Exploring the Unexplored: Worry as a Catalyst for Exploratory Behavior in Anxiety and Depression

Kristin Witte[1,2,*], **Toby Wise**[3], **Quentin J. M. Huys**[4,5,+], **and Eric Schulz**[1,2,+]

[1]Max Planck Institute for Biological Cybernetics, MPRG Computational Principles of Intelligence, Tübingen, 72076, Germany
[2]Helmholtz Munich, Center for Computational Health, Neuherberg, 85764, Germany
[3]King's College London, Department of Neuroimaging, Institute of Psychiatry, Psychology & Neuroscience, London, WC2R 2LS, UK
[4]University College London, Division of Psychiatry, London, W1T 7NF UK
[5]University College London, Max Planck UCL Centre for Computational Psychiatry and Ageing Research and Wellcome Centre for Human Neuroimaging, Institute of Neurology, London, WC1B 5EH, UK
[*]kristin.witte@tuebingen.mpg.de
[+]these authors contributed equally to this work

ABSTRACT

The relationship between anxious and depressive traits and exploration behavior has been examined in several studies with mixed results. While some studies suggest that anxious and depressive traits are related to avoidance and a decrease in exploratory behaviour, others find the opposite to be true. In our studies, we adopt a multi-armed bandit task in which arms that were spatially close to each other have similar rewards, allowing for generalisation from observed rewards. Furthermore, we introduced risks to simulate costs of over-exploration in the real world. In two studies, we investigate the relationship between transdiagnostic symptoms of anxiety and depression, specifically worrying, and task behaviour. While our first study uses a purely correlational design, our second study involves a psychotherapy-inspired intervention to reduce worries and investigate their causal effect on exploration behaviour. The results suggest that worrying may be a causal factor linking anxious and depressive traits to increased exploration behaviour. Specifically, using computational modelling, we show that worrying is related to an increased preference for novel options, as opposed to mere choice stochasticity. These findings enhance our understanding of the complex links between depression, anxiety and exploration behaviour, and highlight the importance of worry in driving increased exploration.

Keywords: exploration-exploitation; anxiety; depression; worry; rumination

Introduction

Imagine you are planning a vacation: You could either go to your usual place where you know the local customs, the best restaurants, and the shortest way to the beach, or you could try out a new destination that you have never been to before, hoping that it might be even better. In this situation, you are faced with an exploration vs. exploitation dilemma: Will you stick with what you know, i.e. exploit, or will you try something new that might be better but might also be worse, i.e. explore)?

In the field of computational psychiatry, previous research applied computational modelling to investigate how people make such decisions and how their strategies relate to psychiatric traits. Traits of particular interest are those involved in depression and anxiety: Depression and anxiety are comorbid and share many symptoms[1]. Improving our understanding of these transdiagnostic traits can therefore improve our understanding of both disorders. One such transdiagnostic trait is worrying, which has been linked to both disorders[2], albeit with potentially different contents[3]. In the exploration-exploitation trade-off, worry content could be either the threat of losing rewards (leading to increased exploitation) or the threat of missing out on even better, currently unobserved rewards (leading to increased exploration). Disentangling these two potential worry contents may shed light on what leads to worry and thereby elucidate disease mechanisms in both disorders. Intuitively, one would expect anxious and depressed individuals to display less exploration behaviour as these disorders have been linked to increased avoidance of negative[4,5] or uncertain[6] outcomes as well as a general lack of motivation to seek rewards[7]. Some studies have found anxious traits to be related to decreased directed exploration[8] and both chronic and acute stress to increased exploitation[9]. However, other studies have found no relationship between the transdiagnostic traits of worrying and intolerance to uncertainty and exploration behaviour[10]. When researchers used tasks where participants could never be entirely certain about the values of all options, for example, by introducing a previously unobserved arm in a bandit task[11] or through volatile underlying rewards[12,13], they found that anxiety and depression were associated with increased exploration behaviour. This suggests that anxious and depressed individuals may have an increased need for information. It has also been proposed that different subtraits of anxiety are differentially related to exploration behaviour. Specifically, Fan et al.[14] found somatic anxiety to be related to decreased directed and undirected exploration but negative affect to be related to increased directed and undirected exploration.

One aspect that has rarely been considered in this line of research is the fact that real-world exploration comes with risks: In our initial example, even though a trip to Mars might be the most exciting and novel vacation, it also has a non-negligible probability of ending fatally. There is some evidence that people generally tend to seek out uncertainty in safe environments but avoid it when there are risks to their exploration[15].

Another aspect of our task that makes it more ecologically valid is that instead of having only a few options to choose from as in classical bandit tasks[16,17], our task offers 121 options with correlated values such that participants can generalise from observations. Thereby we expect our task to be more ecologically valid as its multitude of options and possibility for generalisation mirrors real-world decisions: When planning a vacation, there are usually an immeasurable amount of options. After a trip to a small town in France, one has not only learned about how much one likes vacationing in that particular French town but likely in the entire area and, to a lesser degree, in the entire French countryside (see also[18]).

