0.186(0.065)	2.86	0.004**	0.575(0.488)	1.18	0.239
Low self-esteem	−0.084(0.06)	−1.41	0.159	0.06(0.444)	0.134	0.893

Factor scores were obtained using EFA results in Experiment 1. All trait anxiety factors were standardized and entered into the same model together with age and gender. The definition of subjective RU and V can be found in Experiment 2 Results section. A positive RU:Factor (V:Factor) indicates an overestimation of RU (V). We used t-tests (two-sided) to assess the significance of the coefficients. Multiple comparisons correction was not applied. *$P<0.05$, **$P<0.01$, ***$P<0.001$.

Trait somatic anxiety is associated with underestimation of relative uncertainty

To model the effects of trait anxiety on subjective V and RU, we examined how trait anxiety factors modulate the relationship between subjective and normative estimates of V and RU. The regression results revealed that the Somatic anxiety factor of trait anxiety interacted negatively with the normative estimate of relative uncertainty ($t(12{,}076)=-4.85$, $P<0.001$, $\beta=-0.293$, 95% CI (−0.411, −0.174); see Table 2), suggesting that people high on trait somatic anxiety tend to underestimate relative uncertainty in the environment. The effect of underestimation remained significant when using the STICSA-T somatic subscale as trait somatic anxiety measure ($t(12{,}082)=-4.20$, $P<0.001$, $\beta=-0.211$, 95% CI (−0.309, −0.112); see Supplementary Table 5). This could be a potential explanation for the diminished directed exploration in individuals high on trait somatic anxiety. There was also a significantly negative interaction between Somatic anxiety and the normative estimate of value difference ($t(12{,}076)=-2.50$, $P=0.012$, $\beta=-1.435$, 95% CI (−2.035, −0.246)), implying a tendency to underestimate value difference among those scoring high on Somatic anxiety factor. We failed to detect this interaction when measuring trait somatic anxiety with the STICSA-T somatic subscale ($t(12{,}082)=-1.41$, $P=0.158$, $\beta=-0.532$, 95% CI (−1.270, 0.206)).

Discussion

The current study examined the association between trait anxiety components and exploration during decision making under uncertainty. Across two large-scale online experiments, we found a selective role of trait somatic anxiety—the propensity to experience enduring physical symptoms of anxiety—in modulating exploration in two distinct ways: first, trait somatic anxious individuals showed diminished directed exploration, being less likely to direct exploration towards the more uncertain option. The reward prediction task data offered a potential mechanistic to account for this inverse relationship, showing that people high on trait somatic anxiety tended to underestimate relative uncertainty. Second, people high on trait somatic anxiety exhibited a low level of undirected exploration, choosing less randomly regardless of uncertainty. Interestingly, none of the other trait anxiety components interact with any exploration strategy.

To dissociate and quantify exploration strategies, we manipulated different kinds of uncertainty in a two-armed bandit task. Dovetailing with previous work[23–25], we found behavioural signatures for directed, random and undirected exploration, which were driven by relative uncertainty, total uncertainty and value difference, respectively. Specifically, the current study altered uncertainty by changing volatility (that is, the diffusion rate of an option's generative mean), while holding the riskiness of both options fixed (that is, the variance of the outcome distribution). This design removed the potential interference effect of risk aversion. In other words, people may be reluctant to try an uncertain option due to a desire to avoid risk, therefore offering a cleaner measure of directed and random exploration. More importantly, these findings extend the uncertainty-driven exploration framework by proving that people are sensitive to epistemic uncertainty in their posterior estimates of the choices instead of aleatoric uncertainty caused merely by outcome variability.

The negative association between trait somatic anxiety and directed exploration is consistent with previous work showing that anxious individuals find uncertainty more aversive and tend to avoid it[9,10]. In a volatile environment, this avoidance behaviour is self-reinforcing. As time goes by, the unchosen option grows more uncertain because the information obtained from the last observation is less predictive of its current value, which further drives a somatic anxious individual away from it. In the long term, insufficient directed exploration impedes people from collecting information and updating the value of the uncertain options, which could lead to suboptimal decision strategies and give rise to real-life maladaptive avoidance behaviour. Therefore, reduced directed exploration can be seen as a cognitive risk factor for the later development of anxiety disorders. On the other hand, diminished undirected exploration suggests that people scoring high on trait somatic anxiety made choices in line with the value difference. In tandem with other cognitive biases, reduced indirect exploration could also lead to avoidance behaviour. For example, suppose an anxious individual holds a pessimistic belief towards novel objects and assigns low initial values to them[46], they may tend to avoid the unfamiliar options without a single interaction. Furthermore, if the anxious individual believes that interacting with the novel object will lead to negative outcomes, their degree of uncertainty avoidance could be exaggerated due to an elevated prepotent bias to withhold responding in the face of negativity[47].

