The two faces of anxiety in exploration: Taking risks or playing it safe

Kristin Witte
Max Planck Institute for Biological Cybernetics
72076 Tübingen, Germany
kristin.witte@tue.mpg.de

Toby Wise
King's College London
London, England SE5 8AF
toby.wise@kcl.ac.uk

Eric Schulz
Max Planck Institute for Biological Cybernetics
72076 Tübingen, Germany
eric.schulz@tue.mpg.de

Abstract

While in most lab studies of exploration behaviour, the only potential downside to exploration is forgoing rewards, in real life, there can be actual risks to exploration, making *Safe Exploration* strategies important to avoid catastrophic outcomes. In the present study, we contrast exploration behaviour in safe with exploration behaviour in risky environments, i.e. in environments where reckless exploration can lead to the loss of all previously acquired rewards. Using computational modelling, we show that while people tend to use the same general strategy in both environments, they become more averse to uncertainty in risky environments compared to safe ones. We further investigate individual differences in exploration behaviour and show that while most types of anxiety and depression-related traits were associated with increased subjective uncertainty, there seemed to be two distinct strategies in dealing with this uncertainty: While subjects with high trait somatic anxiety decreased their exploration in risky environments, subjects high on depression and worry-related traits showed increased exploration and decreased aversion to uncertainty in risky environments. Our results reveal different aspects of exploration behaviour in a more ecologically valid task and show how this relates to transdiagnostic psychiatric traits, thereby illuminating potential disease mechanisms.

Keywords: exploration-exploitation, safe exploration, generalisation, uncertainty, cognitive anxiety, somatic anxiety, worry, depression

1 Background

Consider the following scenario: You just moved to a new country and want to try the local cuisine by going to a food market. Classic exploration strategies would tell you to try dishes that you have never heard of and do not know the ingredients of, as these have the highest potential for uncertainty reduction. Yet, intuitively, you might want to make a different call and rather choose a *safe exploration* strategy to avoid ending up with catastrophic outcomes such as a dish you completely dislike or food poisoning. This is because in real-world exploration scenarios there often are risks beyond potentially forgoing rewards that need to be taken into account, making safety a major concern in exploration vs. exploitation scenarios (Sui et al., 2015).

2 Present study

In the present study, we contrasted exploration behaviour in safe and in risky environments. More precisely, we used a spatially correlated multi-armed bandit task (Wu et al., 2018) where participants selected squares on an 11 by 11 grid to reveal the rewards underlying each square, trying to find as many rewards as possible in 10 clicks. Reward magnitudes were spatially correlated, thus allowing for generalisation from observed rewards. Participants played a total of 10 rounds, each of which constituted of a new grid to explore. On half of the rounds, participants were informed that it was a risky round: if they were to click a square with a reward below 50 (median reward in each round), they would lose all rewards from that round and move on to the next round. Half-way through the game, participants played an additional bonus round that involved freely exploring in a safe environment for 5 clicks. After this, 5 randomly chosen and previously unseen squares were highlighted one after the other and participants were asked to provide an estimate of the reward magnitude at that location, alongside indicating how certain they were of their estimate. Afterwards, they were asked to choose one of the 5 locations to sample before continuing with the task normally.

The aims of this study were two-fold: First of all, we wanted to test whether people use different exploration strategies in risky versus safe environments, similarly to a previous study by Schulz et al. (2018). Secondly, we investigated whether exploration behaviour in risky and safe environments relates to self-reported psychiatric traits such as anxiety and depressivity.