We present two studies that relate transdiagnostic traits commonly implicated in anxiety and/or depression, such as cognitive and somatic anxiety, as well as rumination, to behaviour on a task where exploration can (sometimes) lead to the loss of all previously collected rewards, making it a risky undertaking. Specifically, we used a gamified version of a rich multi-armed bandit task[15,19] in which subjects searched for rewards on an 11 x 11 grid, with each square representing an arm in a multi-armed bandit. The rewards associated with nearby arms were similar, allowing for spatial generalisation and thereby making the task more ecologically valid (see Figure1A-B for an example). As the number of arms in this multi-armed bandit task far exceeds the number of trials allowed in a given block, subjects could never explore all the arms, leaving a degree of irreducible uncertainty that might encourage exploration[11]. We use computational modelling to distinguish between exploration that is due to pure choice stochasticity (i.e. random exploration) and exploration that is guided towards uncertain or novel options. To get reliable estimates of these two exploration strategies, we use modern hierarchical Bayesian parameter estimation techniques which have been shown to yield more robust parameter estimates[20–22].

Given the mixed findings in the literature regarding the link between anxious and depressive traits and exploration behaviour, we first took a cross-sectional approach. We contrasted exploration behaviour in a safe and a risky version of the task and related this to measures of traits implicated in anxiety and depression. Based on our initial findings, we conducted a second study to verify our preliminary findings that worrying is a transdiagnostic process that drives increased exploration. To achieve this, we used a psychotherapy-inspired intervention to experimentally reduce worries and assess their causal implications on behaviour in the risky version of our task.

Results

General approach

In the tasks, subjects viewed an 11 by 11 grid, with each square representing an arm in a multi-armed bandit task (see Figure 1A-B;[15,19]). At the beginning of each block, the reward underlying one of the squares was revealed. In Study 1 (Figure 1A), this square was randomly selected from the arms with rewards above the safety threshold (50). Rewards for all arms remained stable throughout a block and were spatially correlated. In Study 1, participants played 10 blocks, half of which were risky, meaning that if subjects clicked a square with a reward below the safety threshold during any of these blocks, they lost all rewards from that block and moved on to the next block. In Study 2 (Figure 1B), participants played 11 blocks, all of which were risky blocks. The first block was an additional practice block and was excluded from all analyses. In addition, we designed the reward structure so that there were only two areas (about 9-12 squares each) with rewards above the threshold, similar to two islands in a sea of unsafe squares. There was always a path of safe squares connecting the two areas. One area had twice as many rewards as the other, but the starting square was always in the lower reward area, and subjects were aware of this structure. Importantly, the exact location of the reward areas was still unknown to the subjects in all rounds and had to be discovered through exploration. In this study, halfway through the task (i.e. after 6 rounds), subjects in the intervention condition received a metacognitive therapy[23]-inspired intervention aimed at reducing worry. In contrast, subjects in the control condition received a control intervention.

Model-agnostic results

Behaviour in safety vs risk

Overall, subjects learned and performed well in both versions of the task (see Figure 1C-D). As a model-agnostic index of exploration behaviour we use the subjects' proportion of novel squares selected which was significantly reduced in the risky condition ($\beta = -1.00$, 95%CI: $-1.06, -0.95$, see Figure 2A), suggesting that, when in a risky environment, subjects explored

