Lower levels of directed exploration and reflective thinking are associated with greater anxiety and depression

Ryan Smith¹, Samuel Taylor¹, Robert C. Wilson², Annie E. Chuning¹, Michelle R. Persich³, Siyu Wang², William D.S. Killgore^{2,3}

¹ Laureate Institute for Brain Research

² Department of Psychology, University of Arizona

³ Department of Psychiatry, University of Arizona

Corresponding author:

Ryan Smith
Laureate Institute for Brain Research
6655 S Yale Ave, Tulsa, OK 74136, USA
Email: rsmith@laureateinstitute.org

Abstract

Anxiety and depression are often associated with strong beliefs that entering specific situations will lead to aversive outcomes – even when these situations are objectively safe and avoiding them reduces well-being. A possible mechanism underlying this maladaptive avoidance behavior is a failure to reflect on: 1) appropriate levels of uncertainty about the situation, and 2) how this uncertainty could be reduced by seeking further information (i.e., exploration). To test this hypothesis, we asked a community sample of 417 individuals to complete measures of reflective cognition, exploration, and symptoms of anxiety and depression. Consistent with our hypotheses, we found significant associations between each of these measures in expected directions (i.e., positive relationships between reflective cognition and strategic information-seeking behavior or “directed exploration”, and negative relationships between these measures and anxiety/depression symptoms). Further analyses suggested that the relationship between directed exploration and depression/anxiety was due to an ambiguity aversion promoting exploration in conditions where information-seeking was not beneficial (as opposed to under-exploration when more information would aid future choices). In contrast, reflectiveness was associated with greater exploration in appropriate settings and separately accounted for differences in reaction times, decision noise, and choice accuracy in expected directions. These results shed light on the mechanisms underlying information-seeking behavior and how they may contribute to symptoms of emotional disorders. They also highlight the possibility that reflectiveness and exploration could represent novel treatment targets for those who show low levels of these tendencies – consistent with the need to develop individualized precision medicine approaches within computational psychiatry.

Keywords: Directed Exploration; Random Exploration; Cognitive Reflection; Anxiety;

Depression; Explore-Exploit Dilemma

Introduction

Adaptive behavior requires striking an optimal balance between seeking desired outcomes and gathering information to best acquire those outcomes. The difficult task of finding this optimal balance is known as the explore-exploit dilemma. If one is confident in their current understanding of the world – and believes the contingencies in the world are stable – it makes sense to exploit current knowledge when making decisions to achieve one's goals. In contrast, if one is uncertain about the expected outcomes of different choices – or believes those outcomes can change over time – it is more adaptive to first gather more information. Research in psychology and neuroscience has produced a substantive literature characterizing both optimal and sub-optimal solutions to this dilemma that are exhibited by humans and other animals (reviewed in (1, 2)). One common strategy proposed in the reinforcement learning literature is random exploration, in which the drive to exploit high-value options is countered by randomness in the decision process (3). Another strategy that derives from optimal solutions to the explore-exploit dilemma is directed exploration, in which decisions are biased towards more uncertain and/or more informative options (4, 5). By now, empirical evidence for the existence of both types of exploration in human behavior is overwhelming (e.g., (6-13)).

The importance of solving the explore-exploit dilemma appropriately has more recently motivated the hypothesis that psychiatric disorders could be characterized by a poor balance in seeking information vs. reward (14). One possibility is that psychiatric disorders involve under-exploration via less information-seeking behavior in novel situations. This under-exploration promotes avoidance of feared situations that are in fact safe and prevents individuals from learning more adaptive patterns of behavior (15-17). Another possibility is over-exploration, in which individuals remain uncertain for too long (perhaps because they do not learn from their

experience) and fail to use current knowledge to efficiently seek out reward. These deficit profiles could also differ in different disorders and could be specific to random or directed exploration.

A few studies to date have begun to examine this hypothesis within substance use disorders, schizophrenia, and, to a lesser extent, emotional disorders (for a review, see (14)). For example, one study found that, relative to healthy participants, individuals with alcohol use disorder (AUD) displayed less exploratory behavior in the context of both gains (i.e., trying to maximize reward) and losses (i.e., trying to minimize negative outcomes), while individuals with binge eating disorder showed greater exploration than those with AUD in the context of losses (18). Reduced exploration was also observed across all participants in the context of losses relative to gains. A second study in a heterogeneous population of substance use disorder (SUD) patients found no difference in directed exploration compared to healthy participants, but found that SUDs were associated with slower learning rates from losses and greater randomness in choice (however, this was not clearly tied to random exploration; (19)). Nicotine smokers have also been found to make fewer exploratory choices and to possess higher learning rates relative to healthy participants (20); also see (21, 22).

Recent work in Schizophrenia patients found reduced directed exploration relative to healthy participants, while random exploration did not differ between groups (23). Exploratory behavior was not associated with negative symptoms, but random exploration was positively correlated with psychotic symptoms and cognitive impairment. Faster learning rates were also associated with more severe psychotic symptoms. Another recent study in healthy participants showed that normative levels of trait anxiety were positively associated with exploratory behavior and that an inverted-U pattern characterized performance differences; namely,

moderate anxiety and associated levels of exploration led to higher levels of performance than those with very high anxiety or very low anxiety (i.e., who explored too much or too little; (24)). This finding was in the context of a volatile environment with changing reward or loss probabilities and appeared to be driven by reduced reward-seeking behavior and an elevated drive to reduce uncertainty. This is consistent with previous work showing over-exploration in depression (25), and also consistent with suggestive evidence of greater exploration with higher anxiety and depression symptoms in one of the aforementioned studies on SUDs (19). However, greater anhedonia (in patients with Schizophrenia) has also been associated with reduced exploratory behavior in prior work (26). Therefore, while more research is needed, this growing body of work suggests that a range of potential differences in explore-exploit behavior may be present in the context of poor mental health – and could reflect state differences or vulnerability factors.

A further question pertains to the origins of variability in exploratory behavior. Normatively speaking, information-seeking should scale with estimates of one's own uncertainty. Therefore, under identical conditions, one would expect greater exploratory behavior in those who are more uncertain – yet this need not entail that this level of uncertainty is warranted given past observations. One may “jump to conclusions” and cease exploration prematurely, or uncertainty could persist despite having gathered sufficient information to infer the best course of action. One source of these differences in uncertainty estimation could be trait differences in cognitive reflection. This refers to the degree to which an individual tends to “go with their gut” and trust initial intuitions about how to respond vs. engaging effortful cognitive processes to reflect on the reliability of those intuitions. In recent years, such differences have been studied using performance measures such as the cognitive reflection test (CRT-7; (Toplak,

West, & Stanovich, 2014a)) and the comprehensive assessment of rational thinking (CART; (Stanovich, West, & Toplak, 2016)), as well with self-report measures of epistemic attitudes (Toplak, West, & Stanovich, 2014b). However, these reflective tendencies have not been studied in the context of explore-exploit behavior, nor have they been thoroughly investigated in relation to psychopathology.

In this study, we first test the competing hypotheses that variation in symptoms of depression and anxiety in a community sample will be associated with higher vs. lower levels of exploratory behavior, and we examine whether directed vs. random exploratory strategies are affected. We then test the hypothesis that differences in explore-exploit behavior are predicted by differences in cognitive reflection. Finally, we examine whether cognitive reflection predicts symptoms of depression and anxiety and whether this relationship may be accounted for by differences in explore-exploit behavior.

Methods

Participants

A convenience sample of students at the University of Arizona as well as individuals from the surrounding community (mean age = 23.75, SD = 5.61 years; minimum age of 18), 115 male and 301 female, was recruited from Tucson, AZ. Participants gave informed consent and were paid for their participation. This study was approved by the University of Arizona Institutional Review Board (Protocol #1607696724) and the Human Research and Protection Office of the U.S. Army Medical Research and Development Command (Protocol #A-19136.a and A-19136.b).

Measures

Depression and anxiety symptoms

Participants completed the Beck Depression Inventory (BDI-II), a 21-item scale that measures symptoms of depression (27). Each item on the BDI is scored from 0 to 3, with higher scores consistent with depressive symptomatology. Participants also completed the state-trait anxiety inventory (STAI), a 40-item scale that measures anxiety symptoms in the present moment (state) and in life more generally (trait) (28). This data on depression and anxiety symptoms was also utilized in a recent study examining questions unrelated to exploration or cognitive reflection (29).

Levels of directed and random exploration

To measure directed and random exploration, participants completed a previously validated task called the Horizon Task (13). In this task, participants are asked to repeatedly choose between one of two options that can win the player an unknown number of points (see **Figure 1A**). In this paper, we used a version of the task with 80 games, where each game includes either 5 or 10 repeated choices (40 games of each length, interleaved). Each option can pay out between 1 and 100 points drawn from a Gaussian distribution (rounded to the nearest integer) with a fixed standard deviation of 8 points. The generative means of the underlying Gaussians are different for the two options and remain stable within a game. In each game, the mean of one option is set to either 40 or 60 points and the mean of the other is set such that the difference between the two means is 4, 8, 12, 20, or 30 points (counterbalanced across games with respect to left vs. right option and mean differences; see **Figure 1B** and **1C**). Through a set of illustrated onscreen instructions, participants are explicitly told that the means of the two options are constant across trials within a game and that the variability in the values for both options is constant across all

games. Participants are told that one option, in any given game, is always better on average, and they are instructed to maximize the points they earn.