Complementary to avoiding the uncertain option, trait somatic anxiety was linked to a reduced ability to track relative uncertainty with self-report confidence. In other words, when holding the uncertainty level of one option unchanged, an increase in the uncertainty level of the other option transforms to a smaller change in confidence difference between two options in somatic anxious individuals. This altered representation of uncertainty converges with previous research showing that state anxious individuals hold a more precise belief about the reward contingencies in the environment and update less[30]. Moreover, recent work uniquely linked the somatic component of anxiety to a smaller uncertainty estimate during aversive learning, while the direction of the relationship was reversed for cognitive anxiety[41]. Building upon these findings, the current study further showed that somatic anxious individuals were poor at sensing the uncertainty difference between options. The blunted response to uncertainty could serve as an explanation for the reduced directed exploration: an anxious person may favour the uncertain option less due to a smaller uncertainty bonus, which is proportional to subjective relative uncertainty. However, the current data could not directly test this hypothesis because confidence ratings were only acquired at the end of each block instead of on a trial-by-trial basis. The task was designed this way due to the concern that eliciting reward prediction and confidence rating too often would interfere with the exploration process. Because there is no clear link between directed exploration during the block and relative uncertainty at the end of the block, we could not directly tie subjective report in with exploration in the current prediction task. Future studies could use different paradigms to formally examine whether subjective relative uncertainty mediates the relationship between trait somatic anxiety and directed exploration. Besides underestimating uncertainty difference, trait somatic anxious individuals underestimate value difference as well. The latter relationship is not as robust as the former since it depends on the way trait somatic anxiety is measured (that is, factor score or STICSA-T questionnaire). Nonetheless, it suggests that reduced undirected exploration cannot be explained from the perspective of biased subjective estimation, which would predict overestimating value difference. In this case, we conjecture that trait somatic anxiety is associated with an increased sensitivity to the value difference between options. The current task design cannot directly test this hypothesis, which is one limitation of Experiment 2. Future studies interested in testing this hypothesis could explicitly instruct the participants about the value difference between options and examine the relationship between brain activity during the choice phase and trait somatic anxiety after controlling for the value difference.

The current study utilizes a dimensional approach to examine the interaction between anxiety and exploration. Among all latent factors of trait anxiety identified in the present study, only Somatic anxiety was associated with a change in exploration. In contrast, the others did not reliably modulate the exploration process. These differential effects are critical in understanding the inconsistencies in past work on the anxiety-exploration relationship. Most existing studies treated trait anxiety as a unitary construct, while their measurements might tap onto distinct components. According to our factor structure, STAI-T primarily reflects the Negative affect factor, which was associated with an increase in directed exploration in Experiment 1. This finding matches recent work[11]. However, this relationship disappeared in Experiment 2 when the reward prediction task was added, implying that the effect of Negative affect might not be robust to task modification. Supporting this argument, another recent study did not find a significant relationship between trait anxiety, measured by STAI-T, and uncertainty-driven exploration strategies[48]. Another recent study[49] administered the Penn State Worry Questionnaire[50], which primarily measures cognitive symptoms of anxiety[51], and failed to detect a relationship between anxiety and directed exploration. This finding aligns well with our finding that trait cognitive anxiety did not interact with exploration. On the other hand, past literature on physical exploratory behaviour focused on patients diagnosed with panic disorder or agoraphobia[6], both of which are anxiety disorders primarily characterized by elevated somatic symptoms[52]. Specifically, Kallai and colleagues[53] found that compared to those diagnosed with generalized anxiety disorder, people who have panic disorder with agoraphobia demonstrated more disturbed exploratory behaviour, hinting at a unique contribution of somatic anxiety to suppressed spatial exploration. Since physical exploration usually requires going through novel environments, it is possible that it shows a flavour of directed exploration, which is decreased in trait somatic anxious individuals. More studies and analysis need to be done to understand the similarity and difference between spatial exploration and exploration during decision making. Nonetheless, our dimensional approach showed the promise of unifying past work on exploration across task domains.

An important question remains unanswered: why does trait somatic anxiety interact with exploration? Despite the prevalence of somatic symptoms in anxiety disorders, it is only recently that researchers have begun to examine its specific impact on learning and decision making in the general population[41,54]. Drawing inspirations from work on panic disorder, we conjecture that altered exploration in trait somatic anxiety could be attributed to enhanced interoception, that is, increased sensitivity to arousal signals[55–57]. These signals are important because arousal systems are involved in the processing of uncertainty and value, which could be used to guide decisions[58–61]. Past work on the concordance between trait somatic anxiety and physiological arousal has documented mixed findings, suggesting that high trait somatic anxiety is not necessarily accompanied by an increase in physiological indices of arousal[62–64]. Instead, we speculate that heightened interoception in somatic anxious individuals might amplify the impact of physiological arousal on exploration[65], which could account for the altered subjective computation of relative uncertainty in the same population. This would be in line with past work documenting a mediating role of trait anxiety on the relationship between physiological responses and risk decision-making[60]. It is left for future studies to provide a holistic picture of the dynamics between somatic anxiety, physiological and subjective arousal, and exploration.

This study has a few limitations. Although we demonstrated the specificity of the associations between trait somatic anxiety and exploration strategies by controlling for other trait anxiety components in the same analysis, we do not know whether these effects are robust to including other psychiatric symptoms such as compulsivity. We believe that our results will still hold since no clear relationship between compulsive behaviour and exploration has been documented in the literature. However, given the high comorbidity of obsessive-compulsive disorder and anxiety[66], it is important for future studies to take compulsivity into account to obtain cleaner measures of the effects of anxiety on decision-making components. Besides, the current study examined the anxiety-exploration relationship in an online general population where the distribution of trait somatic anxiety was positively skewed. Therefore, future work should test the generalizability of our findings in a clinically diagnosed population, such as panic disorder patients who usually demonstrate more somatic symptoms.