3 Results

To compare exploration behaviour in risky and safe environments, we contrasted the distances between subsequent clicks and the probability of choosing a unique option, i.e. clicking a square only once in an entire round, between the two conditions. Bayesian mixed effects regressions accounting for age and gender showed that both the average distances between subsequent clicks and the probability of choosing a unique option were decreased in the risky condition ($\beta = -0.97$, 95%CI: $-1.13, -0.65$; $\beta = -3.63$, 95%CI: $-4.41, -2.78$, respectively, see Figure 1A-B), suggesting that, when in a risky environment, subjects were exploring less and sticking to options closer to the previously selected one. To reveal the decision strategies being used in the different task conditions, we fit five different computational models to the data from both conditions separately. All models contained Gaussian Process (GP) learning model, with a Radial Basis Function Kernel including a λ parameter governing generalisation. The GP learns a distribution over potential generating functions based on past observations and returns a mean estimated reward and uncertainty about this estimate for each location on the grid. The GP learning model was combined with one of two types of decision policies, based on the winning models in Schulz et al. (2018): In the Confidence Bound models (CB), the choice probability of an option is the function of the estimated reward and the uncertainty about said reward. Whether uncertainty increases or decreases the probability of choosing an option is governed by the β parameter of the decision model. For positive β, uncertainty increases the probability of choosing an option, whereas for negative β it decreases it. In addition to the full version of the model (CB), we fit two reduced versions: in one version, β was fixed to 0, thus assuming indifference to uncertainty (CB_β0). In the other version, the λ parameter from the GP learning model was fixed to 0, thus assuming no generalisation from observed rewards (CB_λ0) and purely tracking the mean of individual options. We further fit two versions of the Probability Of staying Safe model: In the standard model (POS), subjects choose options with a high probability of their rewards being above a safety threshold, i.e. higher than 50 in our task. We also fit a version where the safety threshold is a free parameter, thus allowing for individual variation in how close subjects were willing to get to the actual safety threshold determined by the task (POS-T model). To assess comparative model fit we used a full cross-validation, thus for each condition fitting the model repeatedly to four of the five blocks and evaluating the accuracy on the held out block. This revealed that the winning model for both task conditions was the full version of the CB model (see Figure 1F-G), thus indicating that in both environments subjects' exploration was guided by uncertainty. While this suggests that subjects used the same general strategy in both conditions, one should note, that while in the safe condition, CB_β0 is the second-best fitting model, in the risky condition both POS models outperform the reduced versions of CB. This suggests that uncertainty has a stronger role in driving exploration in the risky condition than in the safe condition.

This idea also finds support when comparing the subject-level β parameters in the safe and the risky condition (Figure 1C): A linear mixed effects regression accounting for age and gender revealed that β values in the risky condition were significantly more negative than in the safe condition, thus suggesting increased aversion to uncertainty when there were risks involved ($\beta = -0.21$ 95%CI: $-0.23, -0.19$). With regard to generalisation, λ was slightly but significantly increased in the risky condition ($\beta = 0.16$ 95%CI: $0.06, 0.27$), suggesting increased generalisation from observed rewards. The τ parameter from the GP learning model governing the randomness of exploration is slightly but significantly elevated in the risky condition, when accounting for age and gender ($\beta = 0.001$ 95%CI: $0.00001, 0.002$), indicating a slight increase in random exploration in risky environments (see Figure 1E).