Figure 1. (A) Depiction of example Horizon 1 (H1) and Horizon 6 (H6) games at the first free choice. (B-C) Probability of choosing the option on the right depending on differences in the generative means of the two options. (D-E). Estimates of decision noise in H1 and H6 games (within equal and unequal information conditions), showing the predicted pattern of greater noise in H6 games. (F) Estimates of information bonus in H1 and H6 games, showing the predicted pattern of information bonus in H6 games. (G) Fraction of correct responses (i.e., choice of the option with the higher generative mean) for each trial in H1 and H6 games. (H-I) Inverted-U relationship between information bonus and choice accuracy (i.e., choice of the option with the higher true mean) in H1 and H6 games.

In each game, the first four choices are forced (i.e., participants are told which option to choose). In “horizon 1” (H1) games they are then allowed to make one free choice, while in “horizon 6” (H6) games they are allowed to make six free choices (see **Figure 1A**). The slot machine display for each option shows the point value for the chosen option and an “XX” for the unchosen option, and outcomes of the whole sequence of choices remain on the screen throughout the game (i.e., each slot machine has either 5 or 10 blank boxes that are sequentially filled with each choice in the game; see **Figure 1A**). The forced choices could provide either equal information (two outcomes from each option [2 2]) or unequal information (one outcome from one option and three from the other [1 3]). This setup allows assessment of the first free choice under equal vs. unequal information when participants expect to make one or six future choices. Previous studies using this task have shown that, during unequal information games, participants are more likely to choose the more uncertain option (i.e., where they have only seen a single outcome) in H6 games compared to H1 games (13). This is interpreted as a form of directed exploration in which the initial choice is expected to improve the ability to maximize reward on the subsequent five choices (i.e., whereas this would have no advantage in H1 games). In both equal and unequal information conditions, participants are also more likely to choose the option with the lower mean (of the forced choice outcomes) in H6 games than H1 games. This is interpreted as a form of random exploration, in which the lower mean choice also facilitates information gain and can benefit future choices (i.e., which would also have no advantage in H1 games).

Computational modelling is used to quantify individual differences in levels of directed exploration and random exploration based on the option chosen on the first free choice. In this model, the value Q_a of each option a is used to make probabilistic choices (i.e., a higher value

leads to a higher probability of selection). The values of Q_a are computed as follows:

$$Q_a = R_a + \alpha I_a + B s_a$$

Here, the experimental parameters are: R_a the expected reward, I_a the information value, and s_a the side on which option a is presented. Subject-specific parameters are: α , the value of information, which acts as an information bonus, and B the spatial bias. The relationship between Q_a and the probability of selecting each option is also influenced by decision noise within a logistic function with subject-specific variance parameter: σ_d . The probability of choosing option a over option b is then:

$$P_a = \frac{1}{1 + \exp\left(\frac{R_b - R_a + \alpha(I_b - I_a) + B(s_b - s_a)}{\sigma_d}\right)}$$

The values of R_a and R_b are specified as the observed means of the forced-choice outcomes of the respective options. I_a is specified such that if option b is more informative (i.e., when only one outcome has been observed for that option in unequal information games) then $I_b - I_a = 1$, whereas this value is made negative if option a is more informative. $I_b - I_a = 0$ within equal information games. The spatial variable (i.e., identifying the option being on the left vs. the right) is set to $s_b - s_a = 1$ when option b is on the right; this value is made negative when option a is on the right. Fitting this model to participants' behavior allows estimation of the values of the information bonus (α), decision noise (σ_d), and spatial bias (B) that best account for patterns of decision-making at the individual level (here, this was accomplished using the gradient-based 'fmincon' function in MATLAB to find parameters that maximize the probability of participants' behavior under the model). These values can then be used as individual difference measures to

assess how each may relate to other variables of interest. The information bonus is estimated separately for H1 and H6 conditions, and the difference in these two estimates ($\alpha_{H6} - \alpha_{H1}$) is taken as the measure of directed exploration. The decision noise is estimated separately for H1 and H6 conditions as well as for equal and unequal information conditions. This allows calculation of two measures of random exploration – one for equal and one for unequal information conditions – by taking the difference between σ_d in H6 and H1 games for each condition. In this study, we did not perform further analyses on spatial bias estimates, which are included primarily to remove this as a confounding influence on estimating the parameters of interest.

General reflectiveness measures

We used three related measures to assess general reflective cognition: the cognitive reflection test (CRT-7; (30)), the actively open-minded thinking scale (AOMTS; (31)), and the 2-subscale version of the comprehensive assessment of rational thinking (CART; (32)). The CRT-7 asks seven short questions designed such that there is an immediately intuitive, but incorrect, answer, and a correct answer that, while not logically difficult, requires the individual to devote effortful cognitive resources instead of immediately choosing the intuitively appealing response.

Example item:

“If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?” (intuitive answer: 100 minutes; correct answer: 5 minutes)

It tests the tendency to “stop and think” before immediately trusting one’s intuition.

The AOMTS is a self-report scale which asks individuals to rate 30 statements, from 1 (strongly disagree) to 6 (strongly agree), which describe more or less reflective or “rational” attitudes. Example item:

“I like to gather many different types of evidence before I decide what to do.”

Higher scores indicate more open-minded, reflective attitudes.

The CART assesses vulnerability to various common reasoning biases that arise (in part) from insufficient engagement of reflective capacities. The 2-subscale version we used includes statistical and scientific reasoning problems. Example item:

“Dice game: Even numbers win and odd numbers lose on a die throw. The fair die has six sides, with three even and three odd numbers. Jan has thrown seven odd numbers in a row. What are her chances of throwing an even number on her next throw?” (correct answer: 3/6).

Higher scores indicate a greater tendency to engage effortful cognition and avoid common reasoning biases during problem-solving.

When assessing differences on these reflective cognition measures, it is important to first account for individual differences in general intelligence. To do so, we also asked participants to complete the 2-subscale Wechsler abbreviated scale of intelligence (WASI-II; (33)), a common measure of IQ.

Analyses

To assess each of the hypothesized relations between exploration, cognitive reflection,

and depression/anxiety symptoms, we ran JZS Bayes factor analyses with default prior scales in R (BayesFactor package (34, 35)) comparing null (intercept only) models to the space of models that included all combinations of main effects of age, sex, IQ, and the predictor of interest (and interactions between sex and the predictor of interest) on the relevant target variables.

Interactions with sex were included because sex differences have been observed in both cognitive reflection measures (36-41) and symptoms of emotional disorders (42-46). Models included exploration measures as predictors of depression/anxiety symptoms, cognitive reflection measures as predictors of exploration measures, and cognitive reflection measures as predictors of depression/anxiety symptoms. To assess whether cognitive reflection measures and exploration measures accounted for shared vs. independent variance in symptoms, we ran further models including both types of measures as predictors of symptom measures in cases where both measures were predictors of symptoms in independent models.

When comparing models, a Bayes factor (BF) represents the ratio of the probability of the data under one model vs. another, indicating the relative evidence for different models. If H_0 indicates the null hypothesis, H_1 indicates the alternative hypothesis, and d indicates the data, then:

$$BF = \frac{p(d|H_1)}{p(d|H_0)}$$

A higher BF value indicates greater evidence for the alternative hypothesis; e.g., $BF = 3$ indicates three times as much evidence for the alternative hypothesis than for the null hypothesis. We adopt the guidelines described in Lee and Wagenmakers (47) for interpreting the strength of evidence provided by different BF values: $BF = 1-3$, poor/anecdotal evidence; $3-10$, moderate evidence; $10-30$, strong evidence, $30-100$, very strong evidence, >100 , extremely strong

evidence. Empirical means and credible intervals for each variable in winning models were found by sampling from the posterior of the model, using Markov chain Monte-Carlo (MCMC) sampling for 10,000 iterations.

After identifying the winning models, we conducted post-hoc Pearson correlations to better interpret the strength of relationships between each predictor variable and each target variable. To better interpret findings regarding directed and random exploration, we also examined additional Horizon Task performance measures. This included the information bonus (α) and decision noise (σ_d) parameters for H1 and H6 trials (see **Figure 1D-F**), total choice accuracy (i.e., number of overall free choices corresponding to the arm with the higher true mean value; see **Figure 1G**), and average reaction time for the first free choice on each trial. We also calculated the distance between information bonus values and their optimal values. This was motivated by inspection of scatterplots (shown in **Figure 1H-I**) relating exploration measures to choice accuracy, which suggested inverted-U effects for information bonus values for both H1 and H6 trials (IBH1 and IBH6). In each case, this was captured by the fact that a quadratic model fit the data better than a linear model (Akaike information criterion [AIC] = -606.78 vs. AIC = -507.69 for IBH1; AIC = -2209.45 vs. AIC = -2168.41 for IBH6). As expected theoretically, choice accuracy was highest for IBH1 values near 0. The empirically optimal value for IBH6 (based on the maximum of a quadratic model) was 6.78. Therefore, we computed absolute deviations from optimal values for both conditions (i.e., from 0 for IBH1 and 6.78 for IBH6) and used these values in some secondary analyses.

Results

The descriptive statistics for all variables are shown in **Table 1** and **Table 2**. These tables

also include comparisons between males and females, motivated by the reliable sex differences found for these variables in previous work described above. As seen there, depression and anxiety were primarily within the subclinical range, but values indicative of moderate to severe symptoms ($BDI > 20$, $STAI > 40$) were also represented. STAI Trait scores were greater in females, while CRT-7, AOMTS, and CART scores were greater in males. Females showed faster first-choice RTs in the Horizon Task as well as moderate evidence for greater decision noise in H1 trials with unequal information. No sex differences were found in choice accuracy or measures of directed or random exploration.