Methods

The study was approved by the Harvard University Committee on the Use of Human Subjects (IRB19-0789). We report how we determined our sample size, all data exclusions, all manipulations and all measures for both experiments in the study. All data and analysis code for both experiments are available online (see Data availability statement). The experiment design, sample size, exclusion criteria and primary data analysis plan for Experiment 1 were pre-registered at https://aspredicted.org/vi8wg.pdf.

Participants

The study was approved by the Harvard University Committee on the Use of Human Subjects. Participants were recruited via the Prolific Platform. Informed consent was given before testing. Before Experiment 1, we conducted a power analysis by simulating the pilot data, which revealed that a sample size of 500 (after exclusion) would be necessary to detect an effect with 80% power. The exclusion criteria were pre-registered and participants were excluded if they: reported an age outside the range 18–65, chose the more rewarding option in <60% of trials, or did not complete anxiety-related questionnaires. The accuracy criterion was set to make sure that participants paid attention to the task and have shown effective learning, which is in line with previous studies using a similar task setting[67]. We recruited 1,097 participants in total (Experiment 1: $N = 531$; Experiment 2: $N = 566$) and 985 participants were included in the final analysis of the bandit task (Experiment 1: $N = 501$, 219 women, 277 men, 5 unreported; age $M = 36.1$ years, s.d. = 10.9; Experiment 2: $N = 484$, 197 women, 279 men, 8 unreported; Age $M = 35.3$ years, s.d. = 10.5). As for the data analysis of the prediction task in Experiment 2, we excluded prediction trials in which the reward prediction was >50 or <−50 (the maximum and minimum possible experienced reward in the task were 47 and −44). This resulted in the exclusion of 10.52% of all trials. Participants who always selected the same trial-by-trial confidence rating were also excluded (29 participants, 4.03% of all trials). In the final analysis of the prediction task, 446 participants were included (190 women, 248 men, 8 unreported; age $M = 34.6$ years, s.d. = 10.4). Participants received monetary compensation (Experiment 1: US\$7; Experiment 2: US\$15) and could earn a bonus depending on their performance in the bandit task (up to US\$3; Experiment 1 and 2) and their reward prediction accuracy (up to US\$1; Experiment 2).

Experiment design

Two-armed bandit task. In Experiment 1 and 2, participants performed 30 blocks of a two-armed bandit task adapted from ref. [23]. Each block consists of 10 trials. In each block, participants encountered a new pair of arms and chose between them. Both the fluctuating and the stable arms delivered rewards (rounded to the nearest integer) drawn from a Gaussian distribution (variance $\tau^2 = 1$) around its current generative mean. At the beginning of each block, the generative means of both arms were reset, drawing randomly from a Gaussian distribution (variance $\tau_0^2 = 100$) with mean 0. The generative mean of the fluctuating arm diffused in a Gaussian random walk, that is, its mean on trial $t+1$ was drawn from a Gaussian distribution (variance $\tau_{mean}^2(F) = 4$) centred on its mean on trial t. In contrast, the mean of the stable arm was fixed within a block, that is, the variance of the mean $\tau_{\mathrm{mean}}^2(S)$ is 0. A demonstration of the bandit task can be found online: https://9kqpbf3ddo.cognition.run.

Prediction task. In Experiment 2, participants played the two-armed bandit task for 30 blocks and completed a reward prediction task at the end of each block. During the prediction task, participants reported their estimate of the number of points that each machine will generate in the next trial by entering an integer (see Supplementary Methods for details). Feedback on their prediction was not provided. Participants were encouraged to use the full range of the scale (0-guess randomly; 10-very confident). After the prediction task, participants encountered two new slot machines and started a new block of the bandit task. Whether they made a prediction for the left or the right arm first is randomized. At the end of the experiment, one prediction task trial was randomly chosen. If the prediction made on that trial was within one of the real generative mean value of the arm, the participant received a monetary bonus, which is in addition to the bonus depending on their performance.

Belief update process

In line with previous literature[14,15,23,68], we assumed that participants approximate an ideal Bayesian observer and track both the expected value and uncertainty in the estimation. Given the Gaussian distributional structure underlying our task, the posterior over the value of arm k is Gaussian with mean $Q_t(k)$ and variance $\sigma_t^2(k)$. Both can be recursively updated on each trial t using the Kalman filtering equations:

$$Q_{t+1}(a_t) = Q_t(a_t) + \alpha_t(r_t - Q_t(a_t)), \tag{2}$$

$$\sigma_{t+1}^2(k) = \sigma_t^2(k) - \alpha_t\sigma_t^2(k) + \tau_{\mathrm{mean}}^2(k) \text{ if } arm\ k \text{ is chosen}, \tag{3}$$

$$= \sigma_t^2(k) + \tau_{\mathrm{mean}}^2(k) \text{ if arm } k \text{ is not chosen}, \tag{4}$$

where α_t is the chosen arm ($\alpha_t = 1$ if the left arm is chosen, $\alpha_t = 2$ if the right arm is chosen), r_t is the delivered reward, and the Kalman gain α_t is given by:

$$\alpha_t = (\sigma_t^2(k) + \tau_{\mathrm{mean}}^2(k)) / (\sigma_t^2(k) + \tau_{\mathrm{mean}}^2(k) + \tau^2(a_t)) \tag{5}$$