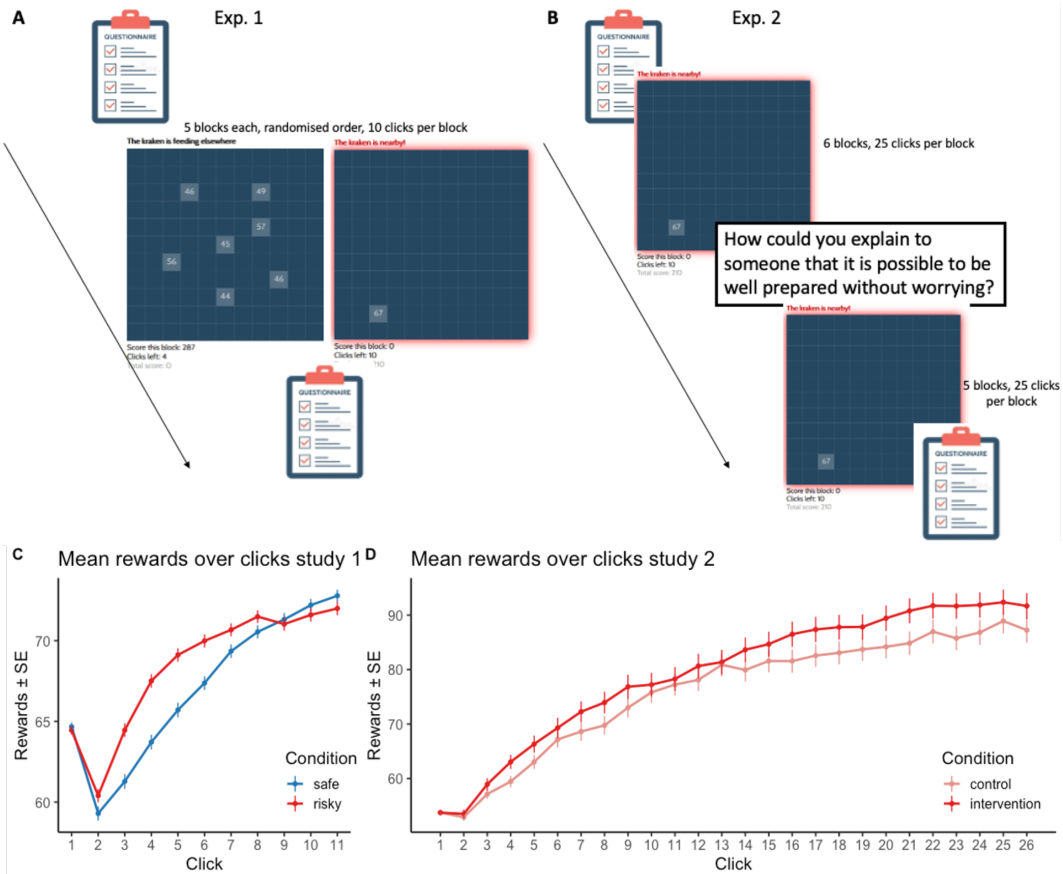

Figure 1. Overview of the study designs. **A**: In study 1, subjects completed a range of questionnaires, then performed 10 blocks of the task, five safe, five risky, interleaved. They then answered some more questions. **B**: Subjects again completed a range of questionnaires and then only performed the risky version of the task. After 6 blocks of the task, subjects underwent either the metacognitive therapy intervention or a control intervention and then proceeded to play 5 more blocks of the task and then answered some more questions. **C-D**: Subjects learned to find high rewards in all conditions and both tasks. In study 2, subjects were given more trials to encourage exploration.

less. Nevertheless, on average, participants clicked a square below the safety threshold on 46% of the risky rounds causing them to loose all rewards on that round.

Bayesian mixed-effects regressions revealed that traits related to worrying such as rumination and cognitive anxiety as well as negative affect and depressivity were related to increased probabilities of choosing a novel option (Figure 2B), particularly in the risky condition (Figure 2C). This thus suggests that, in the risky condition, traits related to cognitive worrying and rumination but not somatic anxiety were related to increased exploration behaviour.

Experimentally reducing worries reduces exploration behaviour

The metacognitive therapy-based intervention successfully reduced subjects' self-reported nervousness by almost 20% while there was no change in self-reported nervousness in the control condition (β = -18.88, 95%CI: -30.32, -7.24, see Figure 2D). This reduction in nervousness was not accompanied by a change in task performance, as indexed by a lack of a statistically significant effect on the average rewards collected per click (β = 0.16, 95%CI: -5.30, 5.61), or the proportion of subjects clicking a square with a reward below the safety threshold and thereby losing all rewards (β = 1.75, 95%CI: -0.60, 4.09). Overall, participants clicked a square with a reward below the safety threshold on approximately 57.3% of the rounds.

The intervention did, however, have the predicted effect on subjects' exploration behaviour: Subjects in the intervention condition decreased their exploration behaviour following the intervention (β = -0.98, 95%CI: -1.92, -0.03; see Figure 2F). When accounting for subjects' self-reported nervousness, it was nervousness beyond the intervention itself that predicted subjects' exploration behaviour (β = 0.09, 95%CI: 0.02, 0.16; see Figure 2E). It thus seems that the intervention reduced subjects' exploration behaviour by reducing their nervousness. A causal mediation analysis using non-parametric bootstrapping confirmed this relationship as the indirect effect of the intervention on exploration via nervousness reached significance (-0.11,