Table 1. Descriptive statistics and tests of sex differences for age, symptoms, and reflectiveness

	Total	Female	Male	Usable Data (N)*	Effect of Sex**
Age	23.75 (5.61)	23.43 (5.49)	24.61 (5.84)	Female: 301 Male: 115 Total: 416	$t(414) = -1.93$ $d = -0.19$ $BF = 0.71$
BDI	7.31 (6.79)	7.65 (6.92)	6.4 (6.36)	Female: 301 Male: 115 Total: 416	$t(414) = 1.69$ $d = 0.17$ $BF = 0.47$
STAI State	30.12 (8.22)	30.21 (8.14)	29.87 (8.46)	Female: 300 Male: 115 Total: 415	$t(413) = 0.38$ $d = 0.04$ $BF = 0.13$
STAI Trait	36.55 (10.98)	37.51 (10.75)	34.06 (11.25)	Female: 300 Male: 115 Total: 415	$t(413) = 2.89$ $d = 0.28$ $BF = 6.37$
WASI	108.37 (11.47)	107.92 (11.14)	109.55 (12.24)	Female: 300 Male: 115 Total: 415	$t(413) = -1.3$ $d = -0.13$ $BF = 0.27$
CART	10.52 (2.44)	10.1 (2.26)	11.62 (2.56)	Female: 299 Male: 114 Total: 413	$t(411) = -5.91$ $d = -0.58$ $BF > 100$

AOMTS	134.78 (13.44)	133.57 (12.98)	137.93 (14.16)	Female: 300 Male: 115 Total: 415	t(413) = -2.98 d = -0.29 BF = 8.37
CTR-7	2.65 (2.05)	2.35 (1.95)	3.44 (2.13)	Female: 291 Male: 109 Total: 400	t(398) = -4.84 d = -0.49 BF > 100

*Not all data were available from all participants. We therefore report Ns for each measure.

**Results of Bayesian two-sample t-tests. Bayes factors (BFs) supporting a group difference with values greater than 3 are bolded. Effect sizes (Cohen's d) are also shown.

Table 2. Descriptive statistics and tests of sex differences for Horizon Task variables

	Total	Female	Male	Usable Data (N)*	Effect of Sex**
Dir. Exp.	6.26 (11.79)	5.79 (11.85)	7.48 (11.57)	Female: 301 Male: 115 Total: 416	t(414) = -1.3 p = 0.19 d = -0.13 BF = 0.27
Info. Bonus H1	-0.73 (8.82)	-1.01 (9.16)	0 (7.86)	Female: 301 Male: 115 Total: 416	t(414) = -1.04 p = 0.3 d = -0.1 BF = 0.2
Info. Bonus H6	5.53 (11.31)	4.78 (11.4)	7.47 (10.86)	Female: 301 Male: 115 Total: 416	t(414) = -2.18 p = 0.03 d = -0.21 BF = 1.17
Rand. Exp. (Equal Info)	3.23 (6.37)	3.13 (6.4)	3.51 (6.32)	Female: 301 Male: 115 Total: 416	t(414) = -0.55 p = 0.59 d = -0.05 BF = 0.14
Dec. Noise H1 (Equal Info)	4.4 (7.25)	4.69 (7.38)	3.64 (6.88)	Female: 301 Male: 115 Total: 416	t(414) = 1.33 p = 0.19 d = 0.13 BF = 0.28
Dec. Noise H6 (Equal Info)	7.63 (7.78)	7.82 (7.65)	7.15 (8.12)	Female: 301 Male: 115 Total: 416	t(414) = 0.79 p = 0.43 d = 0.08 BF = 0.16

Rand. Exp. (Unequal Info)	2.63 (6.51)	2.51 (6.76)	2.94 (5.82)	Female: 301 Male: 115 Total: 416	$t(414) = -0.6$ $p = 0.55$ $d = -0.06$ $BF = 0.14$
Dec. Noise H1 (Unequal Info)	4.96 (6.74)	5.5 (7.17)	3.54 (5.2)	Female: 301 Male: 115 Total: 416	$t(414) = 2.68$ $p = 0.008$ $d = 0.26$ BF = 3.73
Dec. Noise H6 (Unequal Info)	7.59 (7.12)	8.01 (7.38)	6.48 (6.28)	Female: 301 Male: 115 Total: 416	$t(414) = 1.98$ $p = 0.05$ $d = 0.19$ $BF = 0.78$
Info. Bonus – Deviation from Optimal H1	5.71 (6.76)	6.21 (6.8)	4.4 (6.5)	Female: 301 Male: 115 Total: 416	$t(414) = 2.46$ $p = 0.01$ $d = 0.24$ $BF = 2.16$
Info. Bonus – Deviation from Optimal H6	9.18 (6.7)	9.34 (6.82)	8.76 (6.39)	Female: 301 Male: 115 Total: 416	$t(414) = 0.79$ $d = 0.08$ $BF = 0.16$
First-Choice RT	1.46 (0.73)	1.38 (0.62)	1.67 (0.92)	Female: 301 Male: 115 Total: 416	$t(414) = -3.69$ $d = -0.36$ BF = 76.3
First-Choice RT H1	1.41 (0.78)	1.32 (0.67)	1.62 (0.97)	Female: 301 Male: 115 Total: 416	$t(414) = -3.56$ $d = -0.35$ BF = 49.52
First-Choice RT H6	1.51 (0.74)	1.44 (0.63)	1.72 (0.95)	Female: 301 Male: 115 Total: 416	$t(414) = -3.51$ $d = -0.34$ BF = 41.73
Choice Accuracy	0.23 (0.03)	0.22 (0.03)	0.23 (0.03)	Female: 301 Male: 115 Total: 416	$t(414) = -2$ $d = -0.2$ $BF = 0.82$

Choice Accuracy H1	0.8 (0.13)	0.79 (0.13)	0.82 (0.13)	Female: 301 Male: 115 Total: 416	$t(414) = -1.9$ $d = -0.19$ BF = 0.68
Choice Accuracy H6	0.13 (0.02)	0.13 (0.02)	0.13 (0.02)	Female: 301 Male: 115 Total: 416	$t(414) = -1.83$ $d = -0.18$ BF = 0.6

*Not all data were available from all participants. We therefore report Ns for each measure.

**Results of Bayesian two-sample t-tests. Bayes factors (BFs) supporting a group difference with values greater than 3 are bolded. Effect sizes (Cohen's d) are also shown.

Initial validity analyses

Consistent with previous work, IBH6 values were greater than IBH1 values ($t(415) = 10.83$, $d = .53$, BF > 100). Likewise, decision noise in H6 (DNH6) trials was also greater than in H1 (DNH1) trials in both equal ($t(415) = 10.34$, $d = .51$, BF > 100) and unequal ($t(415) = 8.25$, $d = .40$, BF > 100) information conditions (see **Figure 1D-E**). Plotting the proportion of optimal choices as a function of choice number showed the expected pattern in which 1) the lower mean option was chosen more often at the first free choice in H6 than H1 trials, and 2) performance improved with each subsequent choice in H6 trials (**Figure 1G**).

Intercorrelations between Horizon Task measures were consistent with expectations and supported the validity of model parameters. Decision noise estimates in all conditions showed negative correlations with choice accuracy (r s between $-.57$ and $-.83$, all BFs > 100) and first-choice RTs (r s between $-.15$ and $-.30$, BFs between 13.5 and > 100). First-choice RTs were positively correlated with choice accuracy (r s between $.17$ and $.28$, BFs between 57.5 and > 100). Random exploration scores for equal and unequal information conditions were both positively associated with directed exploration (r s = $.175$ and $.183$, BFs = 67.8 and > 100) and choice accuracy in H1 trials (r s = $.21$ and $.18$, BFs > 100). IBH1 values were negatively

correlated with choice accuracy (r s between $-.16$ and $-.25$, BFs between 26.5 and > 100).

Directed exploration was positively correlated with choice accuracy (r s between $.22$ and $.24$, all BFs > 100). Neither directed nor random exploration was associated with first-choice RTs (r s between $-.06$ and $.06$, BFs between 0.12 and $.26$).

Cognitive reflectiveness measures were correlated in the expected directions:

CART:CRT-7 ($r = .66$, BF > 100); CART:AOMTS ($r = .37$, BF > 100); CRT-7:AOMTS ($r = .30$, BF > 100). Each measure also showed the expected moderate correlation with WASI IQ scores (CART: $r = .54$, BF > 100 ; CRT-7: $r = .54$, BF > 100 ; AOMTS: $r = .32$, BF > 100).

Symptom measures were also correlated in the expected directions: STAI State:BDI ($r = .60$, BF > 100); STAI Trait:BDI ($r = .80$, BF > 100); STAI State:STAI Trait ($r = .68$, BF > 100).

Exploration as a predictor of symptoms

As shown in **Table 3**, the best models included directed exploration as a predictor of both STAI State scores (negative relationship; BF = 7.56 compared to an intercept-only model; moderate evidence) and BDI scores (negative relationship; BF = 16.62 compared to an intercept-only model; strong evidence; see **Figure 2**). STAI Trait scores were best explained by a model that only included sex.