Note that the posterior mean is updated only for the chosen arm regardless of its type, whereas the posterior variance for the fluctuating arm is updated every trial irrespective of the participant's choice. The diffusion noise for the fluctuating arm $\tau_{\mathrm{mean}}^2(F)$ and for the stable arm $\tau_{\mathrm{mean}}^2(S)$ are set to 4 and 0, respectively. The posterior means are initialized with the prior mean, $Q_1(k) = 0$ for all k, and posterior variances are initialized with the prior variance, $\sigma_1^2(k) = \tau_0^2 = 100$. Kalman filtering is an idealization of learning from noisy observations, and past research has shown that it can account for human choices in bandit tasks well[14,68].

Choice probability analysis

The parameters of the full hybrid model (equation (1)) were estimated using maximum likelihood estimation in MATLAB (function fitglme, version R2020a). In line with models used in previous work, the current model included fixed and random effects for V, RU and V/TU. In Wilkinson notation, the model specification was Choice ~ V+RU+V/TU+(V+RU+V/TU|SubjectID). All fixed effects were unbounded, and the random effects were restrained to be coming from a multivariate Gaussian distribution (see Supplementary Fig. 3 for the distributions of random effects). The residual plot of the model and the distribution of random effects (Supplementary Fig. 4) were visually inspected to confirm that the assumptions of generalized linear mixed models were met (same for all mixed probit regression models mentioned below) and was therefore not formally tested for normality or equal variance. To understand what exploration strategies participants used, we compared the full hybrid model (with regressors V, RU and V/TU) with models nested within it, resulting in a set of 7 models. Given that these models differed in their number of parameters, we used the Bayesian Information Criterion (BIC) to compare models; this metric penalizes models on the basis of the number of parameters (Supplementary Table 2). A model recovery analysis was carried out to confirm that BIC was a reliable model selection metric, accurately selecting the true generative model from the whole set of model candidates (Supplementary Fig. 11). A parameter recovery test was conducted on the winning model to confirm its ability to recover true parameters that were used to generate behaviour. Specifically, we sampled w_1, w_2 and w_3 from $N(0, 10 \times \mathrm{I})$, simulated behaviour, fitted the model and compared the generative and fitted parameters (Supplementary Fig. 6). This process was repeated 1,000 times. Data simulated using the winning model were analysed using the condition-based choice probability analysis to examine whether the simulated data showed behavioural signatures of directed and random exploration and can be distinguished from data generated by other models (see Supplementary Methods and Results for details).

To obtain psychometric curves of choice behaviour across conditions, we modelled choices as a function of experimental condition (SF, FS, SS, FF):

$$P(a_t = 1|\mathbf{w}) = \Phi\left(\sum_j w_4^j \pi_{tj} + w_5^j \pi_{tj} V_t\right), \tag{6}$$

where j is the experimental condition, and $\pi_{tj} = 1$ if trial t belongs to condition j and $= 0$ otherwise. We refer to the w_4 terms as intercepts and the w_5 terms as slopes. w_4 and w_5 were estimated using a generalized linear mixed-effect model with a probit link function. In Wilkinson notation, the model specification was Choice ~ V + V:cond + (V+V:cond|Subject ID). All t-tests and F-tests for fixed effect coefficients are two-sided.

Trait anxiety measure

STAI-T and STICSA-T were used to measure trait anxiety. STAI-T is the most widely used trait anxiety measure[42] and has a high internal consistency (Cronbach's $\alpha = 0.90$). Compared with STAI-T, STICSA-T includes more somatic items and has superior divergent validity[69,70]. To obtain latent factors of trait anxiety, we conducted an exploratory factor analysis on all items from STICSA-T and STAI-T ($N = 82$). The number of factors was selected on the basis of parallel analysis, which compared the eigenvalues generated from the data matrix to those generated from a simulated matrix created from random data of the same size and retained the eigenvalues that are above the 95% quantile of the simulated ones[71]. Factor scores were then calculated using the Bartlett method[72] implemented in the psych package[73] in R (version 4.0.2). In Experiment 2, latent factors of trait anxiety were extracted using the factor loading structure obtained in Experiment 1.