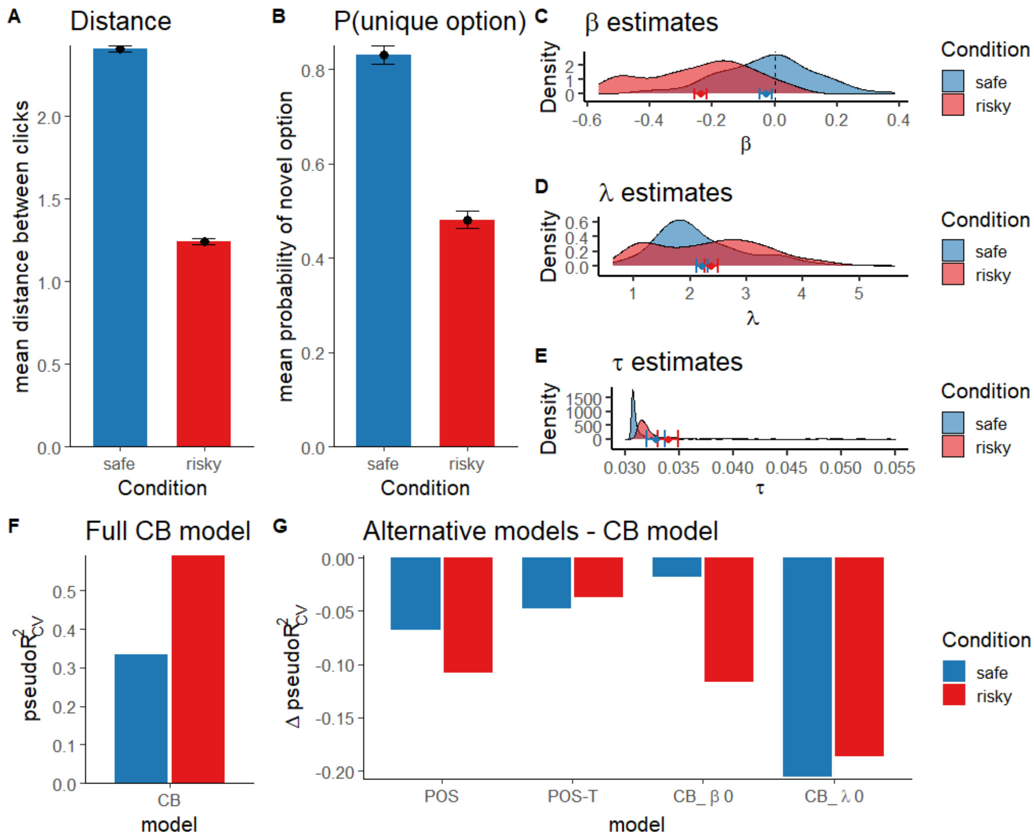

Figure 1: **A:** Average euclidean distances between subsequent clicks (± SE) in the safe vs the risky condition. **B:** Average probability of choosing a unique option (± SE), i.e. the probability of choosing a square only once over the course of an entire block. **C-E:** Distribution of subject-level parameter estimates from the CB model including means and 95%CIs for both conditions. **F:** Comparative model fit in the Safe condition. While CB is the best-fitting model, the reduced version with $\beta = 0$ is a close second, outperforming both POS models. **G:** Comparative model fit in the risky condition. CB is again the best-fitting model. In contrast to the safe condition, both POS models outperform the reduced versions of the CB model.

For our second aim, we administered questionnaires investigating trait levels of cognitive and somatic anxiety (State Trait Inventory for Cognitive and Somatic Anxiety, Ree et al., 2008), intolerance to uncertainty (Intolerance to Uncertainty Scale, short form, Carleton et al., 2007), rumination (Reflection Rumination Questionnaire, Trapnell and Campbell, 1999), depressivity (Community Assessment of Psychic Experience, depressivity subscale, Mossaheb et al., 2012) and negative affect (Personality Inventory for DSM-5, negative affect subscale, Fossati et al., 2013). The advantage of this selection of questionnaires is that they do not simply investigate anxiety and depression as single, independent disorders but instead test for transdiagnostic traits that can be present in both disorders, thus taking into account the immense comorbidity between the two disorders (Gillan et al., 2016; Kaufman & Charney, 2000). In the bonus round, all

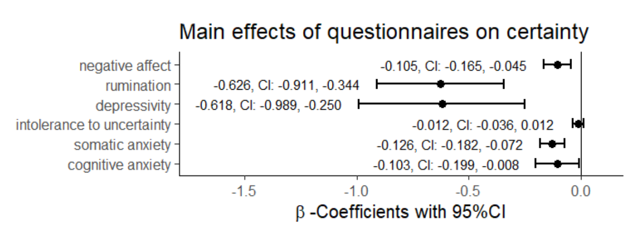

Figure 2: β-estimates (95%CI) of the different questionnaire scores in predicting subjective certainty estimates.

traits except for intolerance to uncertainty were related to decreased certainty estimates (see Figure 2), thus suggesting that individuals that had high tendencies to worry, feel down or to have an anxious feeling in their gut (somatic anxiety), also felt less certain about the rewards at unobserved locations.

How this uncertainty translated to exploration behaviour was, however, distinct between subjects with high somatic anxiety compared to the other, more cognitive, traits: Bayesian mixed effects regressions relating distances between clicks and the probability of choosing a unique option to the questionnaire scores revealed that all traits except for somatic anxiety were related to increased exploration, especially in the risky condition (see Figure 3 for effect sizes and credible sets). For somatic anxiety, the opposite was the case: distances between clicks as well as the probability of choosing a new option were significantly reduced in the risky condition (see Figure 3B and D). This suggests that subjects with a high tendency for cognitive worry and negative affect dealt with their uncertainty by exploring and thereby reducing it, while subjects with high somatic anxiety were more likely to explore less and thereby avoid uncertainty.

We confirmed this idea through computational modelling: all psychiatric trait scores except for somatic anxiety were related to increased β estimates in the risky condition, although for cognitive anxiety, this effect did not reach significance (see Figure 4A-B for effect sizes and 95%CIs). This indicates that subjects with a tendency to worry or feel depressed, also tend to be less averse to uncertainty when risks are involved, thereby explaining their increased exploration behaviour in the risky condition. When it comes to the generalisation parameter λ, only the depression related questionnaires on negative affect, rumination and depressivity were significantly related to increased parameter estimates in the risky condition (see Figure 4C-D). Lastly, the τ parameter from the GP learning model part of the CB model was slightly but significantly increased in subjects with higher tendencies for negative affect and intolerance to uncertainty, however this effect was reversed in the risky condition, indicating that the increased exploration in the risky condition and decreased uncertainty aversion are not due to random exploration (see Figure 4E-F). Taken together, these findings suggest that while a high tendency for somatic anxiety led individuals to decrease their exploration, the other, more worry related traits, led individuals to reduce their uncertainty by seeking out uncertain options and, in the cases of the more depression-related traits, even generalise more strongly from observed rewards.