To interpret these results, we ran several post-hoc analyses. We first ran correlations between BDI/STAI State scores and IBH1 and IBH6 values separately. Correlations were weakly positive with IBH1 (STAI State: $r = .14$, BF = 6.66 ; BDI: $r = .11$, BF = 1.7) but absent in IBH6 (STAI State: $r = -.04$, BF = $.17$; BDI: $r = -.08$, BF = $.37$), suggesting that the negative relationship

between directed exploration and BDI/STAI State scores was primarily accounted for by greater information-seeking in H1 trials, as opposed to lower information-seeking in H6 trials. As the relationship between IBH1 and choice accuracy reported above was negative, this suggested that stronger symptoms were associated with a suboptimal strategy. However, choice accuracy did not appear directly related to these symptoms in IBH6 (STAI State: $r = -.08$, $BF = .39$; BDI: $r = -.11$, $BF = 1.29$). Symptoms were also not associated with deviation from optimal IBH1 or IBH6 values (r s between $.02$ and $.08$, BF s between $.12$ and $.46$) or with first-choice RTs (r s between $-.03$ and $-.10$, BF s between $.14$ and $.98$).

Figure 2. Zero-order relationships between directed exploration and depression (BDI) and state anxiety (STAI). Relationships between information bonus in H1 and H6 trials suggested this was explained by greater information seeking in H1 trials with more severe symptoms.

Random exploration values were not included in the winning models predicting any symptom measure (see **Table 3**). There was also no relationship between symptoms and decision noise in H1 or H6 trials (r s between $-.02$ and $.07$, BF s between $.11$ and $.29$), except for possible anecdotal

evidence of a positive relationship between BDI and decision noise in H6 trials with unequal information ($r = .12$, $BF = 2.56$).

Table 3. Exploration as a predictor of symptoms

	Dir. Exp.	Rand. Exp. (Equal Info)	Rand. Exp. (Unequal Info)
Outcome Variable*			
STAI State	<p>Best Model (M1): Dir. Exp.</p> <p>$b = -0.1$, $CI = [-0.17, -0.03]$</p> <p>M1/Intercept: BF = 7.56</p> <p>2nd Best Model (M2): Dir. Exp. + Age</p> <p>M1/M2: $BF = 2.84$</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414 	<p>Best Model: Intercept only</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414 	<p>Best Model: Intercept only</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414
STAI Trait	<p>Best Model (M1): Sex</p> <p>$b = 1.67$, $CI = [0.53, 2.81]$</p> <p>M1/Intercept: $BF = 6.75$</p> <p>2nd Best Model (M2): Sex + Dir. Exp.*Sex</p> <p>M1/M2: $BF = 1.28$</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414 	<p>Best Model (M1): Sex</p> <p>$b = 1.68$, $CI = [0.49, 2.86]$</p> <p>M1/Intercept: $BF = 6.75$</p> <p>2nd Best Model (M2): Sex + Rand. Exp. (Equal Info)*Sex</p> <p>M1/M2: $BF = 6.42$</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414 	<p>Best Model (M1): Sex</p> <p>$b = 1.67$, $CI = [0.52, 2.84]$</p> <p>M1/Intercept: $BF = 6.75$</p> <p>2nd Best Model (M2): Sex + Rand. Exp. (Unequal Info)</p> <p>M1/M2: $BF = 7.27$</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414

BDI	<p>Best Model (M1): Dir. Exp. $b = -0.09$, $CI = [-0.14, -0.03]$ M1/Intercept: BF = 16.62</p> <p>2nd Best Model (M2): Dir. Exp. + WASI</p> <p>M1/M2: BF = 1.67</p> <p>N: - Female: 300 - Male: 115 - Total: 415</p>	<p>Best Model: Intercept only</p> <p>N: - Female: 300 - Male: 115 - Total: 415</p>	<p>Best Model: Intercept only</p> <p>N: - Female: 300 - Male: 115 - Total: 415</p>
-----	---	---	---

Note: For ease of inspection, the model predictors and BF's consistent with a priori hypotheses have been bolded. *Not all data were available from all participants. We therefore report Ns for each measure.

Cognitive reflection as a predictor of symptoms

As shown in **Table 4**, the best models included AOMTS as a predictor of STAI State scores and BDI scores (negative relationship in both cases; BF's = 3.3 and 3.7, respectively; see **Figure 3**). STAI Trait scores remained best accounted for by a model including only sex as a predictor.

Figure 3. Relationships between self-reported reflectiveness (AOMTS) and depression (BDI) and state anxiety (STAI). Greater reflectiveness was associated with lower symptom levels.

Table 4. Cognitive reflection as a predictor of symptoms

Outcome Variable*	CART	CRT	AOMTS
STAI State	Best Model: Intercept only N: - Female: 298 - Male: 114 - Total: 412	Best Model: Intercept only N: - Female: 290 - Male: 109 - Total: 399	Best Model (M1): AOMTS $b = -0.08$, $CI = [-0.13, -0.02]$, M1/Intercept: BF = 3.30 2nd Best Model (M2): AOMTS + AOMTS*Sex M1/M2: BF = 2.70 N: - Female: 299 - Male: 115 - Total: 414

STAI Trait	Best Model (M1): Sex $b = 1.66$, CI = [0.5, 2.81] M1/Intercept: BF = 6.05 2nd Best Model (M2): Sex + CART M1/M2: BF = 6.81 N: - Female: 298 - Male: 114 - Total: 412	Best Model (M1): Sex $b = 1.64$, CI = [0.43, 2.82] M1/Intercept: BF = 4.78 2nd Best Model (M2): Sex + CRT M1/M2: BF = 3.25 N: - Female: 290 - Male: 109 - Total: 399	Best Model (M1): Sex $b = 1.68$, CI = [0.51, 2.82] M1/Intercept: BF = 6.75 2nd Best Model (M2): Sex + AOMTS M1/M2: BF = 2.20 N: - Female: 299 - Male: 115 - Total: 414
BDI	Best Model (M1): WASI $b = -0.06$, CI = [-0.12, -0.01] M1/Intercept: BF = 1.02 2nd Best Model (M2): CART M1/M2: BF = 1.20 N: - Female: 298 - Male: 114 - Total: 412	Best Model (M1): WASI $b = -0.07$, CI = [-0.13, -0.01] M1/Intercept: BF = 1.71 2nd Best Model (M2): WASI + Sex M1/M2: BF = 2.98 N: - Female: 290 - Male: 109 - Total: 399	Best Model (M1): AOMTS $b = -0.06$, CI = [-0.11, -0.02] M1/Intercept: BF = 3.71 2nd Best Model (M2): AOMTS + WASI M1/M2: BF = 2.62 N: - Female: 299 - Male: 115 - Total: 414

Note: For ease of inspection, the model predictors and BF's consistent with a priori hypotheses have been bolded. *Not all data were available from all participants. We therefore report Ns for each measure.

Cognitive reflection as a predictor of exploration

As shown in **Table 5**, when CART was included as the predictor of interest for directed exploration, the winning model included effects of CART (positive relationship) and an

interaction between CART and sex (stronger relationship between directed exploration and CART in males; $BF = 5.94$; moderate evidence; see **Figure 4**). The main effect of the CART was most important, as the data became 26.22 times less probable when it was removed, but only 1.43 times less probable when the interaction was removed. When assessing CRT-7 as the predictor of interest for DE, the winning model only included CRT-7 (positive relationship; $BF = 3.79$; moderate evidence). Similarly, when assessing AOMTS as the predictor of interest for DE, the winning model only included AOMTS (positive relationship; $BF = 5.12$; moderate evidence).

To interpret these results, we did several post-hoc analyses. We first ran correlations between CART/CRT-7/AOMTS and IBH1 and IBH6 values separately. For CART and CRT-7, positive correlations were present with IBH6 (CART: $r = .16$, $BF = 20.18$; CRT-7: $r = .17$, $BF = 49.82$). AOMTS was not correlated with IBH6, and no reflectiveness measures were correlated with IBH1 (r s between $-.11$ and $.06$, BF s between $.13$ and 1.23). This suggested that greater reflectiveness led to greater increases in exploration on H6 trials. We next examined whether reflectiveness measures predicted deviations from optimal IBH1 and IBH6 values – finding that higher CART and CRT-7 scores predicted less deviation from optimal IBH1 values (CART: $r = -.25$, $BF > 100$; CRT-7: $r = -.27$, $BF > 100$) and that higher CRT-7 scores also predicted less deviation from optimal IBH6 values ($r = -.17$, $BF = 29.58$). We also assessed how cognitive reflectiveness values related to choice accuracy and first-choice RTs. Choice accuracy was higher with greater reflectiveness scores on all measures (CART: $r = .44$, $BF > 100$; CRT-7: $r = .40$, $BF > 100$; AOMTS: $r = .20$, $BF > 100$), and those with higher reflectiveness scores also showed slower RTs (CART: $r = .24$, $BF > 100$; CRT-7: $r = .23$, $BF > 100$; AOMTS: $r = .27$, $BF > 100$).

Figure 4. Relationships between self-reported reflectiveness (CART, CRT-7, AOMTS), directed exploration, and decision noise. Greater reflectiveness was associated greater exploration and reduced decision noise.