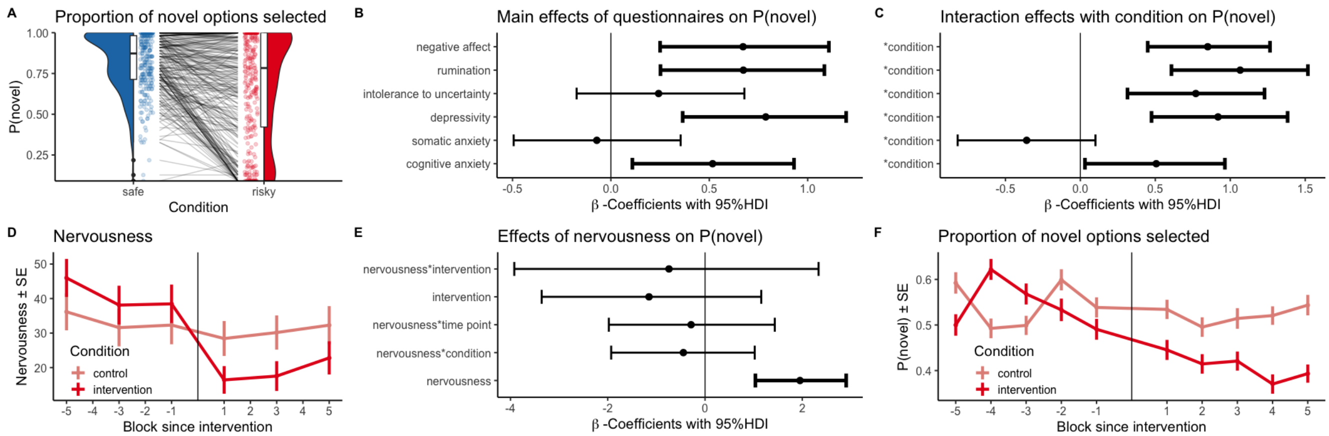

Figure 2. Model-agnostic results. **A**: Study 1, proportion of novel options out of all options selected in the safe vs. the risky condition. Dots and lines indicate single participants. **B-C**: Study 1, effects of questionnaire scores on probability of choosing a novel option on each trial. Regressions were performed separately for each questionnaire and accounted for reward on the previous square, trial number, and block number. For precise statistics, see SI. **D**: Study 2, nervousness throughout the task. Vertical line indicates when the intervention happened. **E**: Study 2, effects of nervousness on the probability of choosing a novel option on each trial. All predictors were estimated in the same regression. The predictor intervention encodes the interaction effect of time point (Pre vs Post-intervention) by condition (intervention vs control). **F**: Study 2, the proportion of novel options selected throughout the task, vertical line indicates when the intervention happened.

95%CI: -0.19, -0.03) while the direct effect did not (-0.11, 95%CI: -0.358, 0.14), suggesting that the effect of the intervention on exploration was fully mediated by nervousness.

To get an idea of how the intervention might have impacted participants' explicit strategies, we added an explicit question to the intervention part. We asked participants right before the intervention and after they had finished the game whether they think it would be a better strategy to try to click as many squares as possible or to re-click the best ones. A Wilcoxon signed rank test revealed a significant shift towards re-clicking squares that give the highest rewards following the intervention ($V = 71$, $p = .014$), a strategy that is likely more adaptive.

Computational modelling results

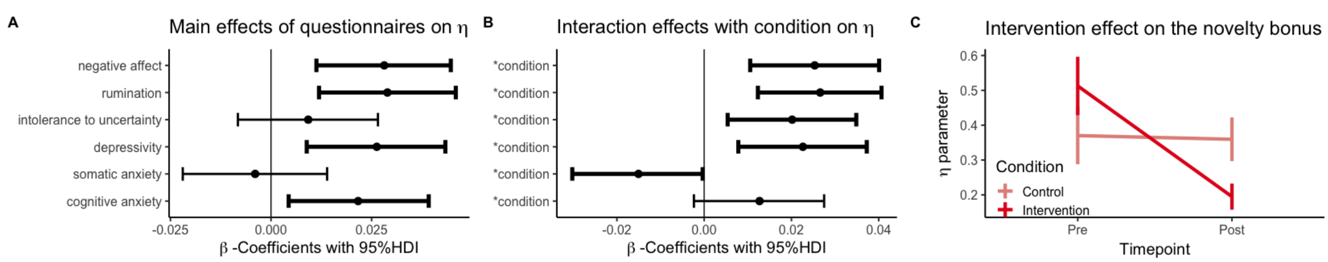

Figure 3. Computational modelling results. **A-B**: Study 1, effects of questionnaires on the novelty bonus parameter η. All questionnaire effects were estimated in separate regressions. For precise statistics, see SI. **C**: Study 2, effect of the intervention on the novelty bonus parameter η

The observed association between worry and increased exploration behaviour could theoretically be due to either an increase in choice stochasticity (i.e. reduced reward sensitivity) or an increased valuation of novel or uncertain options. We therefore used computational modelling to better understand participants' strategies in the task. The best fitting model, inspired by Dubois et al.[11], was a Gaussian Process learning model combined with a valuation model that added a value bonus to previously unobserved options. The Gaussian Process learning model uses a Radial Basis Function kernel to learn a distribution over possible generating functions of the data, thereby non-parametrically approximating the reward estimate and uncertainty for each arm at each time point and for each subject. It can generalise smoothly from observations, with a free parameter λ controlling the degree of generalisation from observations. In the valuation model, a free parameter η, called the novelty bonus, scaled a dummy-coded variable indicating whether an option had been previously observed (0) or was yet to be explored (1).

Finally, we estimated choice probabilities using a Softmax choice policy with a free parameter τ governing choice stochasticity, often interpreted in this context as random exploration. For a full model comparison, see SI.

In Study 1, linking the transdiagnostic traits to the model parameters revealed that the increased exploration behaviour in worry-prone individuals is most likely driven by an increased valuation of novel options, as opposed to an increase in choice stochasticity or generalisation. Specifically, worry-related traits such as negative affect, depressivity and rumination were related to an overall increase in the novelty bonus parameter η, particularly when risks were involved. Somatic anxiety, in turn, was related to a decreased η parameter in the risky condition (see Figure 3A-B). We found no relationships between the questionnaire traits and generalisation and only small and inconsistent relationships with choice stochasticity (see SI).