Modelling the effects of trait anxiety factors on exploration strategies

We extended equation (1) to include trait anxiety factor scores and their interactions with V, RU and V/TU:

$$P(a_t = 1|\mathbf{w}) = \Phi\left(w_1 V_t + w_2 \mathrm{RU}_t + w_3 V_t/\mathrm{TU}_t + \sum_n w_{11}{}^n V : \mathrm{Anx}^n + w_{21}{}^n \mathrm{RU}_t : \mathrm{Anx}^n + w_{31}{}^n V_t/\mathrm{TU}_t : \mathrm{Anx}^n + w_6{}^n \mathrm{Anx}^n\right), \tag{7}$$

where n denotes the n^{th} trait anxiety factor. The factor scores were standardized before entering in the model. Age and gender were also entered as covariates. A positive w_{21} indicates that people scoring high on the trait anxiety factor are more sensitive to relative uncertainty and engage in more directed exploration. Similarly, a positive w_{31} implies elevated random exploration (that is, behaving more randomly when facing the same total uncertainty). On the other hand, a positive w_{11} is interpreted as a negative association between the trait anxiety factor and undirected exploration due to an increased sensitivity to the value difference between options. The parameters were estimated using generalized linear mixed-effect models with a probit link in MATLAB. In Wilkinson notation, the model specification was Choice ~ V + RU + V/TU + V:Anx + RU:Anx + V/TU:Anx + Anx + (V + RU+V/TU|SubjectID).

Modelling the effects of trait anxiety factors on subjective value and uncertainty estimates

We used the following linear mixed-effect regressions to examine how trait anxiety factors modulate the relationship between subjective and normative estimates of V and RU:

$$\text{Subjective } V_t = w_7 V_t + \sum\nolimits_n \left(w_{71}{}^n V_t : \mathrm{Anx}^n + w_8 \mathrm{Anx}^n\right), \tag{8}$$

$$\text{Subjective } \mathrm{RU}_t = w_9 \mathrm{RU}_t + \sum\nolimits_n \left(w_{91}{}^n \mathrm{RU}_t : \mathrm{Anx}^n + w_0 \mathrm{Anx}^n\right) \tag{9}$$

The factor scores were standardized and age and gender were entered as covariates. Across two equations, a positive coefficient for the interaction term, that is, w_{71} in equation (8) and w_{91} in equation (9), means that given the same normative estimates of V (RU), participants high on the trait anxiety factor report higher subjective V (RU), which is equivalent to an overestimation. The parameters were estimated using linear mixed-effect models in MATLAB. In Wilkinson notation, the model specification for equation (8) was Subjective V ~ V + V:Anx + Anx + (*V* | SubjectID), and Subjective RU ~ RU + RU:Anx + Anx + (RU | SubjectID) for equation (9).

Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

All de-identified data are publicly available at the Open Science Framework website: https://osf.io/y6urc/.

Code availability

The code used to fit belief update model, generate regression models and generate figures are publicly available at the Open Science Framework https://osf.io/y6urc/.

References

1. Lister, R. G. The use of a plus-maze to measure anxiety in the mouse. *Psychopharmacology* **92**, 180–185 (1987).
2. Prut, L. & Belzung, C. The open field as a paradigm to measure the effects of drugs on anxiety-like behaviors: a review. *Eur. J. Pharmacol.* **463**, 3–33 (2003).
3. Zweifel, L. S. et al. Activation of dopamine neurons is critical for aversive conditioning and prevention of generalized anxiety. *Nat. Neurosci.* **14**, 620–626 (2011).
4. Britton, D. R. & Britton, K. T. A sensitive open field measure of anxiolytic drug activity. *Pharmacol. Biochem. Behav.* **15**, 577–582 (1981).
5. Biedermann, S. V. et al. An elevated plus-maze in mixed reality for studying human anxiety-related behavior. *BMC Biol.* **15**, 125 (2017).
6. Walz, N., Mühlberger, A. & Pauli, P. A human open field test reveals thigmotaxis related to agoraphobic fear. *Biol. Psychiatry* **80**, 390–397 (2016).
7. Schulz, E. & Gershman, S. J. The algorithmic architecture of exploration in the human brain. *Curr. Opin. Neurobiol.* **55**, 7–14 (2019).
8. Levinthal, D. A. From arms to trees: opportunity costs and path dependence and the exploration-exploitation tradeoff. *Strategy Sci.* **6**, 331–337 (2021).
9. Grupe, D. W. & Nitschke, J. B. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nat. Rev. Neurosci.* **14**, 488–501 (2013).
10. Charpentier, C. J., Aylward, J., Roiser, J. P. & Robinson, O. J. Enhanced risk aversion, but not loss aversion, in unmedicated pathological anxiety. *Biol. Psychiatry* **81**, 1014–1022 (2017).
11. Aberg, K. C., Toren, I. & Paz, R. A neural and behavioral trade-off between value and uncertainty underlies exploratory decisions in normative anxiety. *Mol. Psychiatry* **27**, 1573–1587 (2022).
12. Bennett, D., Sutcliffe, K., Tan, N. P. J., Smillie, L. D. & Bode, S. Anxious and obsessive-compulsive traits are independently associated with valuation of noninstrumental information. *J. Exp. Psychol. Gen.* **150**, 739–755 (2021).
13. Aylward, J. et al. Altered learning under uncertainty in unmedicated mood and anxiety disorders. *Nat. Hum. Behav.* **3**, 1116–1123 (2019).
14. Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879 (2006).
15. Gershman, S. J. Deconstructing the human algorithms for exploration. *Cognition* **173**, 34–42 (2018).
16. Schwartenbeck, P. et al. Computational mechanisms of curiosity and goal-directed exploration. *eLife* **8**, e41703 (2019).
17. Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A. & Cohen, J. D. Humans use directed and random exploration to solve the explore–exploit dilemma. *J. Exp. Psychol. Gen.* **143**, 2074–2081 (2014).
18. Wilson, R. C., Bonawitz, E., Costa, V. D. & Ebitz, R. B. Balancing exploration and exploitation with information and randomization. *Curr. Opin. Behav. Sci.* **38**, 49–56 (2021).
19. Dubois, M. et al. Human complex exploration strategies are enriched by noradrenaline-modulated heuristics. *eLife* **10**, e59907 (2021).
20. Auer, P. Using confidence bounds for exploitation-exploration trade-offs. *J. Mach. Learn. Res.* **3**, 397–422 (2002).
21. Srinivas, N., Krause, A., Kakade, S. M. & Seeger, M. W. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Trans. Inf. Theory* **58**, 3250–3265 (2012).
22. Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* **25**, 285–294 (1933).