The effect of CART/CRT-7 on IBH6 was not accounted for by differences in IQ when models were compared that also included age, sex, or WASI scores (winning models only included CART or CRT-7, with BFs of 16.18 and 38.06 relative to an intercept-only model, respectively). The same was true of relationships between choice accuracy and all three measures of reflectiveness, where winning models included both reflectiveness measures and WASI as an additional predictor (BFs > 100 in all cases). CART and CRT-7 were more important, with data becoming > 100 times less probable when each was removed, while removing the WASI only made the data 2.18 and 12.52 times less probable, respectively. However, WASI was more important than AOMTS, with data becoming > 100 times less probable when WASI was removed, but only 1.14 times less probable when AOMTS was removed.

As shown in **Table 5**, random exploration was not predicted by CART, CRT-7, or AOMTS in

the best models. However, there was very strong evidence for negative relationships between CART/CRT-7 and decision noise in both H1 (DNH1) and H6 (DNH6) trials (r s between $-.25$ and $-.39$, all $BFs > 100$; see **Figure 4**). There was also evidence for relationships between decision noise and AOMTS (r s between $-.15$ and $-.21$, BFs between 13.28 and > 100), with the exception of DNH6 on trials with unequal information.

Table 5. Cognitive reflection as a predictor of exploration

Outcome Variable*	CART	CRT	AOMTS
Dir. Exp.	<p>Best Model (M1): CART + CART*Sex</p> <p>- CART: $b = 0.77$, $CI = [0.29, 1.25]$. VI: BF = 26.22</p> <p>- CART*Sex: $b = -0.52$, $CI = [-1, -0.03]$. VI: BF = 1.43</p> <p>M1/Intercept: BF = 5.94</p> <p>2nd Best Model (M2): CART</p> <p>M1/M2: $BF = 1.43$</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 298 - Male: 114 - Total: 412 	<p>Best Model (M1): CRT</p> <p>$b = 0.75$, $CI = [0.2, 1.31]$</p> <p>M1/Intercept: BF = 3.79</p> <p>2nd Best Model (M2): WASI</p> <p>M1/M2: $BF = 1.98$</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 290 - Male: 109 - Total: 399 	<p>Best Model (M1): AOMTS</p> <p>$b = 0.12$, $CI = [0.03, 0.2]$</p> <p>M1/Intercept: BF = 5.12</p> <p>2nd Best Model (M2): AOMTS + WASI</p> <p>M1/M2: $BF = 1.2$</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414
Rand. Exp. (Equal Info)	<p>Best Model: Intercept only</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 298 - Male: 114 - Total: 412 	<p>Best Model: Intercept only</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 290 - Male: 109 - Total: 399 	<p>Best Model: Intercept only</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414

Rand. Exp. (Unequal Info)	Best Model: Intercept only N: - Female: 298 - Male: 114 - Total: 412	Best Model: Intercept only N: - Female: 290 - Male: 109 - Total: 399	Best Model: Intercept only N: - Female: 299 - Male: 115 - Total: 414
---------------------------	--	--	--

Note: VI = Variable importance. This is the BF of the best model compared to a model in which the relevant predictor is removed. For ease of inspection, the model predictors and BFs consistent with a priori hypotheses have been bolded. *Not all data were available from all participants. We therefore report Ns for each measure.

As shown in **Table 6**, the relationships between decision noise and CART/CRT-7 were not accounted for by differences in IQ. Specifically, these measures were not removed as predictors of DNH1 or DNH6 values when WASI, age, or sex were considered as additional predictors, and BFs strongly favored models with these reflectiveness measures included compared to when they were removed. AOMTS was retained as a predictor in the winning model for DNH1 in equal information trials, but all other DN values were better explained by WASI scores.

Table 6. Cognitive reflection as a predictor of decision noise

	CART	CRT	AOMTS
Outcome Variable			

Dec. Noise H1 (Unequal Info)	<p>Best Model (M1): CART + WASI</p> <p>- CART: $b = -0.64$, $CI = [-0.93, -0.35]$, VI: BF > 100</p> <p>- WASI: $b = -0.08$, $CI = [-0.14, -0.01]$, VI: BF = 2.57</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CART</p> <p>M1/M2: BF = 2.57</p> <p>N: - Female: 298 - Male: 114 - Total: 412</p>	<p>Best Model (M1): CRT + WASI</p> <p>- CRT: $b = -0.71$, $CI = [-1.05, -0.35]$, VI: BF > 100</p> <p>- WASI: $b = -0.08$, $CI = [-0.15, -0.02]$, VI: BF = 3.41</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CRT</p> <p>M1/M2: BF = 3.41</p> <p>N: - Female: 290 - Male: 109 - Total: 399</p>	<p>Best Model (M1): Sex + WASI</p> <p>- Sex: $b = 0.81$, $CI = [0.13, 1.49]$, VI: BF = 1.97</p> <p>- WASI: $b = -0.15$, $CI = [-0.2, -0.09]$, VI: BF > 100</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): WASI</p> <p>M1/M2: BF = 1.85</p> <p>N: - Female: 299 - Male: 115 - Total: 414</p>
Dec. Noise H6 (Unequal Info)	<p>Best Model (M1): CART</p> <p>- CART: $b = -0.7$, $CI = [-0.98, -0.43]$</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CART + WASI</p> <p>M1/M2: BF = 2.81</p> <p>N: - Female: 298 - Male: 114 - Total: 412</p>	<p>Best Model (M1): CRT</p> <p>- CRT: $b = -0.83$, $CI = [-1.16, -0.5]$</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CRT + WASI</p> <p>M1/M2: BF = 3</p> <p>N: - Female: 290 - Male: 109 - Total: 399</p>	<p>Best Model (M1): WASI</p> <p>- WASI: $b = -0.12$, $CI = [-0.17, -0.06]$</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): Sex + WASI</p> <p>M1/M2: BF = 2.04</p> <p>N: - Female: 299 - Male: 115 - Total: 414</p>

Dec. Noise H1 (Equal Info)	<p>Best Model (M1): CART + WASI</p> <p>- CART: $b = -0.95$, $CI = [-1.25, -0.64]$, VI: BF > 100</p> <p>- WASI: $b = -0.07$, $CI = [-0.14, -0.01]$, VI: BF = 1.18</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CART</p> <p>M1/M2: BF = 1.18</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 298 - Male: 114 - Total: 412 	<p>Best Model (M1): CRT + WASI</p> <p>- CRT: $b = -1$, $CI = [-1.38, -0.64]$, VI: BF > 100</p> <p>- WASI: $b = -0.09$, $CI = [-0.16, -0.03]$, VI: BF = 4.05</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CRT + Age + WASI</p> <p>M1/M2: BF = 3.2</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 290 - Male: 109 - Total: 399 	<p>Best Model (M1): AOMTS + WASI</p> <p>- AOMTS: $b = -0.06$, $CI = [-0.12, -0.01]$, VI: BF = 2.49</p> <p>- WASI: $b = -0.16$, $CI = [-0.22, -0.1]$, VI: BF > 100</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): AOMTS + WASI + AOMTS*Sex</p> <p>M1/M2: BF = 1.14</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414
Dec. Noise H6 (Equal Info)	<p>Best Model (M1): CART + WASI</p> <p>- CART: $b = -0.92$, $CI = [-1.25, -0.59]$, VI: BF > 100</p> <p>- WASI: $b = -0.1$, $CI = [-0.17, -0.02]$, VI: BF = 4.02</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CART</p> <p>M1/M2: BF = 4.02</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 298 - Male: 114 - Total: 412 	<p>Best Model (M1): CRT + WASI</p> <p>- CRT: $b = -0.93$, $CI = [-1.35, -0.52]$, VI: BF > 100</p> <p>- WASI: $b = -0.13$, $CI = [-0.2, -0.05]$, VI: BF = 25.07</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): CRT + Sex + WASI</p> <p>M1/M2: BF = 5.31</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 290 - Male: 109 - Total: 399 	<p>Best Model (M1): WASI</p> <p>- WASI: $b = -0.2$, $CI = [-0.26, -0.14]$</p> <p>M1/Intercept: BF > 100</p> <p>2nd Best Model (M2): AOMTS + WASI</p> <p>M1/M2: BF = 4.07</p> <p>N:</p> <ul style="list-style-type: none"> - Female: 299 - Male: 115 - Total: 414

Note: VI = Variable importance. This is the BF of the best model compared to a model in which the relevant predictor is removed. For ease of inspection, the model predictors and BF's consistent with a priori hypotheses have been bolded. *Not all data were available from all participants. We therefore report Ns for each measure.

Cognitive reflection and exploration as independent predictors of symptoms

The only cases in which both cognitive reflection measures and exploration measures predicted symptom measures in the winning models were with AOMTS and directed exploration as predictors of STAI State and BDI scores. When including models with both as possible predictors of STAI State, each was retained in the winning model ($BF = 14.58$ relative to an intercept-only model), suggesting they account for unique variance in anxiety symptoms. Directed exploration was a more important predictor, as the full model had a BF of 4.42 compared to a model removing this variable, while it only had a BF of 1.93 compared to a model removing AOMTS. When including models with both as possible predictors of BDI, each was retained in the winning model ($BF = 20.24$ relative to an intercept-only model), suggesting they also account for unique variance in depressive symptoms. Directed exploration was a more important predictor, as the full model had a BF of 5.46 compared to a model removing this variable, while it only had a BF of 2.08 compared to a model removing AOMTS.

Post-hoc correlations and sex differences

For the interested reader, post-hoc correlations between all variables are provided in **Figure 5**. These correlations recapitulated each of the primary results presented above.