We found a similar relationship in Study 2 where the novelty bonus parameter η but not the generalisation parameter λ or the Softmax temperature τ were significantly predicted by an interaction between the timepoint (Baseline vs After intervention) and the condition (intervention vs. control; β = -0.33, 95%CI: -0.53, -0.14), suggesting that the intervention decreased participants' seeking for novel options (see Figure 3C). It should however be noted that while in study 1 the models we tested were perfectly identifiable, this was not the case in study 2. Specifically, at baseline, the winning model was not perfectly distinguishable from an alternative model, consisting of a GP learning model combined with a Softmax, thus not incorporating the novelty bonus. Further, while the winning model showed perfect parameter recoverablility and identifiablity in study 1, this was not perfectly given in study 2, making η and τ somewhat confounded. Thus, for study 2, we cannot say with certainty whether the intervention decreased subjects' seeking of novel options (η) or their choice stochasticity (τ). For details, see SI.

Discussion

We investigated the link between transdiagnostic traits implicated in depression and anxiety and exploration behaviour in risky and safe environments. We further investigated the effect of a reduction in worrying on exploration behaviour in risky environments. Across both studies, we found the transdiagnostic trait of worrying to be related to increased exploration, particularly when there were risks associated with such behaviour. Computational modelling revealed that this increase in exploration might be due to an increased valuation of novel options, as opposed to a mere increase in choice stochasticity or generalisation.

Our findings might seem counter-intuitive at first, as anxiety is usually associated with avoidance[4–6] and a general lack of motivation or drive to seek rewards[7]. Further, some research has found anxiety to be related to decreased directed exploration[8] and stress to over-exploitation[9]. However, as mentioned earlier, some studies did find anxious and depressive traits to be related to increased exploration, particularly using paradigms with a certain amount of irreducible uncertainty through unobserved options or reward volatility[11–13]. In our task, there were 121 arms and only 11 or 25 trials. Subjects could thus never observe all of the arms, leaving a large amount of residual uncertainty about whether they have found the best arm. This uncertainty might be driving exploration in our task as subjects with a stronger tendency to worry might need more information to commit to an option to exploit. This idea also finds support in the information-seeking literature where anxious individuals tended to seek out more information[24], particularly in response to large environmental changes[25], thus matching the findings from reinforcement learning tasks with volatile rewards.

Interestingly, information to be gained was implemented in a very coarse, heuristic way by our winning model, namely as whether a square had been previously observed or not. This model outperformed a model that had access to a more fine-grained estimate of the information to be gained by selecting a given option by using the posterior variance from the GP learning model as a measure of uncertainty about the reward of that option (see[15,19]). It thus seems that participants in our study favored heuristic and less cognitively demanding exploration strategies, a tendency that has been observed to intensify under heightened cognitive loads[26], suggesting that it might be more resource-efficient. This inclination towards efficient strategies is consistent with previous evidence, indicating a prevalent preference for random exploration as opposed to directed exploration[11]. While in our study, participants did not engage in purely random exploration, their exploration was not driven by uncertainty either, making it a middle-ground between random and directed exploration. In a way, participants explored randomly within the constraint of previously unobserved options. Our results thereby possibly reconcile the literature that linked anxious and depressive traits either to an increase in uncertainty-driven exploration[11,14] or random exploration[12–14].

Still, in study 2, novelty seeking could not be perfectly disentangled from choice stochasticity, due to the high task difficulty, making participants frequently select a square below the safety threshold early in a round and thereby leaving a limited amount of data per participant. Further research is thus needed to confirm whether our findings are best explained by an intolerance to uncertainty or an increase in choice stochasticity, often also termed decreased reward sensitivity. For this, it would be important to contrast exploration-exploitation tasks and information-seeking tasks to disentangle reward-related aspects from a drive to seek out information. It would further be interesting to gauge how different aspects of information such as its hedonic value (i.e. how it makes people feel) and its instrumental value (i.e. how useful it is)[27] impact these differences in information seeking.

The fact that worrying was related to novelty seeking and generally increased, not decreased, exploration also suggests that, in our tasks, participants were more worried about not finding the right square than about losing all rewards on that round or selecting a bad square with low rewards. In other words, the behaviour of participants high on worrying could be best described by perfectionism, rather than fear of threats. This is somewhat in line with the literature on worry contents in depression and anxiety[3]: Depressed individuals are described to worry about having a lack of confidence. Anxious individuals in turn were more likely to worry about a loss of control. Our task might relate to both of those worry contents: Not knowing what rewards are hidden under the unopened squares and whether a better reward is out there might induce a lack of confidence. In a similar vein, not having all the information before committing to a square to re-select until the end of the round might feel like a loss of control. Still, more research is needed to clarify how the behaviour we observe on this task relates to different worry contents.