23. Gershman, S. J. Uncertainty and exploration. *Decision* **6**, 277–286 (2019).
24. Gershman, S. J. & Tzovaras, B. G. Dopaminergic genes are associated with both directed and random exploration. *Neuropsychologia* **120**, 97–104 (2018).
25. Tomov, M. S., Truong, V. Q., Hundia, R. A. & Gershman, S. J. Dissociable neural correlates of uncertainty underlie different exploration strategies. *Nat. Commun.* **11**, 1–12 (2020).
26. Piray, P. & Daw, N. D. A model for learning based on the joint estimation of stochasticity and volatility. *Nat. Commun.* **12**, 1–16 (2021).
27. Beltzer, M. L., Adams, S., Beling, P. A. & Teachman, B. A. Social anxiety and dynamic social reinforcement learning in a volatile environment. *Clin. Psychol. Sci.* **7**, 1372–1388 (2019).
28. Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X. & Bishop, S. J. Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nat. Neurosci.* **18**, 590–596 (2015).
29. Gagne, C., Zika, O., Dayan, P. & Bishop, S. J. Impaired adaptation of learning to contingency volatility in internalizing psychopathology. *eLife* **9**, e61387 (2020).
30. Hein, T. P., de Fockert, J. & Ruiz, M. H. State anxiety biases estimates of uncertainty and impairs reward learning in volatile environments. *NeuroImage* **224**, 117424 (2021).
31. Huang, H., Thompson, W. & Paulus, M. P. Computational dysfunctions in anxiety: failure to differentiate signal from noise. *Biol. Psychiatry* **82**, 440–446 (2017).
32. Pulcu, E. & Browning, M. Affective bias as a rational response to the statistics of rewards and punishments. *eLife* **6**, e27879 (2017).
33. Pulcu, E. & Browning, M. The misestimation of uncertainty in affective disorders. *Trends Cogn. Sci.* **23**, 865–875 (2019).
34. Buss, A. H. Two anxiety factors in psychiatric patients. *J. Abnorm. Soc. Psychol.* **65**, 426–427 (1962).
35. Hamilton, M. A. X. The assessment of anxiety states by rating. *Br. J. Med. Psychol.* **32**, 50–55 (1959).
36. Lang, P. J. The mechanics of desensitization and the laboratory study of human fear. *Behavior Therapy: Appraisal and Status. New York: McGraw-Hill,* 160–191 (1969).
37. Schwartz, G. E., Davidson, R. J. & Goleman, D. J. Patterning of cognitive and somatic processes in the self-regulation of anxiety: effects of meditation versus exercise. *Psychosom. Med.* **40**, 321–328 (1978).
38. Wall, A. D. & Lee, E. B. What do anxiety scales really measure? An Item content analysis of self-report measures of anxiety. *J. Psychopathol. Behav. Assess.* https://doi.org/10.1007/s10862-022-09973-9 (2022).
39. DeGood, D. E. & Tait, R. C. The cognitive-somatic anxiety questionnaire: psychometric and validity data. *J. Psychopathol. Behav. Assess.* **9**, 75–87 (1987).
40. Ree, M. J., French, D., MacLeod, C. & Locke, V. Distinguishing cognitive and somatic dimensions of state and trait anxiety: development and validation of the state-trait inventory for cognitive and somatic anxiety (STICSA). *Behav. Cogn. Psychother.* **36**, 313–332 (2008).
41. Wise, T. & Dolan, R. J. Associations between aversive learning processes and transdiagnostic psychiatric symptoms in a general population sample. *Nat. Commun.* **11**, 1–13 (2020).
42. Spielberger, C. D. Manual for the State-trait Anxiety Inventory (STAI Form Y), consuting Psychologists Palo Alto. Consulting Psychologists Press (1983)
43. Aitchison, L., Bang, D., Bahrami, B. & Latham, P. E. Doubly Bayesian analysis of confidence in perceptual decision-making. *PLoS Comput. Biol.* **11**, e1004519 (2015).
44. Fleming, S. M. & Daw, N. D. Self-evaluation of decision-making: a general Bayesian framework for metacognitive computation. *Psychol. Rev.* **124**, 91–114 (2017).