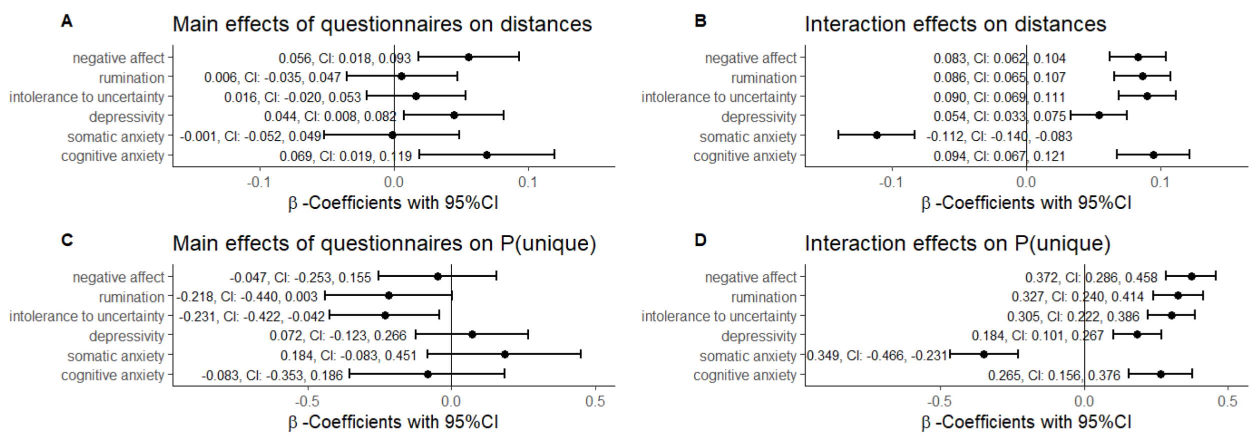

Figure 3: **A-B:** β coefficients (95%CI) of the different questionnaire scores in predicting distances between two subsequent clicks when accounting for the reward on the previous click. **A** shows the main effects, positive β indicate increased distances. **B** shows the interaction effects of questionnaire scores with the condition (safe vs risky). Positive values indicate increased distances in the risky condition compared to safe. **C-D:** β coefficients (95%CI) of the same questionnaire scores predicting the probability of choosing a unique option. **C** shows the main effects, **D** the interaction with the condition.

4 Conclusions

In conclusion, exploration seems to be guided by uncertainty, both in safe environments and in environments where exploration comes with risks. However, subjects differ in whether they seek out uncertain options or avoid it. On average, subjects seem to avoid uncertainty more strongly in risky environments. There are, however, considerable individual differences in this: Subjects with an increased tendency for worry related or depressive traits were less averse to uncertainty in the risky condition, explored more and generalised more from observed rewards. Subjects with increased somatic anxiety in turn explored less in risky environments. One explanation for this behaviour could be that while all of the psychiatric trait scores, except for intolerance to uncertainty, were related to decreased self-reported uncertainty, subjects with high somatic anxiety dealt with this by playing it safe and sticking with more certain options, while subjects with high worry and depression related traits dealt with it by actively seeking out uncertain options and thereby reducing uncertainty.