	STAI State	STAI Trait	BDI	Dir. Exp.	Info. Bon. H1	Info. Bon. H6	Rand. Exp. [2 2]	Rand. Exp. [1 3]	Dec. Noise H1 [1 3]	Dec. Noise H1 [2 2]	Dec. Noise H6 [1 3]	Dec. Noise H6 [2 2]	Choice Acc.	First Choice RT	Dev. from opt. H1	Dev. from opt. H6
CART	-0.01	-0.02	-0.1	0.14	0.02	0.16	-0.01	0.06	-0.32	-0.39	-0.25	-0.38	0.44	0.24	-0.25	-0.08
	BF: 0.12	BF: 0.12	BF: 0.97	BF: 5.17	BF: 0.13	BF: 20.18	BF: 0.12	BF: 0.21	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF: 0.47
CRT	-0.03	0.04	-0.02	0.14	0.04	0.17	-0.01	0.04	-0.31	-0.37	-0.25	-0.35	0.4	0.23	-0.27	-0.17
	BF: 0.14	BF: 0.15	BF: 0.12	BF: 4.89	BF: 0.16	BF: 49.82	BF: 0.12	BF: 0.17	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF: 29.58
AOMTS	-0.13	-0.1	-0.13	0.14	-0.11	0.06	0.05	0.06	-0.16	-0.21	-0.09	-0.15	0.2	0.27	-0.09	-0.08
	BF: 4.38	BF: 1.08	BF: 4.48	BF: 6.76	BF: 1.23	BF: 0.25	BF: 0.19	BF: 0.25	BF: 20.83	BF > 100	BF: 0.7	BF: 13.28	BF > 100	BF > 100	BF: 0.58	BF: 0.45
STAI State	-0.15	0.14	-0.04	-0.05	-0.04	0.07	0.04	0.02	0	-0.08	-0.1	0.07	0.03			
	BF: 9.9	BF: 6.66	BF: 0.17	BF: 0.19	BF: 0.17	BF: 0.29	BF: 0.17	BF: 0.13	BF: 0.11	BF: 0.39	BF: 0.98	BF: 0.33	BF: 0.13			
STAI Trait	-0.11	0.1	-0.03	-0.01	0.03	0.03	-0.01	0.06	-0.02	-0.03	-0.03	0.03	0.02			
	BF: 1.19	BF: 0.94	BF: 0.14	BF: 0.12	BF: 0.14	BF: 0.14	BF: 0.12	BF: 0.24	BF: 0.13	BF: 0.14	BF: 0.14	BF: 0.14	BF: 0.12			
BDI	-0.16	0.11	-0.08	-0.01	0.08	0.06	0.05	0.12	0.03	-0.11	-0.06	0.08	0.04			
	BF: 20.51	BF: 1.7	BF: 0.37	BF: 0.12	BF: 0.37	BF: 0.22	BF: 0.18	BF: 2.56	BF: 0.15	BF: 1.29	BF: 0.24	BF: 0.46	BF: 0.17			

Figure 5. Post-hoc correlations (and associated BFs) between reflectiveness measures, symptoms, and Horizon Task measures. Asterisks indicate BFs greater than 3, 10, 30, and 100, per conventional cutoffs for levels of evidence from moderate to extremely strong.

Given the sex differences observed across several variables, we also ran exploratory correlations separately by sex to see if this might offer additional insights. These supplementary analyses suggested that observed relationships with symptoms were driven more strongly by females and that relationships between reflectiveness and exploration were more strongly driven by males (see **Supplementary Figure S1**). However, the difference in sample sizes between males and females limits the interpretability of these results, and they were not hypothesized. We mention them here simply as preliminary findings that could support generation of future hypotheses.

Discussion

In this study, we tested several predictions about the relationship between cognitive reflectiveness, exploratory behavior, and symptoms of depression/anxiety. We first tested the prediction that differences in depression and anxiety would be predicted by differences in exploratory behavior. This was motivated by two competing hypotheses: 1) a negative relationship would be observed between exploration and symptoms, consistent with maladaptive avoidance behavior, and 2) a positive relationship would instead be observed, consistent with sustained uncertainty about optimal choice and resulting in reduced reward-seeking. We also tested the subsequent prediction that differences in exploratory behavior and symptoms could be accounted for by differences in general reflectiveness.

In support of the first hypothesis, and in line with prior work linking anhedonia to reduced exploration (26), we found that directed exploration was lower in those with more severe depression and state anxiety symptoms. This is consistent with the possibility that symptoms of depression and anxiety are maintained by maladaptive exploratory drives.

However, further analyses indicated that these results were largely explained by greater exploration in H1 games (i.e., when exploration is not beneficial) as opposed to lower exploration in H6 games. This suggests that greater symptoms were associated with an increased drive to reduce ambiguity – that is, to reduce uncertainty at the potential cost of reduced task performance. This is consistent with previous work suggesting that depression and anxiety are associated with intolerance of uncertainty (e.g., see (48)) – but is less consistent with the hypothesis that reduced information-seeking prevents adaptive learning and maintains symptoms. However, it is important to consider that the current version of the Horizon Task is framed with respect to wins (i.e., reward maximization). Avoidance behavior might be more

specifically associated with reduced exploration in the context of avoiding negative outcomes.

As such, it remains an open question whether these same results would be found if this task were instead framed in terms of losses (i.e., where participants begin with a certain amount of money or points and must choose the option that will lead to smaller losses; e.g., see (18, 49)).

Although we find a negative relationship between symptoms and directed exploration, our findings linking elevated symptoms to greater ambiguity aversion share some similarity to the greater exploration in depression (25) and trait anxiety (24) observed in previous work.

However, there are important differences worth considering. First, these studies have used different tasks that do not distinguish directed and random exploration in the way we have done here. Second, the above-mentioned study on depression focused on comparing groups with low vs. high symptom severity using a different measure of depression. Third, our findings are with respect to state anxiety, while those found previously pertain to trait anxiety (for which we did not find significant results). Future studies will therefore be necessary to assess the consistency between these different findings.

With respect to our subsequent hypothesis about reflectiveness, we observed the predicted relationship in which directed exploration was positively associated with all three cognitive reflection measures. One measure of reflectiveness (AOMTS) was also negatively associated with depression and state anxiety. Notably, all results remained present when considering IQ as another explanatory variable in the space of possible models, suggesting that differences in cognitive ability did not account for the relationships between reflectiveness measures and either symptoms or exploration. In further model comparison including both AOMTS and directed exploration as symptom predictors, each predictor was retained in the best model. Omission of directed exploration from the winning model also led to greater reductions in model evidence

than omission of AOMTS, suggesting exploration had higher importance as an explanatory variable. This pattern of results therefore does not support a mediation-like relationship (i.e., reflectiveness and exploration each accounted for unique variance in symptoms). Indeed, while the relationship between directed exploration and symptoms was explained by greater exploration in H1 trials (as mentioned above), the relationship with two of the reflectiveness measures (CART and CRT-7) was better explained by greater exploration in H6 trials, and the third (AOMTS) was explained by lower exploration in H1 trials (i.e., each of which would be beneficial to maximizing long-run reward). Both CART and CRT-7 also specifically predicted lower deviations from optimal levels of exploration.

These results are consistent with the idea that greater trait reflectiveness promotes greater (and more adaptive) levels of information-seeking – offering one potential explanation for why some individuals engage in directed exploration more than others. Interestingly, one recent study using transcranial magnetic stimulation (TMS) found evidence for a causal role of right frontopolar cortex in controlling directed (but not random) exploration (especially in the long horizon condition) (50), and this brain region has also been linked to other reflective processes such as prospection and farsighted planning (51), as well as the management of competing goals (52). This could therefore suggest that directed exploration emerges (in part) as a result of a broader set of reflective processes focused on prospective decision-making in consideration of proximal vs. distal objectives.

In contrast, we found evidence against relationships between random exploration and both reflectiveness and symptoms. However, we observed that higher IQ and higher reflectiveness were associated with less decision noise in both H1 and H6 games within the Horizon Task. Model comparison also suggested that reflectiveness accounted for a greater (and unique)

proportion of the variance in decision noise relative to IQ. This was part of a broader pattern in which reflectiveness measures were also associated with slower RTs and higher choice accuracy, and decision noise was associated with faster RTs and lower choice accuracy. Together, this suggests that less reflection led to faster choices that were more random (across all task conditions) and therefore led to reduced performance. Measures of reflectiveness are therefore not sensitive to changes in randomness between task conditions (i.e., random exploration), but they offer an explanation for the degree to which individuals think through their choices before responding.

This study has a number of strengths and limitations. One strength is the large sample size, which afforded sufficient power to detect small effect sizes. Another was the use of Bayesian analytic approaches that allowed for model comparison and assessment of evidence both for and against the null hypothesis. For example, in addition to supporting some of our primary hypotheses, we also find moderate evidence against relationships between trait anxiety and most other study variables, as well as evidence against relationships between random exploration and either reflectiveness or depression/anxiety symptoms. One limitation is that the sample had a greater number of females than males, which limited interpretability of potential effects of sex. For example, exploratory analyses suggested that some relationships we observed may differ in males and females. However, interactions with sex were most often not included in the winning models predicting symptoms or exploratory behavior. A larger number of males could have provided more evidence for such interactions or, alternatively, could have led such suggestive patterns to disappear. Thus, future work could follow up on this to examine this possibility further. Another limitation is that we used a shorter version of the Horizon Task (80 games), while previous studies have used longer versions (up to 320 games). This smaller number of

games could have made our parameter estimates less reliable, and therefore added noise to the data. This raises the possibility that stronger relationships may have been found if parameters were fit to larger amounts of choice data. Finally, our study focused on young adults. Explore-exploit behavior is known to change throughout the lifespan, especially in adolescence (53), which is a key time for the development of mental health.

To conclude, in this study we found that less directed exploration was associated with greater levels of anxiety and depression, and that this was explained by a type of ambiguity aversion in which information-seeking was elevated in contexts where it would not benefit subsequent choices – and may relate to broader patterns of intolerance of uncertainty. We also found that differences in directed exploration were accounted for by trait differences in cognitive reflectiveness, and that reflectiveness was also lower in those with greater depression/anxiety symptoms; yet, reflectiveness and exploration accounted for unique variance in symptoms. Reflectiveness levels further accounted for decision noise, reaction times, and choice accuracy. Together, these results shed light on the mechanisms underlying information-seeking behavior and how they may contribute to symptoms of emotional disorders. Although many effect sizes were small, they also suggest the possibility that reflectiveness and exploration could represent novel treatment targets in those who show low levels of these tendencies – consistent with the need to develop individualized precision medicine approaches within computational psychiatry. Future research should replicate this work and could extend it by examining the role of reflectiveness in other psychiatric contexts as well as how exploration is affected in situations that require avoidance of negative outcomes.

Funding: The study was supported by funding from the U.S. Army Medical Research and Development Command (W81XWH-16-1-0062) to WDSK. RS is supported by the William K. Warren Foundation, the Stewart G. Wolf scholarship, and the National Institute of General

Medical Sciences (P20GM121312). RCW is supported by the National Institute on Aging (R01AG061888).

Conflict of Interest: None of the authors have any conflicts of interest to disclose.

Informed consent: Informed consent was obtained from all individual participants included in the study.

References

1. Wilson RC, Bonawitz E, Costa VD, Ebitz RB. Balancing exploration and exploitation with information and randomization. *Curr Opin Behav Sci*. 2021;38:49-56.
2. Schulz E, Gershman SJ. The algorithmic architecture of exploration in the human brain. *Curr Opin Neurobiol*. 2019;55:7-14.
3. Sutton R, Barto A. Reinforcement Learning: An Introduction. 1998.
4. Bellman R. A problem in the sequential design of experiments. *Sankhyā: The Indian Journal of Statistics (1933-1960)*. 1956;16(3/4):221-9.
5. Gittins JC. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*. 1979;41(2):148-64.
6. Frank MJ, Doll BB, Oas-Terpstra J, Moreno F. Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nat Neurosci*. 2009;12(8):1062-8.
7. Speekenbrink M, Konstantinidis E. Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science*. 2015;7(2):351-67.
8. Dezza IC, Yu AJ, Cleeremans A, Alexander W. Learning the value of information and reward over time when solving exploration-exploitation problems. *Sci Rep*. 2017;7(1):16919.
9. Ebitz RB, Albarran E, Moore T. Exploration Disrupts Choice-Predictive Signals and Alters Dynamics in Prefrontal Cortex. *Neuron*. 2018;97(2):450-61 e9.
10. Wu CM, Schulz E, Speekenbrink M, Nelson JD, Meder B. Generalization guides human exploration in vast decision spaces. *Nat Hum Behav*. 2018;2(12):915-24.
11. Findling C, Skvortsova V, Dromnelle R, Palminteri S, Wyart V. Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nat Neurosci*. 2019;22(12):2066-77.
12. Gershman SJ. Deconstructing the human algorithms for exploration. *Cognition*. 2018;173:34-42.
13. Wilson RC, Geana A, White JM, Ludvig EA, Cohen JD. Humans use directed and random exploration to solve the explore-exploit dilemma. *J Exp Psychol Gen*. 2014;143(6):2074-81.
14. Addicott MA, Pearson JM, Sweitzer MM, Barack DL, Platt ML. A Primer on Foraging and the Explore/Exploit Trade-Off for Psychiatry Research. *Neuropsychopharmacology*. 2017;42(10):1931-9.
15. Pérez-Edgar K. Attention mechanisms in behavioral inhibition: Exploring and exploiting the environment. *Behavioral inhibition*. Cham: Springer; 2018. p. 237-61.
16. Smith R, Lane R, Nadel L, Moutoussis M. A computational neuroscience perspective on the change process in psychotherapy. In: Lane R, Nadel L, editors. *Neuroscience of Enduring Change: Implications for Psychotherapy*. Oxford University press; 2020.
17. Smith R, Moutoussis M, Bilek E. Simulating the computational mechanisms of cognitive and behavioral psychotherapeutic interventions: Insights from active inference. *Scientific Reports*. 2021;11:10128.

18. Morris LS, Baek K, Kundu P, Harrison NA, Frank MJ, Voon V. Biases in the Explore-Exploit Tradeoff in Addictions: The Role of Avoidance of Uncertainty. *Neuropsychopharmacology*. 2016;41(4):940-8.
19. Smith R, Schwartenbeck P, Stewart JL, Kuplicki R, Ekhtiari H, Investigators T, et al. Imprecise Action Selection in Substance Use Disorder: Evidence for Active Learning Impairments When Solving the Explore-exploit Dilemma. *Drug and Alcohol Dependence*. 2020;215:108208.
20. Addicott MA, Pearson JM, Wilson J, Platt ML, McClernon FJ. Smoking and the bandit: a preliminary study of smoker and nonsmoker differences in exploratory behavior measured with a multiarmed bandit task. *Exp Clin Psychopharmacol*. 2013;21(1):66-73.
21. Addicott MA, Pearson JM, Froeliger B, Platt ML, McClernon FJ. Smoking automaticity and tolerance moderate brain activation during explore-exploit behavior. *Psychiatry Res*. 2014;224(3):254-61.
22. Addicott MA, Baranger DA, Kozink RV, Smoski MJ, Dichter GS, McClernon FJ. Smoking withdrawal is associated with increases in brain activation during decision making and reward anticipation: a preliminary study. *Psychopharmacology (Berl)*. 2012;219(2):563-73.
23. Waltz JA, Wilson RC, Albrecht MA, Frank MJ, Gold JM. Differential Effects of Psychotic Illness on Directed and Random Exploration. *Computational psychiatry (Cambridge, Mass)*. 2020;4:18-39.
24. Aberg KC, Toren I, Paz R. A neural and behavioral tradeoff underlies exploratory decisions in normative anxiety. *bioRxiv*: <https://doi.org/10.1101/20200916300525>. 2021.
25. Blanco NJ, Otto AR, Maddox WT, Beevers CG, Love BC. The influence of depression symptoms on exploratory decision-making. *Cognition*. 2013;129(3):563-8.
26. Strauss GP, Frank MJ, Waltz JA, Kasanova Z, Herbener ES, Gold JM. Deficits in positive reinforcement learning and uncertainty-driven exploration are associated with distinct aspects of negative symptoms in schizophrenia. *Biol Psychiatry*. 2011;69(5):424-31.
27. Beck AT, Steer RA, Brown GK. Beck depression inventory-II. San Antonio. 1996;78(2):490-8.
28. Spielberger CD, Gorsuch RL, Lushene RE. Manual for the State-Trait Anxiety Inventory. Palo Alto, CA: Consulting Psychologists Press; 1970.
29. Persich MR, Smith R, Cloonan SA, Strong M, Killgore WDS. Emotional Intelligence Training as a Protective Factor for Mental Health During the COVID-19 Pandemic. *Depression and Anxiety*. 2021.
30. Toplak ME, West RF, Stanovich KE. Assessing miserly information processing: An expansion of the Cognitive Reflection Test. *Thinking & Reasoning*. 2014;20:147-68.
31. Toplak ME, West RF, Stanovich KE. Rational thinking and cognitive sophistication: Development, cognitive abilities, and thinking dispositions. *Developmental Psychology*. 2014;50:1037-48.
32. Stanovich KE, West RF, Toplak ME. The rationality quotient: Toward a test of rational thinking: MIT press; 2016.
33. Wechsler D. WASI -II: Wechsler abbreviated scale of intelligence - second edition. WASI. 2011.
34. Morey RD, Rouder JN. BayesFactor (Version 0.9.10-2)[Computer software]. 2015.
35. Rouder JN, Morey RD, Speckman PL, Province JM. Default Bayes Factors for ANOVA Designs. *Journal of Mathematical Psychology*. 2012;56:356-74.
36. Brañas-Garza P, Kujal P, Lenkei B. Cognitive reflection test: Whom, how, when. *Journal of Behavioral and Experimental Economics*. 2019;82:101455.
37. Easton C. Women and 'the philosophical personality': Evaluating whether gender differences in the cognitive reflection test have significance for explaining the gender gap in philosophy. *Synthese*. 2021;198(1):139-67.
38. Frederick S. Cognitive reflection and decision making. *Journal of Economic perspectives*. 2005;19(4):25-42.
39. Ring P, Neyse L, David-Barett T, Schmidt U. Gender Differences in Performance Predictions: Evidence from the Cognitive Reflection Test. *Front Psychol*. 2016;7:1680.

40. Zhang DC, Highhouse S, Rada TB. Explaining sex differences on the Cognitive Reflection Test. *Personality and Individual Differences*. 2016;101:425-7.
41. Juanchich M, Sirota M, Bonnefon JF. Anxiety-induced miscalculations, more than differential inhibition of intuition, explain the gender gap in cognitive reflection. *Journal of Behavioral Decision Making*. 2020;33(4):427-43.
42. Baron P, Perron LM. Sex differences in the Beck Depression inventory scores of adolescents. *J Youth Adolesc*. 1986;15(2):165-71.
43. Bryson SE, Pilon DJ. Sex differences in depression and the method of administering the Beck Depression Inventory. *J Clin Psychol*. 1984;40(2):529-34.
44. Hammen CL, Padesky CA. Sex differences in the expression of depressive responses on the Beck Depression Inventory. *J Abnorm Psychol*. 1977;86(6):609-14.
45. King DA, Buchwald AM. Sex differences in subclinical depression: administration of the Beck Depression Inventory in public and private disclosure situations. *Journal of personality and social psychology*. 1982;42(5):963-9.
46. Steer RA, Beck AT, Brown G. Sex differences on the revised Beck Depression Inventory for outpatients with affective disorders. *J Pers Assess*. 1989;53(4):693-702.
47. Lee MD, Wagenmakers EJ. *Bayesian cognitive modeling: A practical course*: Cambridge university press; 2014.
48. Carleton RN, Mulvogue MK, Thibodeau MA, McCabe RE, Antony MM, Asmundson GJ. Increasingly certain about uncertainty: Intolerance of uncertainty across anxiety and depression. *J Anxiety Disord*. 2012;26(3):468-79.
49. Krueger PM, Wilson RC, Cohen JD. Strategies for exploration in the domain of losses. *Judgment and Decision Making*, Vol 12, No 2, March 2017, pp 2017;12(2):104–17.
50. Zajkowski WK, Kossut M, Wilson RC. A causal role for right frontopolar cortex in directed, but not random, exploration. *Elife*. 2017;6.
51. Benoit RG, Gilbert SJ, Burgess PW. A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *J Neurosci*. 2011;31(18):6771-9.
52. Mansouri FA, Koechlin E, Rosa MGP, Buckley MJ. Managing competing goals - a key role for the frontopolar cortex. *Nat Rev Neurosci*. 2017;18(11):645-57.
53. Somerville LH, Sasse SF, Garrad MC, Drysdale AT, Abi Akar N, Insel C, et al. Charting the expansion of strategic exploratory behavior during adolescence. *J Exp Psychol Gen*. 2017;146(2):155-64.

		Total															
		STAI State	STAI Trait	BDI	Dir. Exp.	Info. Bon. H1	Info. Bon. H6	Rand. Exp. [2 2]	Rand. Exp. [1 3]	Dec. Noise H1 [1 3]	Dec. Noise H1 [2 2]	Dec. Noise H6 [1 3]	Dec. Noise H6 [2 2]	Choice Acc.	First Choice RT	Dev from opt. H1	Dev from opt. H6
Total	CART	-0.01	-0.02	-0.1	0.14	0.02	0.16	-0.01	0.06	-0.32	-0.39	-0.25	-0.38	0.44	0.24	-0.25	-0.08
		BF: 0.12	BF: 0.12	BF: 0.97	BF: 5.17	BF: 0.13	BF: 20.18	BF: 0.12	BF: 0.21	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF > 100	BF: 0.47
	CART	-0.03	0.04	-0.02	0.14	0.04	0.17	-0.01	0.04	-0.31	-0.37	-0.25	-0.35	0.4	0.23	-0.27	-0.17
Male	AOMTS	-0.13	-0.1	-0.13	0.14	-0.11	0.06	0.05	0.06	-0.16	-0.21	-0.09	-0.15	0.2	0.27	-0.09	-0.08
		BF: 4.38	BF: 1.08	BF: 4.48	BF: 6.76	BF: 1.23	BF: 0.25	BF: 0.19	BF: 0.25	BF: 20.83	BF > 100	BF: 0.7	BF: 13.28	BF > 100	BF > 100	BF: 0.58	BF: 0.45
	CART	0.15	0.09	0.02	0.3	-0.08	0.27	-0.12	-0.07	-0.29	-0.44	-0.3	-0.46	0.47	0.28	-0.25	0.08
Female	CRT	0.02	0.1	0.02	0.24	-0.02	0.24	-0.04	0.04	-0.37	-0.44	-0.27	-0.41	0.43	0.33	-0.23	0.06
		BF: 0.23	BF: 0.37	BF: 0.22	BF: 4.21	BF: 0.23	BF: 4.05	BF: 0.24	BF: 0.24	BF > 100	BF > 100	BF: 8.88	BF > 100	BF > 100	BF: 68.95	BF: 3.51	BF: 0.26
	AOMTS	-0.03	-0.09	-0.12	0.26	-0.19	0.14	0.07	0.04	-0.23	-0.39	-0.15	-0.27	0.32	0.36	-0.09	-0.08
Total	CART	-0.07	-0.01	-0.12	0.05	0.03	0.08	0.02	0.09	-0.3	-0.37	-0.21	-0.34	0.41	0.15	-0.22	-0.13
		BF: 0.3	BF: 0.14	BF: 1.31	BF: 0.2	BF: 0.16	BF: 0.34	BF: 0.14	BF: 0.45	BF > 100	BF > 100	BF: 89.83	BF > 100	BF > 100	BF: 3.89	BF > 100	BF: 1.8
	CRT	-0.05	0.06	0	0.08	0.05	0.12	0	0.04	-0.27	-0.35	-0.22	-0.33	0.37	0.12	-0.26	-0.24
Male	AOMTS	-0.17	-0.08	-0.13	0.08	-0.09	0.01	0.04	0.06	-0.12	-0.13	-0.06	-0.09	0.13	0.19	-0.07	-0.08
		BF: 12.53	BF: 0.38	BF: 1.42	BF: 0.36	BF: 0.46	BF: 0.14	BF: 0.16	BF: 0.25	BF: 1.13	BF: 1.52	BF: 0.22	BF: 0.47	BF: 1.98	BF: 30.97	BF: 0.25	BF: 0.33
	STAI State	-0.15	0.14	-0.04	-0.05	-0.04	0.07	0.04	0.02	0	-0.09	0.09	-0.06	0.03	-0.1	0.06	0.11
Female	STAI Trait	-0.11	0.1	-0.03	-0.01	0.03	0.03	-0.01	0.06	-0.01	0.06	-0.02	-0.03	-0.03	0.03	0.03	0.02
		BF: 1.19	BF: 0.94	BF: 0.14	BF: 0.12	BF: 0.14	BF: 0.14	BF: 0.12	BF: 0.24	BF: 0.13	BF: 0.14	BF: 0.14	BF: 0.14	BF: 0.14	BF: 0.14	BF: 0.14	BF: 0.12
	BDI	-0.16	0.11	-0.08	-0.01	0.08	0.06	0.05	0.12	0.03	-0.11	-0.06	0.08	0.04			
Total	STAI State	-0.07	0.05	-0.04	0.02	0.09	0	-0.09	0.09	-0.09	0.09	-0.06	0.03	-0.1	0.06	0.11	
		BF: 0.27	BF: 0.24	BF: 0.23	BF: 0.22	BF: 0.34	BF: 0.22	BF: 0.32	BF: 0.32	BF: 0.32	BF: 0.26	BF: 0.23	BF: 0.38	BF: 0.26	BF: 0.44		
	STAI Trait	0	0.05	0.04	0.06	0.07	-0.09	-0.16	-0.01	-0.08	0.11	0.03	0.03	0.1			
Male	BDI	-0.07	0.06	-0.04	0.05	0.19	-0.04	-0.05	0.14	-0.01	0.01	-0.11	0.11	0.2			
		BF: 0.29	BF: 0.26	BF: 0.23	BF: 0.24	BF: 1.5	BF: 0.23	BF: 0.25	BF: 0.67	BF: 0.22	BF: 0.22	BF: 0.43	BF: 0.39	BF: 1.96			
	STAI State	-0.18	0.17	-0.04	-0.07	-0.09	0.08	0.09	0	0.03	-0.12	-0.1	0.07	-0.01			
Female	STAI Trait	-0.14	0.13	-0.04	-0.04	0.02	0.05	0.03	0.07	0	-0.07	-0.03	0.01	-0.02			
		BF: 2.36	BF: 1.77	BF: 0.17	BF: 0.17	BF: 0.15	BF: 0.18	BF: 0.16	BF: 0.25	BF: 0.13	BF: 0.26	BF: 0.15	BF: 0.14	BF: 0.14			
	BDI	-0.18	0.14	-0.08	-0.03	0.05	0.07	0.07	0.11	0.05	-0.14	-0.01	0.06	-0.01			
		BF: 18.59	BF: 2.23	BF: 0.33	BF: 0.15	BF: 0.18	BF: 0.27	BF: 0.3	BF: 0.74	BF: 0.19	BF: 2.71	BF: 0.14	BF: 0.24	BF: 0.14			

Figure S1. Post-hoc correlations (and associated BFs) between reflectiveness measures, symptoms, and Horizon Task measures. This is analogous to **Figure 5** in the main text, but examines males and females separately – and suggests that observed relationships between directed exploration and symptoms were driven by females while relationships between directed exploration and reflectiveness measures were driven by males. Asterisks indicate BFs greater than 3, 10, 30, and 100, per conventional cutoffs for levels of evidence from moderate to extremely strong.