While we do mostly talk about worries when interpreting our results, our findings might not only apply to worrying but also to repetitive negative thoughts more broadly. Repetitive negative thinking includes worrying but also other cognitive processes such as reflection, rumination, or problem-solving and has been suggested to underlie a range of different psychiatric disorders through potentially the same mechanism[28,29]. Both reflection and rumination are somewhat present within the questionnaires we used in study 1 through questionnaire items such as: "I picture some future misfortune". Problem solving however is not an aspect of the psychiatric questionnaires we used. It thus seems like our findings might generalise to some aspects of repetitive negative thinking but less to others. There is a need for future research to directly investigate the link between different facets of repetitive negative thinking and exploration behaviour.

Conclusions

Across two studies, we investigated the link between transdiagnostic traits implicated in rumination and worrying and an increase in exploration behaviour, particularly in environments, where there are risks involved in exploration. We showed that subjects with a higher tendency to worry also showed higher levels of exploration, which seemed to be guided by an increased drive to select unobserved options. Our findings open interesting novel research avenues on the relationship between worrying traits and exploration as well as the potential transdiagnostic nature of worries.

Methods

Participants

For study 1 we gathered data from 300 participants using the platform prolific (prolific.com). Inclusion criteria were residency in the United States and an approval rate of minimum 95% on the platform. We subsequently excluded subjects that failed more than 2 of the 4 attention checks embedded in the questionnaire section or that took more than 11 attempts (mean + 2SD) to complete the comprehension questions. The final sample consisted of 282 subjects (200 females, mean age: 27.78, SD: 9.92).

For study 2, we gathered data from 85 participants using the platform prolific (prolific.com). Inclusion criteria were being a resident of the United Kingdom, between 18 and 65 years old, a native speaker of English and having an approval rate of at least 90% on the platform. We subsequently excluded subjects that failed more than 1 of the 4 attention checks embedded in the questionnaire section or that never clicked a square more than once despite having found the best square in a block (and being made aware of the maximum reward magnitude during the instructions; N = 5) as these subjects likely did not understand the task. The final sample consisted of 75 subjects (40 females, mean age: 31.68, SD: 11.29).

Task

The task was a multi-armed bandit with the arms arranged as squares on an 11-by-11 grid[15,19]. During each block, the mean rewards of each arm remained stable, but a small amount of noise was added. The rewards of neighboring arms were similar to allow for generalization from nearby observations. In study 1, rewards were between 10 and 90, with half of the squares having a reward of 50 or less and thus falling below the safety threshold. In study 2, grids were designed to discourage exploration, with only two regions having rewards above the safety threshold (45 in study 2) that were connected by a narrow path. One of these areas always had rewards around 60 and the other around 120. The location of these areas varied across blocks. In both studies, one square with a value above the safety threshold was revealed at the start of each block. In study 2, this square was always in the less rewarding region such as to strengthen the explore-exploit trade-off. Subjects were aware of all these specifications and needed to correctly answer a range of comprehension questions to proceed to the task. They were allowed 10 clicks (24 in study 2) per block to either receive the rewards from an already revealed square or click a novel square to uncover it and receive the rewards from that square.

Questionnaires

In both studies, we administered questionnaires designed to cover a range of traits involved in depression and anxiety. In Study 1, the focus was on transdiagnostic traits involved in both conditions. We administered questionnaires assessing trait levels of cognitive and somatic anxiety (State Trait Inventory for Cognitive and Somatic Anxiety[30]), intolerance of uncertainty (Intolerance to Uncertainty Scale, short form[31]), rumination (Reflection Rumination Questionnaire[32]), depressivity (Community Assessment of Psychic Experience, depressivity subscale[33]), and negative affect (Personality Inventory for DSM-5, negative affect subscale[34]). We also administered a battery of questionnaires covering a wide range of psychiatric conditions ([35]; see SI).

In Study 2, we administered a range of questionnaires covering anxiety, depression and metacognitive beliefs as well as meta-worries. As these questionnaire scores had no bearing on task behaviour, we will not discuss them further (see SI). Also, after 3 of the task blocks before the intervention and 3 of the task blocks after the intervention, we asked subjects to indicate what was the most nervous they had felt during the previous block of the task using a slider ranging from not nervous at all (0) to very nervous (100).

Intervention

The intervention used in Study 2 had two components. First, there was a psychoeducational component introducing the main ideas in Metacognitive therapy. Second, there was an interactive component motivated by active learning principles[36] which involved asking participants to reformulate the psychoeducational content in a manner such that it could be effectively communicated to participants with similar characteristics as themselves or to give intuitive examples about it from their own lives (see SI for the full intervention script).

To avoid participants understanding the true aim of the study and therefore adjusting their behaviour accordingly, we told them that the aim of the experiment was to get their help in devising an intervention that makes other people less anxious. They would therefore complete the task a first time, to understand why it could make other people nervous, and then a second time after having given their input on the intervention, to give them an opportunity to reflect some more on the way they devised the intervention after which they were allowed to change their answers. Subjects in the control condition were instead asked to provide feedback on the task (see SI for the full control intervention script).