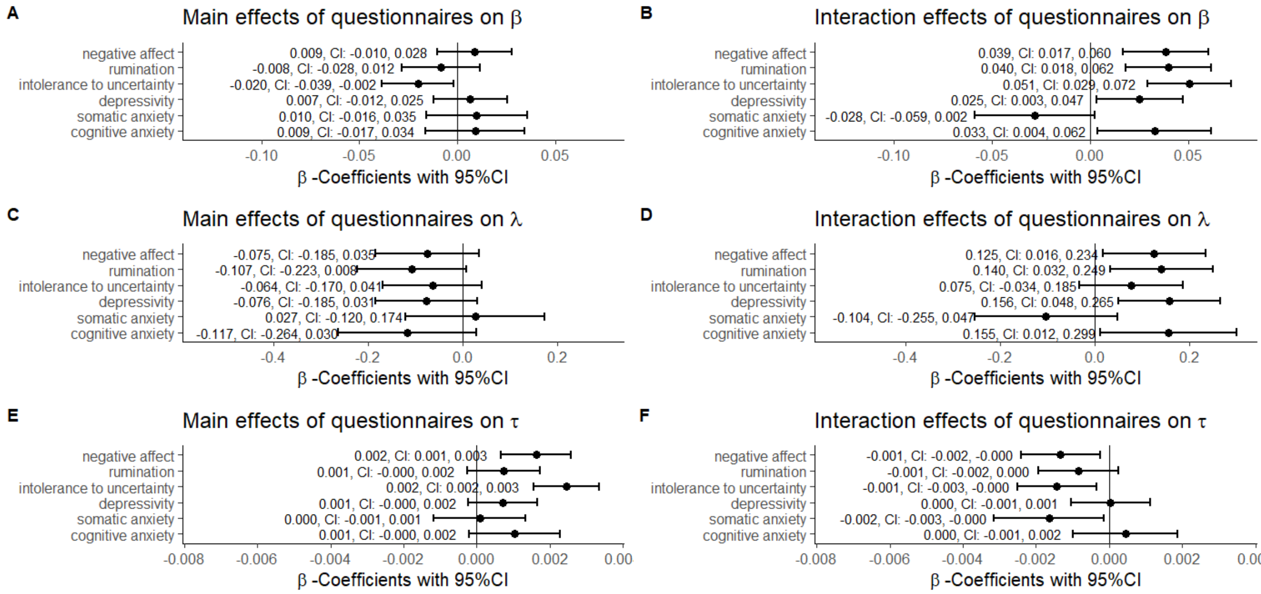

Figure 4: **A-B:** β coefficients (95%CI) of the different questionnaire scores in predicting the β parameter from the CB model. **A** shows the main effects, positive estimates indicate more positive β estimates. **B** shows the interaction effects of questionnaire scores with the condition (safe vs risky). Positive values indicate increased β parameter estimates in the risky condition compared to safe. **C-D:** β coefficients (95%CI) of the same questionnaire scores predicting the λ parameter from the GP learning model portion of the CB model. **C** shows the main effects, positive estimates indicate increased generalisation. **D** shows the interaction with the condition. Positive estimates indicate increased generalisation in the risky compared to the safe condition. **E-F:** Questionnaires predicting the τ parameter from the GP learning model portion of the CB model. **E** shows the main effects of questionnaire scores, positive estimates indicate increased random exploration. **F** shows the interaction effects with condition. Positive estimates indicate increased random exploration in the risky condition compared to the safe condition.

References

Carleton, R. N., Norton, M. A. P. J., & Asmundson, G. J. G. (2007). Fearing the unknown: A short version of the intolerance of uncertainty scale. *J. Anxiety Disord.*, *21*(1), 105–117. https://doi.org/10.1016/j.janxdis.2006.03.014

Fossati, A., Krueger, R. F., Markon, K. E., Borroni, S., & Maffei, C. (2013). 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*(6), 689–708. https://doi.org/10.1177/1073191113504984

Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *Elife*, *5*. https://doi.org/10.7554/eLife.11305

Kaufman, J., & Charney, D. (2000). Comorbidity of mood and anxiety disorders. *Depress. Anxiety*, *12 Suppl 1*, 69–76. https://doi.org/10.1002/1520-6394(2000)12:1+⟨69::AID-DA9⟩3.0.CO;2-K

Mossaheb, N., Becker, J., Schaefer, M. R., Klier, C. M., Schloegelhofer, M., Papageorgiou, K., & Amminger, G. P. (2012). 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*(2-3), 210–214. https://doi.org/10.1016/j.schres.2012.08.008

Ree, M. J., French, D., MacLeod, C., & Locke, V. (2008). 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*(3), 313–332.

Schulz, E., Wu, C. M., Huys, Q. J. M., Krause, A., & Speekenbrink, M. (2018). Generalization and search in risky environments. *Cogn. Sci.*, *42*(8), 2592–2620. https://doi.org/10.1111/cogs.12695

Sui, Y., Gotovos, A., Burdick, J., & Krause, A. (2015). Safe exploration for optimization with gaussian processes. In F. Bach & D. Blei (Eds.), *Proceedings of the 32nd international conference on machine learning* (pp. 997–1005). PMLR.

Trapnell, P. D., & Campbell, J. D. (1999). Private self-consciousness and the five-factor model of personality: Distinguishing rumination from reflection. *J. Pers. Soc. Psychol.*, *76*(2), 284–304. https://doi.org//10.1037//0022-3514.76.2.284

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nat Hum Behav*, *2*(12), 915–924. https://doi.org/10.1038/s41562-018-0467-4