45. Xiang, Y., Graeber, T., Enke, B. & Gershman, S. J. Confidence and central tendency in perceptual judgment. *Atten. Percept. Psychophys.* **83**, 3024–3034 (2021).
46. Zorowitz, S., Momennejad, I. & Daw, N. D. Anxiety, avoidance, and sequential evaluation. *Comput. Psychiatr.* **4**, 1–17 (2020).
47. Mkrtchian, A., Aylward, J., Dayan, P., Roiser, J. P. & Robinson, O. J. Modeling avoidance in mood and anxiety disorders using reinforcement learning. *Biol. Psychiatry* **82**, 532–539 (2017).
48. Smith, R. et al. Lower levels of directed exploration and reflective thinking are associated with greater anxiety and depression. *Front. Psychiatry* **12**, 782136 (2022).
49. Zaller, I., Zorowitz, S. & Niv, Y. Information seeking on the horizons task does not predict anxious symptomatology. *Biol. Psychiatry* **89**, S217–S218 (2021).
50. Meyer, T. J., Miller, M. L., Metzger, R. L. & Borkovec, T. D. Development and validation of the Penn State worry questionnaire. *Behav. Res. Ther.* **28**, 487–495 (1990).
51. Sharp, P. B., Miller, G. A. & Heller, W. Transdiagnostic dimensions of anxiety: neural mechanisms, executive functions, and new directions. *Int. J. Psychophysiol.* **98**, 365–377 (2015).
52. *Diagnostic and Statistical Manual of Mental Disorders: DSM-5*, Vol. 5 (American Psychiatric Association, 2013).
53. Kallai, J. et al. Spatial exploration behaviour in an extended labyrinth in patients with panic disorder and agoraphobia. *Psychiatry Res.* **149**, 223–230 (2007).
54. Gillan, C. M. et al. Experimentally induced and real-world anxiety have no demonstrable effect on goal-directed behaviour. *Psychol. Med.* **51**, 1467–1478 (2021).
55. Domschke, K., Stevens, S., Pfleiderer, B. & Gerlach, A. L. Interoceptive sensitivity in anxiety and anxiety disorders: an overview and integration of neurobiological findings. *Clin. Psychol. Rev.* **30**, 1–11 (2010).
56. Ehlers, A. Cognitive factors in panic attacks: symptom probability and sensitivity. *J. Cogn. Psychother.* **5**, 157–173 (1991).
57. Yoris, A. et al. The roles of interoceptive sensitivity and metacognitive interoception in panic. *Behav. Brain Funct.* **11**, 1–6 (2015).
58. de Berker, A. O. et al. Computations of uncertainty mediate acute stress responses in humans. *Nat. Commun.* **7**, 1–11 (2016).
59. FeldmanHall, O., Glimcher, P., Baker, A. L., NYU PROSPEC Collaboration & Phelps, E. A. The functional roles of the amygdala and prefrontal cortex in processing uncertainty. *J. Cogn. Neurosci.* **31**, 1742–1754 (2019).
60. Miu, A. C., Heilman, R. M. & Houser, D. Anxiety impairs decision-making: psychophysiological evidence from an Iowa gambling task. *Biol. Psychol.* **77**, 353–358 (2008).
61. Nassar, M. R. et al. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* **15**, 1040–1046 (2012).
62. Bacow, T. L., May, J. E., Brody, L. R. & Pincus, D. B. Are there specific metacognitive processes associated with anxiety disorders in youth? *Psychol. Res. Behav. Manage.* **3**, 81–90 (2010).
63. Calvo, M. G. & Miguel-Tobal, J. J. The anxiety response: concordance among components. *Motiv. Emot.* **22**, 211–230 (1998).
64. Hoehn-Saric, R., McLeod, D. R. & Zimmerli, W. D. Symptoms and treatment responses of generalized anxiety disorder patients with high versus low levels of cardiovascular complaints. *Am. J. Psychiatry* **146**, 854–859 (1989).
65. Lenow, J. K., Constantino, S. M., Daw, N. D. & Phelps, E. A. Chronic and acute stress promote overexploitation in serial decision making. *J. Neurosci.* **37**, 5681–5689 (2017).
66. Nestadt, G. et al. Obsessive–compulsive disorder: subclassification based on co-morbidity. *Psychol. Med.* **39**, 1491–1501 (2009).