Statistical tests

For the model agnostic analyses in both studies, we performed a series of Bayesian Linear Mixed Effects Regressions implemented in the `brms` package[37] in R[38]. To ensure convergence and improve the interpretability of the resulting betas, we used the z-scores for all numeric predictors and used effect coding for all binary predictors. We further mean-centered all numeric dependent variables. When analysing the relationships of questionnaire scores with the dependent variables of interest, we performed separate regressions for each questionnaire subscale, to avoid variance inflation.

Model fitting

We fit 6 models to the data in study 1 and 5 of them also to the data in study 2. All models had been used and validated in prior studies[11,15,19] For a full overview, see the SI. All computational models approximated participants' learning of reward estimates in different locations using a Gaussian Process (GP) with a Radial Basis function kernel that included a free parameter λ encoding generalisation[15,19] (see SI for details). The computational models only differed in how the value and uncertainty estimates from the GP were translated into option values. According to a 5-fold cross-validation, the best fitting model only used the mean of the posterior estimate of the GP for each option to determine its value, disregarding uncertainty about these estimates. This mean posterior estimate was increased by a so-called novelty of an option, which was 1 for unobserved options and 0 otherwise. This novelty was scaled by the novelty bonus parameter η, similarly to a recent study by Dubois et al.[11]. To translate the resulting option values into choice probabilities ($p_t(x)$), we used a classic Softmax function with a free temperature parameter τ governing subject-level choice stochasticity. We also fit a random model, which assigns the same probability to each option on each trial (and therefore has the same log likelihood for each subject).

We estimated all computational models using a hierarchical Bayesian procedure, with the models specified using a non-centred parameterisation estimating a group-level parameter and then subject-level offsets where the subject-level parameter is determined by the group-level mean for that parameter group and its standard deviation multiplied by a subject-level offset parameter. For sampling the posterior distribution, we used No U-Turn Markov Chain Monte Carlo sampling implemented in NumPyro[39] in Python[40] (version: 3.8.3). We fit the data separately for each condition and ran two separate Markov chains per model and condition to ensure robustness, with 2000 samples per chain (in addition to 1000 warm-up samples).

References

1. Kaufman, J. & Charney, D. Comorbidity of mood and anxiety disorders. *Depress. Anxiety* **12 Suppl 1**, 69–76, DOI: 10.1002/1520-6394(2000)12:1+<69::AID-DA9>3.0.CO;2-K (2000).
2. Starcevic, V. Pathological worry in major depression: A preliminary report. *Behav. Res. Ther.* **33**, 55–56 (1995).
3. Diefenbach, G. J. *et al.* Anxiety, depression, and the content of worries. *Depress. Anxiety* **14**, 247–250 (2001).

4. 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, DOI: 10.1016/j.biopsych.2017.01.017 (2017).
5. Zorowitz, S., Momennejad, I. & Daw, N. D. Anxiety, avoidance, and sequential evaluation. *Comput. psychiatry (Cambridge, Mass.)* **4**, DOI: 10.1162/cpsy_a_00026 (2020).
6. Grupe, D. W. & Nitschke, J. B. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nat. Rev. Neurosci.* **14**, 488–501, DOI: 10.1038/nrn3524 (2013).
7. American Psychiatric Association, D., Association, A. P. *et al.* *Diagnostic and statistical manual of mental disorders: DSM-5*, vol. 5 (American psychiatric association Washington, DC, 2013).
8. Smith, R. *et al.* Lower levels of directed exploration and reflective thinking are associated with greater anxiety and depression. *Front. psychiatry* **12**, 782136, DOI: 10.3389/fpsyt.2021.782136 (2021).
9. Lenow, J. K., Constantino, S. M., Daw, N. D. & Phelps, E. A. Chronic and acute stress promote overexploitation in serial decision making. *The J. Neurosci.* **37**, 5681–5689, DOI: 10.1523/{JNEUROSCI}.3618-16.2017 (2017).
10. Zaller, I., Zorowitz, S. & Niv, Y. Information seeking on the horizons task does not predict anxious symptomatology. *Biol. Psychiatry* **89**, S217–S218 (2021).
11. Dubois, M. & Hauser, T. U. Value-free random exploration is linked to impulsivity. *Nat. Commun.* **13**, 4542, DOI: 10.1038/s41467-022-31918-9 (2022).
12. 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, DOI: 10.1038/s41380-021-01363-z (2022).
13. Blanco, N. J., Otto, A. R., Maddox, W. T., Beevers, C. G. & Love, B. C. The influence of depression symptoms on exploratory decision-making. *Cognition* **129**, 563–568, DOI: 10.1016/j.cognition.2013.08.018 (2013).
14. Fan, H., Gershman, S. J. & Phelps, E. A. Trait somatic anxiety is associated with reduced directed exploration and underestimation of uncertainty. *Nat. Hum. Behav.* **7**, 102–113, DOI: 10.1038/s41562-022-01455-y (2023).
15. Schulz, E., Wu, C. M., Huys, Q. J. M., Krause, A. & Speekenbrink, M. Generalization and search in risky environments. *Cogn. science* **42**, 2592–2620, DOI: 10.1111/cogs.12695 (2018).
16. Gershman, S. J. Deconstructing the human algorithms for exploration. *Cognition* **173**, 34–42, DOI: 10.1016/j.cognition.2017.12.014 (2018).
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, DOI: 10.1037/a0038199 (2014).
18. Wise, T., Emery, K. & Radulescu, A. Naturalistic reinforcement learning. *Trends Cogn. Sci.* (2023).
19. Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D. & Meder, B. Generalization guides human exploration in vast decision spaces. *Nat. Hum. Behav.* **2**, 915–924, DOI: 10.1038/s41562-018-0467-4 (2018).
20. Atkinson, G. & Nevill, A. M. Statistical methods for assessing measurement error (reliability) in variables relevant to sports medicine. *Sports medicine* **26**, 217–238 (1998).
21. Rouder, J. N. & Lu, J. An introduction to bayesian hierarchical models with an application in the theory of signal detection. *Psychon. bulletin & review* **12**, 573–604 (2005).
22. Scheibehenne, B. & Studer, B. A hierarchical bayesian model of the influence of run length on sequential predictions. *Psychon. bulletin & review* **21**, 211–217 (2014).
23. Wells, A. *Metacognitive therapy for anxiety and depression* (Guilford press, 2011).
24. 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, DOI: 10.1037/xge0000966 (2021).
25. Charpentier, C. J. *et al.* Anxiety increases information-seeking in response to large changes. *Sci. Reports* **12**, 7385, DOI: 10.1038/s41598-022-10813-9 (2022).
26. Cogliati Dezza, I., Cleeremans, A. & Alexander, W. Should we control? the interplay between cognitive control and information integration in the resolution of the exploration-exploitation dilemma. *J. Exp. Psychol. Gen.* **148**, 977–993, DOI: 10.1037/xge0000546 (2019).
27. Sharot, T. & Sunstein, C. R. How people decide what they want to know. *Nat. Hum. Behav.* **4**, 14–19, DOI: 10.1038/s41562-019-0793-1 (2020).

28. Watkins, E. R. Constructive and unconstructive repetitive thought. *Psychol. bulletin* **134**, 163 (2008).

29. Ehring, T. & Watkins, E. R. Repetitive negative thinking as a transdiagnostic process. *Int. journal cognitive therapy* **1**, 192–205 (2008).

30. 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).

31. Carleton, R. N., Norton, M. A. P. J. & Asmundson, G. J. G. Fearing the unknown: a short version of the intolerance of uncertainty scale. *J. Anxiety Disord.* **21**, 105–117, DOI: 10.1016/j.janxdis.2006.03.014 (2007).

32. Trapnell, P. D. & Campbell, J. D. Private self-consciousness and the five-factor model of personality: distinguishing rumination from reflection. *J. Pers. Soc. Psychol.* **76**, 284–304, DOI: 10.1037//0022-3514.76.2.284 (1999).

33. Mossaheb, N. *et al.* The community assessment of psychic experience (CAPE) questionnaire as a screening-instrument in the detection of individuals at ultra-high risk for psychosis. *Schizophr. Res.* **141**, 210–214, DOI: 10.1016/j.schres.2012.08.008 (2012).

34. Fossati, A., Krueger, R. F., Markon, K. E., Borroni, S. & Maffei, C. Reliability and validity of the personality inventory for DSM-5 (PID-5): predicting DSM-IV personality disorders and psychopathy in community-dwelling italian adults. *Assessment* **20**, 689–708, DOI: 10.1177/1073191113504984 (2013).

35. Wise, T., Robinson, O. J. & Gillan, C. M. Identifying transdiagnostic mechanisms in mental health using computational factor modeling. *Biol. Psychiatry* **93**, 690–703 (2023).

36. Yeager, D. S. *et al.* A synergistic mindsets intervention protects adolescents from stress. *Nature* **607**, 512–520, DOI: 10.1038/s41586-022-04907-7 (2022).

37. Bürkner, P.-C. Brms: An R package for bayesian multilevel models using stan. *J. Stat. Softw.* **80**, DOI: 10.18637/jss.v080.i01 (2017).

38. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria (2020).

39. Phan, D., Pradhan, N. & Jankowiak, M. Composable effects for flexible and accelerated probabilistic programming in NumPyro. *arXiv* DOI: 10.48550/arxiv.1912.11554 (2019).

40. Van Rossum, G. & Jr, F. D. Python reference manual. centrum voor wiskunde en informatica. amsterdam, the netherlands.
.

Acknowledgements (not compulsory)

This work was supported by the Max Planck Society and a Jacobs Research Fellowship and an ERC Starting Grant to ES and by the UCLH NIHR BRC. QJMH acknowledges grant funding from Wellcome Trust (221826/Z/20/Z), Carigest S.A. and Koa Health. QJMH has obtained fees and options for consultancies for Aya Technologies and Alto Neuroscience. TW is supported by a Career Development Award from the Wellcome Trust (225945/Z/22/Z).

Author contributions statement

All authors conceived the experiments, K.W. conducted the experiments, K.W. and T.W. analysed the results. All authors reviewed the manuscript.