67. Dorfman, H. M., Bhui, R., Hughes, B. L. & Gershman, S. J. Causal inference about good and bad outcomes. *Psychol. Sci.* **30**, 516–525 (2019).
68. Speekenbrink, M. & Konstantinidis, E. Uncertainty and exploration in a restless bandit problem. *Top. Cogn. Sci.* **7**, 351–367 (2015).
69. Knowles, K. A. & Olatunji, B. O. Specificity of trait anxiety in anxiety and depression: meta-analysis of the state-trait anxiety inventory. *Clin. Psychol. Rev.* **82**, 101928 (2020).
70. Roberts, K. E., Hart, T. A. & Eastwood, J. D. Factor structure and validity of the State-Trait Inventory for Cognitive and Somatic Anxiety. *Psychol. Assess.* **28**, 134–146 (2016).
71. Horn, J. L. A rationale and test for the number of factors in factor analysis. *Psychometrika* **30**, 179–185 (1965).
72. Bartlett, M. S. The statistical conception of mental factors. *Br. J. Psychol.* **28**, 97–104 (1937).
73. Revelle, W. psych: Procedures for Psychological, Psychometric, and Personality Research. R package version 2.2.3 (Northwestern University, 2022); https://CRAN.R-project.org/package=psych

Acknowledgements

We thank T. Rusch for help with data analysis, and members of the Phelps Lab and Gershman Lab for helpful discussions. This work was supported by the National Institute on Drug Abuse (Grant No. R01 DA042855 to E.A.P.). The funders had no role in the study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

H.F., S.J.G. and E.A.P. developed the study concept and designed the study. H.F. collected data and performed data analysis. H.F. interpreted the data under the supervision of S.J.G. and E.A.P. All authors wrote the manuscript and approved its final version for submission.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41562-022-01455-y.

Correspondence and requests for materials should be addressed to Elizabeth A. Phelps.

Peer review information *Nature Human Behaviour* thanks Christoph Korn and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

nature portfolio

Corresponding author(s): Elizabeth Phelps

Last updated by author(s): Aug 18, 2022

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. F, t, r) with confidence intervals, effect sizes, degrees of freedom and P value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☐	☒	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's d, Pearson's r), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Participants were recruited via Prolific platform. They completed an interactive game, which was programmed using jspsych library (version 6.0.5) in javascript and was hosted on cognition.run. They also filled out self-reported psychiatric questionnaires via web forms.
Data analysis	All data analysis code are available at https://osf.io/y6urc/ Analysis were completed in R (version 4.0.2) and Matlab (version R2020a)

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

All data for two experiments are available at https://osf.io/y6urc/

Human research participants

Policy information about studies involving human research participants and Sex and Gender in Research.

Reporting on sex and gender	The study findings apply to all genders. Our analysis included gender and age as demographic covariates, but no gender-based analysis was performed separately because the study did not have a priori hypothesis related to gender difference. Gender was also not considered in the experiment design. The information on gender was collected based on self-report measures and was included in the data.
Population characteristics	See above
Recruitment	Participants were recruited from Prolific platform. The only restrictions placed on the sample were nationality (based in the USA), previous experience on Prolific (have completed at least 10 but no more than 10000 tasks on the Prolific platform before) and an "approval rate" (indicating that the participant pays attention and follows instructions correctly in tasks) of over 95%. While all participants are self-selected due to interest and motivation to participate in research studies, this is unlikely to introduce bias into the sample since the participants recruited through online crowdsourcing platforms are more diverse in age, race, and socioeconomic status than typical undergraduate research samples (Buhrmester, Kwang, & Gosling, 2011; Perspectives on Psychological Science), and studies have shown that participants recruited from the Prolific Platform provides high-quality data (Peer et al., 2017; Journal of Experimental Social Psychology; Peer et al., 2021; Behavior Research Methods)
Ethics oversight	The study was approved by Harvard University Committee on the Use of Human Subjects (protocol number:IRB19-0789)

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Quantitative experimental studies with cognitive modeling component
Research sample	Two general population samples were recruited from Prolific platform (Experiment 1: N=531; Experiment 2: N=576). Participants were required to be based in the USA and to have completed at least 10 tasks but no more than 10000 tasks on the Prolific platform before. They were also required to have successfully completed most previous tasks on the Prolific platform (95% of previous tasks approved). After exclusion (for exclusion criteria, see the 'data exclusion' part below), five hundred and one participants were included in Experiment 1 (219 women, 277 men, 5 unreported; Age M = 36.1 years, SD = 10.9). four hundred and eighty-four participants were included in Experiment 2 (197 women, 279 men, 8 unreported; Age M = 35.3 years, SD = 10.5). Informed consent was given before testing. While not a fully representative sample, the participants recruited from online crowdsourcing platforms are more diverse in age, race and socioeconomic status than typical undergraduate research samples (Buhrmester, Kwang, & Gosling, 2011; Perspectives on Psychological Science; Peer et al., 2017; Journal of Experimental Social Psychology).
Sampling strategy	All experiments used random sampling. The sample size of experiment 1 was determined based on a power analysis. We conducted the power analysis by simulating the pilot data, which revealed that a sample size of 500 would be necessary to detect an effect with 80% power. In experiment 2, we aimed for a comparable sample size of experiment 1 given the similarity in the tasks used.
Data collection	The data were obtained via interactive games and questionnaires displayed in the participants' web browsers on their laptops. The participants were told that in the game, they will see different slot machines and choose between them. The survey will ask you questions about how you feel, think and behave both in the past and in the current situation. The participants were blind to the study hypothesis.
Timing	Experiment 1: Dec 2020 - Jan 2021 Experiment 2: Apr 2021 - May 2021
Data exclusions	Participants were excluded if they reported age outside the range of 18-65, chose the more rewarding option < 60% of trials in the two-armed bandit task, or did not complete anxiety-related questionnaires. The above exclusion criteria has been preregistered at https://aspredicted.org/vi8wg.pdf. Thirty participants (5.6%) were excluded in experiment 1 and seventy-two (12.9%) participants were excluded in experiment 2.
Non-participation	No participants dropped out.

Randomization	In all experiments no between-subjects experimental groups were created.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging