UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title

Towards computational phenotypes of internalizing psychopathology: An investigation of decision-making and learning algorithms

Permalink

https://escholarship.org/uc/item/4pv2t3d2

Author

Senta, Jennifer D.

Publication Date

2024

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Towards computational phenotypes of internalizing psychopathology: An investigation of decision-making and learning algorithms

By

Jennifer D. Senta

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Psychology

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Sonia J. Bishop, Co-Chair
Professor Anne G.E. Collins, Co-Chair
Professor Sheri Johnson
Professor Steven Piantadosi

Summer 2024

© Copyright 2024
Jennifer D. Senta
All rights reserved

Abstract

Towards computational phenotypes of internalizing psychopathology: An investigation of decision-making and learning algorithms

by

Jennifer D. Senta

Doctor of Philosophy in Psychology

University of California, Berkeley

Professor Sonia J. Bishop, Co-Chair
Professor Anne G.E. Collins, Co-Chair

Computational models of cognitive processes have provided deep insights into specific mechanisms of human learning and decision-making. A natural corollary of research into the typical functioning of such mechanisms is to investigate how mental disorders might cause impairment of the underlying algorithmic processes. This line of research, frequently referred to as computational psychiatry, seeks to contribute to more personalized diagnoses and treatments of mental disorders by characterizing behavioral and computational profiles of psychopathology. In the current dissertation we focus on the role of latent dimensions of internalizing psychopathology in computational processes of learning and decision-making to provide insight into dissociable dimensional phenotypes. Chapter 1 of the dissertation provides a primer to computational characterizations of learning and decision-making, including recent and foundational research results in computational psychiatry which motivate the subsequent original research questions. In Chapter 2, we present an investigation of effortful decision-making in reward pursuit and threat avoidance. We use a detailed dimensional profile of depressive and anxious symptoms to characterize unique mechanistic impairments across frequently comorbid symptoms, with individual symptoms of depression relating to specific differences in effortful reward pursuit processes, while a dimensional characterization of anxiety relates to multiple mechanistic differences in effortful threat avoidance. In Chapter 3, we extend the space of psychopathological inquiry to include subclinical levels of mania in an investigation of decision-making under ambiguity in reward pursuit, loss avoidance, and threat avoidance paradigms. Here, individual differences in loss avoidance decisions under ambiguity show opposing effects of anhedonic depression and physiological anxiety in risk sensitivity, with depression showing increased risk sensitivity and anxiety showing reduced nonlinear risk valuation. Ambiguous decisions in threat avoidance reveal a significant dissociation between anhedonic depression and hypomania, with higher levels of hypomania associated with significantly less risk sensitivity in threat avoidance than anhedonic depression. We also find suggestive, although not concrete, evidence that sensitivity to threat magnitude may relate to comorbid variance between anxiety and depression, and we recommend additional work in this area of investigation. In Chapter 4, we present a detailed

investigation of dimensional characterizations of anxiety in reinforcement learning impairment, using an experimental paradigm which differentially manipulates the relative load of working memory versus learning systems across task. Here we find that physiological, but not cognitive, anxiety is related to significant learning impairments which are contributed to both by slower learning processes and by increased rate of working memory decay. These results may have implications for a wide variety of extant reinforcement learning modeling studies in computational psychiatry by emphasizing the importance of an algorithmic account of supporting working memory processes in such investigations. Chapter 5 provides a discussion of the results in the framework of dissociating unique profiles of dimensions of internalizing psychopathology. As additional insights into the relationships of latent dimensions of psychopathology with algorithms of cognition are gained, the field as a whole is enabled to approach the idea of replicable computational phenotypes of psychopathology. The aim of this dissertation is not to operationalize such phenotypes at present, but rather to contribute to the growing body of knowledge within the field to enable and inspire more detailed investigations and replications of the findings contained herein.

Dedication

To having measured out our lives with coffee spoons.

Table of contents

Dedication i

Table of contents ii

List of figures iii

List of tables vi

Acknowledgments vii

Introduction: Seeking to understand the unhealthy mind viii

Chapter 1: Learning and decision-making: a brief orientation from the perspective of computational cognition 1

- Introduction 1
- Value-Based Decision-Making 1
- Reinforcement learning 6
- Final thoughts 9

Chapter 2: Latent subdimensions of anxiety and depression differentially influence exertion of effort in pursuit of reward versus avoidance of threat 11

- Abstract 11
- Introduction 11
- Materials & Methods 13
- Results 24
- Discussion 35
- Chapter 2: Supplemental materials 40

Chapter 3: Ambiguous decision-making and valence-specific framing in dimensional characterization of anxiety, depression, and mania 49

- Abstract 49
- Introduction 49
- Methods 52
- Results 73
- Limitations 83
- Discussion 85
- Chapter 3: Supplemental materials 88

Chapter 4: Reinforcement learning deficits in physiological anxiety are contributed to by both reduced learning rate and increased rate of working memory decay 99

- Abstract 99
- Introduction 99
- Methods 101
- Results 110
- Discussion 118
- Chapter 4: Supplemental materials 121

Chapter 5: Discussion 133

References 136

List of figures

Figure 1 Algorithms of cognition ix
Figure 2.1 Illustrative sigmoidal interpolation in electrical stimulation calibration. 17
Figure 2.2 Illustration of group-level trait offset to posterior parameter distribution in hierarchical Bayesian framework. 20
Figure 2.3 Bifactor analysis of item response data. 26
Figure 2.4 Correlation of questionnaire subscale scores with participant-level bifactor general and specific factor scores in exploratory (online) data participants and in-lab behavioral task participants. 28
Figure 2.5 Effortful decision-making behavioral task. 29
Figure 2.6 Behavioral choice results from participant data and winning computational model. 30
Figure 2.7: Model fit comparison WAIC. 32
Figure 2.8 Effect of factor scores on decision-making parameters in reward pursuit for winning computational model of choice. 33
Figure 2.9 Effect of factor scores on decision-making parameters in reward pursuit for winning computational model of choice. 35
Figure S.2.1: Participant parameter posterior distributions for winning model 40
Figure S.2.2: Parameter recovery for winning model 42
Figure S.2.3: Model recovery 44
Figure S.2.4: Full list of item loadings for specific factors 45
Figure 3.1 Bifactor analysis results. 59
Figure 3.2 Decision-making under ambiguity task. 60
Figure 3.3 Schematic of decision-making under ambiguity. 62
Figure 3.4 Illustration of nonlinear subjective valuation of outcome magnitude. 66
Figure 3.5 Influence of optimistic/pessimistic bias parameter α when 30 tokens are hidden. 71
Figure 3.6 Ambiguity avoidance measured by percentage of choices for unambiguous urn. 73
Figure 3.7 s-ILDAA slope. 74
Figure 3.8 Model comparison. 75
Figure 3.9 Simulated vs. actual data. 76
Figure 3.10 Relationship of categorical ambiguity aversion and sensitivity to level of missing information. 77

Figure 3.11 Participant risk-preference parameters λ for reward pursuit, loss avoidance, and threat avoidance decision-making under ambiguity. 77
Figure 3.12 Participant optimistic/pessimistic inferential bias parameters α for reward pursuit, loss avoidance, and threat avoidance decision-making under ambiguity. 79
Figure 3.13 Relationship of hypomania symptom levels with ambiguity and risk sensitivity. 80
Figure 3.14 Difference between effects of hypomania versus anhedonic depression on probability sensitivity in threat avoidance. 81
Figure 3.15 Trait anxiety and internalizing factor scores vs. sensitivity to magnitude difference in threat avoidance. 82
Figure 3.16 Physiological anxiety and anhedonic depression vs. measures of risk sensitivity in loss avoidance. 83
Figure S.3.1a: Subscale distributions for final in-lab behavioral participants (n=71) 88
Figure S.3.1b: Within-group subscale score correlations for final in-lab behavioral participants (n=71) 89
Figure S.3.2a: Subscale distributions for online participants (exploratory bifactor analysis; n=470) 90
Figure S.3.2b: Within-group subscale score correlations for online participants (exploratory bifactor analysis; n=470) 91
Figure S.3.3a: Subscale distributions for combined in-lab participants (confirmatory bifactor analysis; n=179) 92
Figure S.3.3b: Within-group subscale score correlations for combined in-lab participants (confirmatory bifactor analysis; n=179) 93
Figure S.3.4: Spearman correlations between posterior mean parameter values and scores on factors of the winning bifactor model. 94
Figure S.3.5: Parameter recovery for winning model (Model FI_EUMP_OP_c_A) in each valence paradigm 96
Figure 4.1 RLWM task 103
Figure 4.2 Model fit comparison using Akaike Information Criterion (AIC) 109
Figure 4.3 Differences between learning and testing performance by set size. 111
Figure 4.4 Participant behavioral data versus data simulated from winning model (Model #5) 112
Figure 4.5 Median split of learning and testing performance on MASQ AA and PSWQ 115
Figure 4.6 Relationship of parameters of learning and working memory with individual scores on MASQ AA and PSWQ 117
Figure S.4.1: Distributions of questionnaire metrics in final participant group (n=164) 121
Figure S.4.2: Parameter value distributions: Model #5 122
Figure S.4.3: Parameter recovery analyses: Model #5 123
Figure S.4.4: Model recovery analysis 124

Figure S.4.5: Learning phase reaction times by set size as a function of anxiety scores 125
Figure S.4.6: Testing phase reaction times by set size as a function of anxiety scores 126
Figure S.4.7: Exploratory parameter testing Model #5 127
Figure S.4.8: Exclusion statistics 128
Figure S.4.9: Correlation of PSWQ and MASQ AA subscale scores with overall learning performance 129
Figure S.4.10: Correlation of PSWQ and MASQ AA subscale scores with overall testing performance 131

List of tables

Table 3.1 Model key for decision-making under ambiguity analysis. 71
Table 4.1 Computational models of reinforcement learning evaluated. 109

Acknowledgments

I would like to express my deep gratitude to my advisor, Dr. Sonia Bishop, for giving me the opportunity to study cognitive neuroscience in her lab and for sharing her brilliant research mind with me. The journey to a PhD is undoubtedly complex, and I am incredibly thankful for your encouragement, openness, and generosity in helping to shape this journey.

I am also deeply grateful to my honorary second advisor and dissertation co-Chair, Dr. Anne Collins, who has unfailingly been a sounding board and a wonderful teacher and mentor. Thank you for all of your support!

In addition, I am incredibly grateful to Dr. Peter Dayan for his generous mentorship and superhuman availability.

My deepest thanks to Dr. Sheri Johnson and Dr. Steve Piantadosi; I deeply admire you both, and your roles on my dissertation committee have foundationally shaped my research.

Thank you also to Dr. Joanne Zinger for giving me the first opportunity to trade computational models of risk for computational models of cognition.

Additionally, I would like to thank the Seattle-based actuarial triumvirate of Nick Collier, Mark Olleman, and Dan Wade (otherwise known as NickMarkDan), plus Matt Larrabee (from another office but a kindred spirit) and Karen Steffen, for being formative role models of great intellect with great care and deep morality. I'm thankful for my time as an actuary for helping me develop persistence and grit in pursuit of the actuarial credentials, but I am even more thankful to have worked with and looked up to you all. I also thank Norva Osborn for inspiration on how to be unfailingly lovely, and for bringing beauty and a continual supply of your fresh garden roses to my life.

I am forever grateful for my sisters and built-in besties Kristina Rose and Kaela Anne, who would certainly be two of my favorite people even if I wasn't related to them, and for our beautiful mom who fills our lives with love and care and delight. Thank you for your love and inspiration.

I am deeply thankful to my brilliant dad for inspiring me to love math and science from childhood, from building circuits in the garage to ordering photos from NASA to clarifying all my questions regarding electrons. You are my role model for intelligence and scientific curiosity with a good dash of mischievous twinkle thrown in.

To my treasured friend Dr. Milena Rmus; we've had some delightful adventures all around the world, but the most delightful of them all is simply having you be a part of my journey.

To the irrepressible Dr. Đorđe Gluhović, for always bringing laughter and light and support; your excellence is an inspiration and your presence in my life is a gift.

And finally, to Carl, who has brought me a large quantity of joy and only a small quantity of not-joy with his unexpected emotional depths and tiny fluffy murder mittens.

Introduction

Seeking to understand the unhealthy mind

The human brain is a computational marvel; a decision-making engine with trillions of neurochemical interaction points communicating locally and globally in highly sophisticated parallelized fashion. Our experience of the world and of ourselves is integrally tied to the healthy operation of this biological machine. Even small amounts of damage to some areas of the brain can cause catastrophic changes in behavior and experience. Imbalances in neurochemistry or neural communication can cause a variety of unwanted effects which alter our moods, preferences, and abilities. The sheer complexity of the human brain yields a combinatorial explosion of possible errors. To adapt liberally from Tolstoy (Tolstoy, 1980), perhaps every unhealthy mind is literally unhealthy in its own way.

Such realizations have led to the increasing popularity of the concept of personalized medicine and psychiatry, in which diagnoses and treatments for psychiatric pathology might be more highly individualized and consequently more effective for any given individual. Within the field of psychiatry, the diagnosis and treatment of psychopathology has traditionally been defined by two major sources of classification: the Diagnostic and Statistical Manual of Mental Disorders (DSM) published by the American Psychiatric Association and currently in its 5th edition (DSM-V) (American Psychiatric Association, 2013) and the International Classification of Diseases (ICD) manual published by the World Health Organization, currently in its 11th revision (ICD-11) (World Health Organization, 2019).

These manuals define psychiatric diagnoses as functions of symptom presentation, either self-reported or clinically observed; as such, they typically delineate specific thresholds for binarized diagnosis. One contemporary concern for the field of mental health has been that such diagnoses may be overly binary; while one individual may be diagnosed with Major Depressive Disorder (MDD) based on current symptoms, another individual with all save one of the same symptoms may not meet the threshold for medical diagnosis and therefore not be eligible for the same level of treatment under current medical care guidelines.

In addition to potentially rendering the continuum of psychopathology into overly binary diagnostic categories, the current system also admits of high levels of comorbidity between different diagnoses. According to summary statistics across numerous global studies, individuals diagnosed with depression in a given 12 month period have a greater than 40% chance of being diagnosed with an anxiety disorder in the same 12 months, and comorbidity between lifetime depression and anxiety disorders is also greater than 40% (Kessler et al., 2015). This diagnostic comorbidity presents challenges for research into the cognitive and neural mechanisms involved in each disorder by confounding the ability to characterize effects as unique to one or the other of the co-occurring disorders.

Another facet of the extant diagnostic nosology which presents difficulties for research is the high level of heterogeneity possible within each diagnostic category. Within a single diagnosis such as MDD, differing combinations of underlying symptoms may cause substantial variation in the

individual experience of the diagnosis. For example, depression may present with either hyper- or hyposomnia as a symptom. It is reasonable to posit that depression characterized by too much sleep might relate at least somewhat differently to neural mechanisms than depression characterized by too little sleep, but research relying on diagnostic categorizations cannot make this distinction between symptom level variation between individuals or between time points for the same individual.

In light of these concerns, research into the relationship between psychopathology and cognition has begun to transition to use of a dimensional framework which considers continuums of symptom-level dimensions spanning clinical and subclinical levels of pathology. A focus on individual differences in symptom dominance and combination, as well as a variety of mathematical techniques to parse common from specific variance within and across currently distinct diagnoses, has helped to identify more targeted directions for study in the cognitive systems affected by such pathology.

Within the field of cognitive science, human behavior and its underlying neural mechanisms are increasingly being studied via computational models of putative cognitive processes. Such models propose algorithmic structures for complex behaviors like decision-making or learning and aim to provide a more nuanced view of the underlying cognitive mechanisms than classic behavioral analysis can yield. While these analyses have the potential to provide cognitive scientists, neuroscientists and artificial intelligence researchers with enormous amounts of information about the algorithmic nature of the human brain, they may also be used as a tool to explore the varying ways in which cognition may be computationally abnormal. The growing research field of computational psychiatry aims to apply such cognitive models to the investigation of psychiatric disorders.

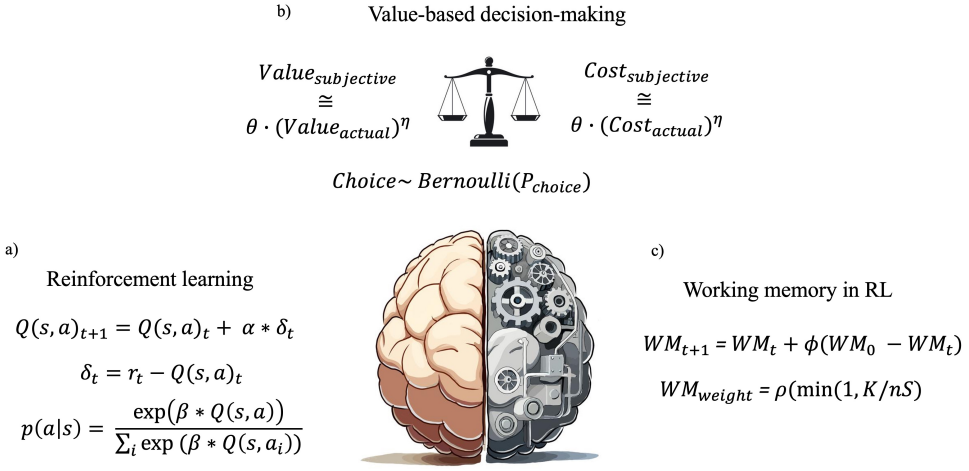


Figure 1 Algorithms of cognition Computational phenotyping explores the mathematical relationships between algorithms of cognition and dimensions of psychopathology. a) Sample algorithms of RL illustrate the updating of outcome values given states and actions according to an individual's learning rate and the extent to which the actual outcome differs from expectation. b) Sample algorithms of value-based decision-making reflect linear and first-order nonlinear subjective valuation of both costs and outcomes, with resulting probability of choice distributed according to a Bernoulli distribution for each trial. c) Sample algorithms for the role of WM in RL reflect rapid forgetting of information stored in WM and limitations on its role in RL processes based upon individual differences in both WM capacity and confidence.
Original center artwork courtesy of Dr. Milena Rmus.

Many associations between psychopathology and abnormal cognition are well-known by the public in broad strokes. For example, thought disorders such as schizophrenia are widely associated with abnormal cognition, which may involve delusions or hallucinations in addition to more subtle symptoms such as increased risk-taking in decision-making (Andreasen and Flaum, 1991; Pedersen et al., 2017) and impaired learning (Hyojin Kim et al., 2007; Paulus et al., 2003). As more sophisticated cognitive models and study methodologies emerge, we can ask more detailed questions about the precise cognitive mechanisms which are altered in individuals with high levels psychopathology. For example, computational models of the interaction between reinforcement learning and working memory systems in the brain were recently able to reveal that a long-held apparent deficit in reinforcement learning among individuals with schizophrenia is attributable to working memory systems and not the reinforcement learning process itself (Collins et al., 2014). Since the neural instantiation of working memory involves brain circuits which are distinct from those involved in reinforcement learning, such insights yield potential new ways in which to press the question of specificity even further. This interactive model of reinforcement learning and working memory systems is the foundation for original research presented in Chapter 4 of this dissertation, in which we investigate the relationship between different subdimensions of anxiety and problems with reinforcement learning or the elements of working memory that support it. Further refining our conceptualization of subdimensions of internalizing psychopathology, the original work presented in Chapter 2 of this dissertation examines the individual role of distinct symptoms of depression in effortful decision-making, finding that although different symptom-level dimensions may result in the same behavior, they may act through distinct cognitive mechanisms.

As the level of collective knowledge with respect to the effects of latent dimensions of psychopathology on algorithms of cognition grows, and as the areas of new inquiry become increasingly nuanced, we are enabled to make progress as a field towards the notion of a computational psychiatric phenotype. A computational psychiatric phenotype explicates a predictive quantitative relationship between psychiatric diagnoses or symptoms and computational models of measurable behavioral or neural outcomes. The idea is that such phenotyping could eventually provide easily evaluated personal and behavioral metrics which would grant immediate rich insight into individual specific mechanistic impairment and contribute to more highly personalized diagnosis and treatment.

Such computational phenotypes are dependent on the efficacy of the selected computational model in describing and predicting behavior, the reliability and reproducibility of the experimental paradigm used, and the reliability and specificity of the measure of trait psychiatric symptomatology used. In the present work I aim to use computational models of cognitive processes to help clarify how latent dimensions of internalizing psychopathology, particularly in relation to dimensions of anxiety and depression, influence mechanisms of learning and decision-making. In particular, it is the goal of this work not to operationalize any specific computational phenotype, but to add to and enhance the body of scientific research necessary to complete the definition of any such phenotype in the future. The research presented here focuses on model categories which have been widely studied and well-validated through a combination of behavioral sciences, economics, and artificial intelligence research.

Chapter 1 of the dissertation provides further background and a broad orientation to the original research subsequently presented. The purpose of this chapter is to provide a relatively targeted overview of computational cognitive decision-making and learning relevant to the research presented in later chapters, along with enough broad background to empower clear understanding of the models presented throughout and their conclusions. Along the way, I will present key foundational and recent findings with respect to the relevant computational psychiatry research which motivated the research questions addressed by further chapters.

In Chapter 2, I present the results of original work investigating the relationship of subdimensions of anxiety and depression with effortful decision-making. Here we use an in-lab behavioral task assessing willingness to exert physical effort to pursue monetary reward and to avoid physical threat of electrical stimulation to explore individual differences both within and across valences. Hierarchical modeling of decision-making incorporating latent trait dimensions reveals a mechanistic specificity behind the similar behavioral influences of two key symptoms of depression. The role of cognitive and physiological anxiety in effortful decision-making is also investigated.

Chapter 3 presents original research into the relationship of dimensional characterizations of anxiety, depression, and hypomania with decision-making under ambiguity. Here we use an in-lab experimental paradigm manipulating levels of masked information in a probabilistic choice task in which magnitude also varies. Each participant completed 3 variants of the task: decisions made in pursuit of reward; decisions made to avoid loss of reward; and decisions made to avoid physical threat. We use hierarchical Bayesian modeling of computational mechanisms and individual differences in dimensional trait scores to assess the influence of psychopathology on ambiguous decision-making and valence framing effects.

In Chapter 4, I present an investigation of the interplay of reinforcement learning and working memory systems during learning tasks. Using an experimental paradigm designed to differentially manipulate the relative load of working memory and reinforcement learning, we explore the relationship of dimensional characterizations of anxiety with learning and its supporting processes.

Keeping in mind the purpose of each model (Actuarial Standards Board, 2019) and with clear explication of methods and approaches, this dissertation aspires to contribute to the scientific understanding of the mind via the computational modeling of cognition and psychopathology.

Chapter 1

Learning and decision-making: a brief orientation from the perspective of computational cognition

Introduction

There are many ways to computationally characterize human cognitive processes like decision-making or learning; in fact, computational models of cognitive science can be as diverse as the scientists who create them. Although the space of possible models is wide, over time successful research in economics, econometrics, mathematical psychology, computer science and cognitive science has coalesced on some useful algorithms which have been well-validated and provide reliable insight into latent cognitive processes. Some have been shown to mirror specific neural processes, while others are highly predictive with respect to behavior.

The purpose of this chapter is to provide a relatively targeted overview of computational cognitive decision-making and learning relevant to the original research presented in later chapters, along with enough broad conceptual background to empower clear understanding of the models presented throughout later chapters and their conclusions. Along the way, I will present key foundational and recent findings in computational psychiatry research that motivate the research questions of focus throughout the dissertation.

Value-Based Decision-Making

Our daily lives are characterized by a constant string of decisions in which we choose between possible actions available to us at any given moment. The notion of value-based decision-making stems from the idea that action selection involves a tradeoff between the costs and benefits associated with possible actions. Generally, actions are selected to maximize net reward for the decision-maker; however, the subjective value of actions, and consequently the actions selected, may vary widely across individuals. Various psychiatric disorders including depression, mania, and schizophrenia have been associated with aberrant decision-making behavior (Clark et al., 2009; Mukherjee et al., 2020; Mukherjee and Kable, 2014; Sevy et al., 2007).

In order to properly investigate the computational underpinnings of value comparison, we must decompose the process into its distinct algorithmic components, each of which may vary across individuals. In particular, computationally determining the subjective value of an action will involve at minimum 1) subjective valuation of outcome benefits, 2) subjective valuation of costs, and 3) the subjective process of comparing costs to benefits. Each component process may involve distinct neural correlates, may vary in degree across

individuals, and may be differentially compromised across latent dimensions of psychiatric disorder.

Subjective reward valuation

The idea that individuals value objectively identical rewards differently is intuitive; as a basic example, a rich man may not value a reward of $100 as highly as a poor man (or graduate student). In fact, a complex interplay of circumstances, values, and preferences will contribute to individual differences in subjective reward valuation. Such concepts may contribute both to straightforward linear differences in reward valuation across magnitudes as described above, as well as nonlinear dynamics resulting from differences in the valuations of relative reward increases (Buechel and Morewedge, n.d.; Hsee et al., 1999; Kahneman and Tversky, 1979; Menger, n.d.; Tversy and Kahneman, 1974). Differences in the valuation of relative reward increases are also conceptually intuitive; we might illustrate this with the idea that a gain of $100 on a starting balance of $0 will likely be more highly valued than the same gain of $100 on a starting balance of $1 million, even by the same individual.

To mathematically operationalize these subjective value considerations, we may describe an individual's linear and nonlinear adjustments to actual value as:

$$V_s = \theta \cdot (V_a)^{\lambda} \quad 1$$

Here both θ and λ represent parameterizations of individual reward valuation differences: θ describes the linear difference between actual value (V_a) and subjective value (V_s) across magnitudes, while λ describes the nonlinear adjustment in an individual's subjective value.

The neural circuitry underlying subjective valuation of reward has been extensively studied across animals and humans. While the circuitry underlying subjective outcome value is complex, it is generally agreed that the ventral striatum (VS) and the orbitofrontal cortex (OFC), particularly its ventromedial aspect (vmOFC), are strongly involved in reward value representations (Kable and Glimcher, 2009; Padoa-Schioppa and Assad, 2006; Peters and Büchel, 2010). Meanwhile, dopaminergic neurons in the ventral tegmental area (VTA) of the brain project bidirectionally to VS, and play a key role in assessing incentive salience (Berridge and Robinson, 1998) and errors in our predictions of outcomes (Montague, Dayan, and Sejnowski 1996; see Diederen and Fletcher 2021 for a more recent review). The amygdala, an area of the basal ganglia central to emotion processing and threat detection (among a multitude of other functions) also strongly projects to an area of the VS known as the nucleus accumbens (Nacc), and is also likely involved in the processing and learning of rewarding outcome values (Fudge et al., 2002; Wassum, 2022).

The system complexity allows for myriad issues to arise during reward valuation and processing, and various mechanistic impairments may be differentially associated with various symptoms of psychopathology. Individuals diagnosed with depression have impaired reward processing behavior relative to controls (Halahakoon et al., 2020) and pursue rewarding outcomes less when more effort is required (Treadway et al., 2012).

Importantly, this reward devaluation may be linked to specific symptoms of depression rather than diagnosed depression itself, such as anhedonia (a reduction in or absence of pleasure) (Geaney et al., 2015; Treadway et al., 2009; Treadway and Zald, 2011) and apathy (a reduction in motivation) (American Psychiatric Association., 2013; Bonnelle et al., 2015a, 2015b). Recent work has highlighted the potentially unique contributions of these symptom dimensions to behaviorally inferred reward valuation, and has even provided further nuance by dividing anhedonia into dimensions itself, separating the experience of pleasure from the anticipation of pleasure (Treadway and Salamone 2022; Geaney, Treadway, and Smillie 2015). In Chapter 2 of the dissertation, we dissociate computational mechanisms of decision-making in effortful reward pursuit that are impaired by symptoms related to depression versus anxiety, finding evidence that indeed distinct mechanisms are impaired by distinct latent psychopathological dimensions, and providing a framework for future neuroimaging studies to explore in more detail.

Despite the high level of comorbidity between depression and anxiety, few studies have explored the relationship of anxiety with reward valuation and pursuit. Many previous analyses have studied the covariance structure across self-report questionnaire measures of anxiety and depression using bifactor analysis, a method of factor analysis which mathematically separates shared versus distinct dimensions of variance, and have found them to be well-fit by a tripartite model which estimates comorbidity as a separate dimension (Clark et al., 1994; Clark and Watson, 1991; Gagne et al., 2020). This hints at the importance of dimensional characterization of highly comorbid disorders in computational research; the common variance between the disorders may itself be a dimension which relates to various cognitive impairment (as in, for example, Gagne et al. 2020).

Avoiding negative outcomes

In research, rewarding outcomes are often represented via monetary gain or the receipt of a pleasurable stimulus (for example, highly desirable juice or sucrose for animals, or preferred foods for humans). However, humans (and animals) also make decisions in order to avoid negative outcomes. The subjective value of outcomes such as the avoidance of monetary loss or the escape from physical threat may be algorithmically represented in a manner analogous to reward. Importantly, individuals may utilize different subjective valuation schemas across outcome valences; for example, a person may assign a very high value to escaping physical threat but not care much about receiving money, resulting in different subjective value characterizations for reward versus threat in that individual.

In particular, threat processing relies heavily on amygdala responses to sensory information, and likely comprises both reactive and evaluative circuitry (Faul et al., 2020; Levy and Schiller, 2021). While sensitivity and response to threat are healthy behaviors, some forms of psychopathology such as anxiety and post-traumatic stress disorder (PTSD) may involve maladaptive responses to threat (Ball and Gunaydin, 2022; Bishop and Gagne, 2018a; Etkin et al., 2004; Robinson et al., 2013, 2019). Consideration of individual differences in subjective value across valence paradigms is therefore likely to show valence-specific relationships to dimensions of psychopathology. In Chapters 2 and 3 of the dissertation,

we use experimental paradigms across multiple valences to test such differences in decision-making.

Decision-making under risk

Econometrics and behavioral economics have historically researched value-based decision-making using both monetary reward and loss. In such research, outcomes are typically awarded probabilistically so that the participant must weigh both the objective magnitude and the likelihood of the outcome in determining its subjective value. Decisions in which the outcomes are uncertain are referred to as "risky" decisions and represent the realistic scenario in which we must make decisions without a guaranteed outcome. Individuals may have subjective values of outcome probabilities (e.g., they may estimate them optimistically or pessimistically), and when the subjective value of outcome magnitude and probability are both taken into account we refer to the "expected value" of an outcome or action. To incorporate a probabilistic outcome into the algorithmic representation of subjective value, we may use economic prospect theory (Kahneman and Tversky, 1979; Tversy and Kahneman, 1974) to write:

$$EV = \theta \cdot {V_a}^{\lambda} \cdot P^{\omega} \qquad 2$$

Evaluated in this way, the parameter ω represents individual subjective valuation of outcome probability, and the parameter λ may be interpreted as the subjective risk tolerance for a given individual (see Chapter 3 of the dissertation for additional detail on the consideration of nonlinear magnitude valuation as risk preference in the context of expected value).

Clinical diagnoses, subsyndromal symptoms, and even induced mood states of anxiety have been associated with increased risk aversion in reward decision-making (Giorgetta et al., 2012; Maner et al., 2007; Raghunathan and Pham, 1999), while induced sadness has been associated with decreased risk aversion (Raghunathan and Pham, 1999). In contrast, mania has been behaviorally associated with increased risky behavior, but trait level differences across questionnaire measures of mania have not been reliably found to relate to increased risky decision-making in a research environment (Ramírez-Martín et al., 2020). However, risk is not the only type of uncertainty we face in decision-making.

Decision-making under second-order uncertainty

Risky decision-making is often referred to as first-order uncertainty, in which we are uncertain about the outcome of an action. Decisions must sometimes also be made under second-order uncertainty, in which we are uncertain about the first-level uncertainty of the outcome (Bach et al., 2011). Second-order uncertainty may arise in multiple ways: One common occurrence is when some of the necessary information regarding an outcome probability is not available to us, requiring us to estimate an ambiguous outcome probability. Individuals differ in their tolerance of ambiguity, and it is likely that here, too, outcome valence plays a role; research has found that while individuals are generally

ambiguity avoidant in moderate likelihood gain outcomes, they may be ambiguity seeking in loss outcomes (Kocher et al., 2018).

A study exploring ambiguity sensitivity in threat avoidance, where participant decisions were made to avoid electrical stimulation, showed that higher levels of trait anxiety were related to a higher sensitivity to level of missing information in driving ambiguity avoidant behavior (Lawrance et al., 2022). Research has also shown a heightened allocation of resources to vigilance of uncertain threat (Geng et al., 2018; Grupe and Nitschke, 2013; Robinson et al., 2013) among individuals with anxious symptoms. This raises several research questions which are explored in greater detail in Chapter 3 of the dissertation. There we seek to investigate whether this relationship holds specifically for dimensional anxiety, or for shared variance between anxiety and depression. We also investigate relationships of dimensions of depression and mania, in addition to anxiety, in second-order uncertainty across reward, loss, and threat decisions.

Cost valuation and the cost/benefit tradeoff

A key aspect of real-world decision-making is the potential cost of taking an action. Costs may take various forms, including physical effort exertion, cognitive effort exertion, and delay of outcome. Costs generally serve to subjectively devalue action outcomes, and different forms of cost may use distinct neural processes in their subjective valuation. For example, physical effort costs may require the use of the cingulate motor area (CMA) and supplemental motor area (SMA) (Hosokawa et al., 2013; Klein-Flügge et al., 2016), whereas cognitive effort costs (such as computation or sustained attention) may require the use of executive control systems in the prefrontal cortex (Parr et al., 2023; Sayalı and Badre, 2021).

An additional and perhaps more subtle aspect of the decision-making process which may differ across individuals is the subjectivity of the tradeoff itself between outcomes and costs during the consideration process. Information regarding subjective value of outcome must be neurally integrated with information regarding subjective value of cost for a decision to be made. A large body of animal and human literature has investigated the neural underpinnings of value comparison in decision-making. For decisions involving non-physical costs like delayed outcome, the ventromedial and orbitofrontal areas of the prefrontal cortex (vmPFC and OFC, respectively), as well as neurons in the anterior cingulate cortex (ACC), have been widely observed to be involved specifically in the comparison of cost and benefit in decision-making (Bailey et al., 2016; Kennerley et al., 2011; Padoa-Schioppa and Assad, 2006). When the valuation of physical effort costs is involved, the supplementary motor area (SMA) has been found to encode subjective cost information and value comparison to take place in ACC (Klein-Flügge et al., 2016). These neural points of comparison between various outcomes and costs provide yet additional possibilities for mechanistic impairment in psychopathology.

Reinforcement learning

A related and highly relevant topic of research is the investigation of how decision outcomes shape our future behavior, often referred to as reinforcement learning (Sutton and Barto, 1998). In the simplest sense, the heart of the reinforcement learning update process may be expressed via the delta rule, a simplified version of the Rescorla-Wagner algorithm which represents the updating of action values V_t at time *t* based on feedback *r* over time (Rescorla and Wagner, 1972). This may be characterized as:

$$V_{t+1} = V_t + \alpha * \delta_t \quad (3)$$

$$\delta_t = r_t - V_t \quad (4)$$

where δ_t represents our prediction error for the most recent action outcome, α describes a learning rate which controls the speed at which we incorporate new information into our stored values, and r_t represents the outcome (typically characterized as reward) at time *t*.

When we learn the value of multiple possible actions a in response to multiple possible states s, the values may be represented algorithmically in a matrix of dimension [N_{states}, $N_{actions}$]. In modern research, these stored action values are typically referred to as Q-values, in reference to the Q-learning algorithm (Sutton and Barto, 1998; Watkins, 1989; Watkins and Dayan, 1992). Here, we may write the modified delta rule updating algorithm as:

$$Q(s,a)_{t+1} = Q(s,a)_t + \alpha * \delta_t \quad (5)$$

where $Q(s,a)_t$ now represents our value at time *t* for action a given state s. Our prediction error is defined as before, with:

$$\delta_t = r_t - Q(s,a)_t \quad (6)$$

Action selection

While our Q-values tell us the relative subjective value of each action at a given time-point, there are multiple ways we may define our policy of action selection using these values. For example, we might implement a policy to always select the action with the highest relative Q-value at a given time point; this policy is referred to as "greedy" action selection. In general, we may assume that there will be some noise in the process of comparing action values. One action-selection policy broadly used in RL research uses a softmax rule to assign probabilities to action selections based on a tradeoff of value-based evidence and decision stochasticity as follows:

$$p(a|s) = \frac{\exp(\beta * Q(s,a))}{\sum_i \exp(\beta * Q(s,a_i))} \quad (7)$$

Here β is a parameter referred to as the softmax inverse temperature, and governs how noisy the process of value comparison is. This may be clearly illustrated by considering the result of the equation above when the parameter β is set to 0. In this case, the probability of selecting an action reduces mathematically to 1/*nA*, where *nA* is the number of possible actions; in other words, $\beta = 0$ implies action selection by pure chance. As β increases, we increasingly weight the *Q*-values of actions in the probabilistic decision rule.

$$p(a|s) = \frac{\exp(0*Q(s,a))}{\sum_i \exp(0*Q(s,a_i))} = \frac{1}{1+1+\cdots} = \frac{1}{nA} \quad 8$$

In addition to stochasticity in the value comparison process, humans may sometimes choose actions randomly without regard to learned value; this behavior is sometimes referred to as "exploration," and contrasts with the "exploitation" of current values to select an action.

A common policy for action selection which incorporates these features is the ε-greedy (epsilon greedy) softmax policy, which assumes that actions are selected randomly with probability ε and according to stochastic value comparison with probability $1 - \varepsilon$.

$$p(a|s) = \varepsilon \cdot \left(1/n_A\right) + (1 - \varepsilon) \cdot \frac{\exp(\beta*Q(s,a))}{\sum_i \exp(\beta*Q(s,a_i))} \quad 9$$

Temporal difference learning

While the algorithms outlined above enable the simultaneous updating of multiple state and action values over time based on past experience, humans might also consider all the future rewards that they expect an action to yield when determining value. We may mathematically characterize this by enhancing the previously defined learning algorithm to use a temporal projection structure, where the reward value at each time is computed as a discounted projection of future expected rewards. This discounting reflects a higher value for rewards received sooner than those received later, and may be written as follows:

$$R_t = r_t + \gamma \cdot r_{t+1} + \gamma^2 \cdot r_{t+2} + \ \ldots \ = \sum_{k=t}^{m} \gamma^{k-t} \cdot r_k \quad 10$$

where γ represents our discount factor for each future time step. The present value of future reward is then updated at each time step to reflect new information. Although the projected value of reward here includes an expectation of future rewards, at this point it is still just a scalar reflection of future expectation based on past experience.

One common parametric mechanism of research interest in reinforcement learning is the learning rate α, which controls the rate at which value estimates are updated by experience. Within a single learning task there will typically be a wide span of individual differences in learning rate (Haile et al., 2024; Zerr et al., 2018), and individuals may learn differently in differing learning tasks or differently valenced paradigms (Fleishman, 1978; Smith et al., 2022).

Despite much interest in the role of psychopathology in learning, there have been mixed findings regarding a direct relationship between learning rate and dimensions of anxiety or depression. Although studies have observed relationships between overall learning impairments and depression, anxiety, and bipolar disorder (Esfand et al., 2024; Kumar et al., 2018; Pizzagalli et al., 2008), there have been no consistently replicated findings linking these deficits to changes in learning rate (Huys et al., 2013, 2015). Higher levels of trait anxiety have been linked to a difficulty in adjusting learning rate between stable and volatile stochastic environments (Browning et al., 2015), which more recent research has shown might actually be linked to shared variance between anxiety and depression rather than trait anxiety alone (Gagne et al., 2020). However, there may be more to impaired learning than RL processes alone, as we will discuss in subsequent sections.

Multiple systems of reinforcement learning

Reinforcement learning as described above stores and updates action values based upon recent experience, and is largely thought to be neurally instantiated through dopaminergic projections to areas of cortex and the basal ganglia (Collins and Frank, 2012). This system of learning through reinforcement is often referred to as "model-free" learning because of its simple reliance on past outcomes and lack of requirement for a sophisticated model of the world. However, in many circumstances humans show the ability to perform more complex learning, for example by integrating anticipated changes in future states based on devaluation of future reward in current decisions. This more complex approach to learning requires the maintenance and refinement of a model of the world, and is generally referred to as "model-based" learning (Daw et al., 2005; Daw and Dayan, 2014). In model-based learning, we are able to project forward possible paths and their related outcomes, and it is sometimes referred to as "goal-driven" learning or behavior. Using such projections is computationally expensive for the brain and cannot be performed exhaustively given the sheer number of possible future alternatives which would need to be considered. It is thought that the brain uses a trade-off between model-free and model-based learning to optimize the balance between metabolic needs and directed behavior (Bishop and Gagne, 2018a).

Research has shown that our models of the world are updated not just via the experience of action outcomes, but also through hippocampal "replay" of past action/outcome sequences (Johnson and Redish, 2005; Liu et al., 2021; Wise et al., 2021), and "simulation" of possible future action/outcome sequences (Foster and Wilson, 2006; Johnson and Redish, 2007; Schacter et al., 2007). Intriguingly, these modes of offline updating each seem to relate naturally to symptoms of psychopathology; for example, depression often involves rumination, a repetitive replay of (typically negative) experiences or thoughts (Berg et al., 2022; Lavender and Watkins, 2004). Both anxiety and depression have been associated with differences in simulation of future possible outcomes; individuals with higher levels of anxiety are apt to predict more negative events occurring in the future, whereas individuals with higher levels of depression predict both more negative and fewer positive future events (MacLeod and Byrne, 1996; MacLeod and Salaminiou, 2001).

Supporting systems during learning

The process of learning via reinforcement, whether model-free or model-based, does not happen in isolation. Rather, an ensemble of brain processes is likely interactively involved as our brain learns. For example, heightened activity in the prefrontal cortex (PFC), an area associated with executive function and attention, has been observed during learning, particularly during initial learning (Collins and Frank, 2012). It has been suggested that this early-process heightened activity of PFC during learning may at least in part reflect the use of working memory systems supporting initial learning (Collins, 2018; Collins et al., 2014; Collins and Frank, 2012). In contrast to slow but long-lasting reinforcement learning processes, working memory is thought to provide brief access to relevant information to guide immediate decision-making (Bays et al., 2024; Van Ede and Nobre, 2023), and is characterized by both a limited effective capacity and a short duration of effect . Behavior during learning is influenced by both RL and WM systems, and working memory may directly influence reinforcement learning values during learning acquisition (Collins and Frank, 2012).

In light of this, recent investigations in computational cognition have expanded the RL modeling paradigm in human research to reflect an interactive involvement of WM and RL systems during learning (Collins, 2018; Collins and Frank, 2012). This enhanced modeling paradigm proved immediately useful with respect to psychopathology, when long-observed reinforcement learning deficiencies in individuals with schizophrenia were found to be entirely attributable to computational parameters of the working memory system rather than RL systems (Collins et al., 2014). Anxiety has been associated both with working memory and learning deficits (Balderston et al., 2017; Forster et al., 2015; LaFreniere and Newman, 2019), directly motivating our research question in Chapter 4 of the dissertation: are learning deficits in anxiety subtypes attributable to WM systems, RL systems, or both?

Final thoughts

As with any mathematical modeling endeavor, certain assumptions are required to analyze and interpret the computational representation of unobserved cognitive processes. From classical to Bayesian statistical approaches, random effects to fixed effects to hierarchical formulations of algorithms, maximum likelihood estimators to maximum a posteriori estimators to full Markov chain Monte Carlo sampling of posterior parameter distributions, experts in the field of computational cognitive modeling bring varying skill sets and preferences to bear on computational psychiatry research.

One important maxim to keep in mind in this and all other mathematical modeling is the legendarily paraphrased statement attributed to mathematician George Box that “all models are wrong, but some are useful” (see below for full quote).

“Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful.” (Box, George E. P. and Draper, Norman R., 1987)

We seek knowledge of the brain through models which represent cognitive and neural processes, and therefore it is in our best interest to ensure through whatever means available to us that our models are as not-wrong as possible. However, our present research goal is not to explicate a full and precise mathematical model of the entire human brain, taking into account every possible neural connection and interaction; indeed, this would be beyond our capabilities (at present). We are not striving to propose any single "correct" model of cognitive process.

To ensure the scientific reliability of our findings, we must emphasize a clear adherence to scientific principles in our iterative model-building approach and interpretation of the resulting insights to cognition. Among these scientific best practices: focus on building models which are accurately predictive to a high degree of out-of-sample data when inverted to be generative; replication of model results by ourselves and other researchers; and transparent elocution of the methodology used and the reasons for using it. The body of knowledge being uncovered is a scientific group effort; as more nuance is understood about the processes of interest, we are able to propose more accurate models, which may then provide additional insight.

Chapter 2

Latent subdimensions of anxiety and depression differentially influence exertion of effort in pursuit of reward versus avoidance of threat

Chapter based upon and adapted from manuscript in prep: Latent subdimensions of anxiety and depression differentially influence exertion of effort in pursuit of reward versus avoidance of threat (Jennifer Senta, Anne G.E. Collins, Peter Dayan, Sonia Bishop)

Abstract

Anxiety and depression are notoriously comorbid. However, empirical studies in animal models suggest that willingness to exert effort may reveal important differences between anxious and depressed phenotypes. In humans, symptoms of depression have been linked to reduced willingness to exert effort to gain rewards. In rodents, decreased exertion of effort to avoid threat characterizes depressive phenotypes and is responsive to antidepressant agents while increased effortful threat avoidance has been linked to a more anxiety-like phenotype responsive to anxiolytic agents. Here we sought to investigate whether subdimensions of anxious and depressed affect differentially impact effortful reward pursuit versus threat avoidance in humans. Hierarchical Bayesian modeling revealed that component mechanisms underlying effortful reward pursuit were impacted by latent dimensions capturing depressive symptoms, with distinct effects of factor scores capturing apathy versus anhedonia. Meanwhile, only the latent dimension tapping worry-related symptoms significantly impacted effortful avoidance of threat, increasing willingness to engage across all threat magnitudes and decreasing sensitivity to threat magnitude changes.

Introduction

Everyday life is characterized by decisions in which we must evaluate the costs versus benefits of various actions. One frequently encountered form of cost is a necessity to exert physical effort in the execution of an action. We might select such effortful actions with the goal of obtaining a positive reward, such as monetary gain, or with the goal of avoiding a negative outcome, such as physical threat. For example, we might ask ourselves if the effort required to prepare an elaborate meal is worth the reward of enjoying it, or if an effortless snack might be preferable. Similarly, we might ask ourselves whether avoiding a dangerous neighborhood is worth walking some distance out of our way, or if we prefer the shortcut which is less physically demanding. Individuals vary in how much physical effort they are willing to exert in pursuit of specific outcomes, and aberrant willingness to exert effort has been associated with a variety of psychopathological conditions (Bishop and Gagne, 2018b; Culbreth et al., 2018; Nord et al., 2017; Treadway and Salamone, 2022; Whitton et al., 2015).

A large body of experimental literature in rodents has established a relationship between depressive phenotypes and reduced effort exertion in reward pursuit. Effortful reward pursuit in such paradigms is typically elicited via physical lever presses to gain a highly desirable food option, and is typically measured in terms of choice selection between a low effort/low reward option and a high effort/high reward option (Canonica and Zalachoras, 2022; Haluk and Wickman, 2010). Findings link depressive phenotypes to low levels of high effort/high reward choice selection (Bergamini et al., 2016a, 2016b; Nunes et al., 2014; Vichaya et al., 2019; Xu et al., 2017), with this behavior being normalized by the administration of antidepressant psychopharmacology (Bergamini et al., 2016a, 2016b).

In humans, effortful reward pursuit is usually studied with rewards consisting of monetary gain (Cléry-Melin et al., 2011; Culbreth et al., 2018; Geaney et al., 2015; Treadway et al., 2009; Treadway and Salamone, 2022). A human analog of the rodent task, the Effort Expenditure for Rewards Task (Treadway et al., 2009) , which pits a low effort/low reward option against a high effort/high reward option, has been used in studies of individuals with depressive disorders. Findings from these studies have been broadly consistent with the animal literature, with those diagnosed with major depressive disorder (MDD) showing lower selection of high effort/high reward options than healthy control participants (Subramaniapillai et al., 2019; Treadway et al., 2012; Wang et al., 2022; Yang et al., 2014). Among medicated MDD patients in remission, reduced high effort/high reward choice selection has also been shown to be predictive of future depressive relapse (Berwian et al., 2020).

In addition to the use of binary psychiatric diagnoses such as MDD, researchers have used well-validated standardized questionnaires to examine the relationship between continuous measures of depressive symptomatology and effortful reward pursuit. Here a more complex picture has emerged. Both apathy and anhedonia have been associated with reduced effortful reward pursuit (Bonnelle et al., 2015a, 2015b; Geaney et al., 2015; Treadway et al., 2012; Treadway and Salamone, 2022). However, it remains unclear whether these dimensions of depressive symptomatology influence effortful reward pursuit via common or distinct mechanisms. Further complicating this question, scores on self-report measures of depressive symptomatology are highly correlated with each other and with scores on self-report measures of anxiety. This raises the question of how the various precise mechanisms of impairment in decision-making are related specifically to symptom-level, syndrome-level, or comorbidity-based characterizations of depression and anxiety.

Despite the high comorbidity between anxiety and depression (Hopwood, 2023; Pollack, n.d.; Sartorius et al., 1996), there has been little human research investigating the relationship of anxiety with individual differences in effortful reward pursuit. However, a large basic neuroscience literature on the instrumental avoidance of threat (Ball and Gunaydin, 2022; LeDoux et al., 2017; Rachman, 1976; Ramirez et al., 2015) has led to instrumental avoidance paradigms being used in investigations of human anxiety (see Pittig et al., 2018 for a review). Here, however, effort costs are typically not manipulated.

Intriguingly, research done in rodents (Cryan et al., 2005; Lee et al., 2017; Porsolt et al., 1977; Ramboz et al., 1998) suggests that willingness to exert physical effort to avoid threat in a forced

swim test (FST) paradigm may differentiate animals with anxious versus depressive phenotypes. The FST involves the placement of rodent subjects into a body of water with no rest platforms or exits and measures the amount of time the animal spends exerting physical effort to stay afloat or find an escape (Porsolt et al., 1977). Rodents showing increased anxiety-like states (for example, as a result of 5HT1a receptor knockout or induced alcohol withdrawal; (Lee et al., 2017; Ramboz et al., 1998) have been found to exert significantly more physical effort (i.e., longer swim times) to stay afloat in the FST than controls, and this behavior may be normalized by anxiolytic psychopharmacology (Lee et al., 2017). This contrasts with classical findings that rodents with depressive phenotypes typically show reduced swim times on the FST (Cryan et al., 2005; Porsolt et al., 1977; Weiss et al., 1998, 1981), which is reversed by administration of antidepressant agents (see Cryan et al., 2005 for review). This leads to the question of whether anxiety and depression might show differential impacts on effortful avoidance of threat in human decision-making.

Our study hence sought to fully characterize the relationship between common and unique aspects of anxious and depressive symptoms and alterations or biases in the computations underlying effort-based decision-making. Our specific aims were as follows: First, we sought to investigate whether separable latent subdimensions of depression impacted the same or different component mechanisms of decision-making involved in effortful reward pursuit. Second, given the lesser literature on the influence of anxiety on effortful decision-making, we sought to determine whether cognitive or physiological symptoms of anxiety were associated with differences in effortful decision-making, and if so whether this varied between effortful reward pursuit and effortful threat avoidance. Finally, we sought to determine whether either, or both, effortful decision-making in pursuit of reward and in avoidance of threat differentiated individuals with elevated anxious versus depressive symptomatology.

Materials & Methods

Participants: Behavioral task

Sixty-one participants completed the behavioral experiment over two sessions on site at the Computational Psychiatry and Affective Neuroscience Lab at UC Berkeley. Eleven participants were ultimately excluded: 5 due to equipment failure and 6 due to incomplete data collection. Final participant group n=50, 57% female, mean age 22 +/- 2.3 years. Participants were recruited from the UC Berkeley student research participation pool (RPP) and the local community. Participants completed written informed consent before each study session. The tasks and consent processes were approved by the UC Berkeley Committee for the Protection of Human Subjects. Participants could choose whether to receive psychology course participation credit in the amount of 1 credit per hour as their main source of study compensation (UC Berkeley students only, credited in accordance with UC Berkeley RPP policies), or to receive cash compensation of $15 per hour for each of the two study sessions. Study sessions were scheduled and paid for a minimum of 2 hours each.

In addition to hourly compensation for participation, participants were eligible for bonuses paid in cash regardless of recruitment source. Participants received a $2.50 bonus during each session to incentivize maximal effort during effort calibration. Participants received a $15 bonus upon completing both sessions to incentivize completion over multiple days. During the reward session,

participants accumulated points represented as apples; at the end of the session, the number of apples collected relative to the total possible number of apples to be collected during the task was converted pro-rata to a cash bonus of between $0 and $7.50. Participants were not informed in advance of how the apples collected during the task would be converted into a cash bonus but were aware they would earn a greater bonus by collecting more apples.

Participants: Latent dimensions of psychopathology

Two separate sets of data were collected for dimensionality reduction analysis via online questionnaire administration through Qualtrics. The first data set, used for exploratory analysis, consisted of 453 participants (285 participants recruited from UC Berkeley RPP and 168 participants recruited via Amazon Mechanical Turk (AMT) [n=453 total; 264 female; age range 18-45; specific age not collected]. Participants recruited through AMT were compensated $12 for responding to a number of questionnaires including those used in the current study. Participants through RPP were compensated with 1.5 RPP psychology course credits per UC Berkeley policy for responding to a number of questionnaires including those used in the current study.

The second data set, used for confirmatory analysis, consisted of 252 participants recruited via AMT [155 female; average age 42.8 +/- 12.1 years] to participate in a separate online behavioral study administered by our lab. This online behavioral study included the questionnaire responses used in the current study. These participants also completed an online behavioral learning task; however, only the questionnaire data for the relevant scales was used in the current study. Participants completed online informed consent before participation in the study. The study and consent process were approved by the UC Berkeley Committee for the Protection of Human Subjects. Participants were compensated $20 for participating in the study, plus a $5 bonus to incentivize performance during the behavioral task (not analyzed here).

Bifactor modeling of item responses was used to parse shared vs. latent dimensions of unique variance. This approach has shown considerable value in modeling comorbid psychopathological conditions such as depression and anxiety, which often share some overlapping variance but also feature distinctly separable dimensions (Gagne et al., 2020; Reise, 2012; Reise et al., 2010; Simms et al., 2008). The number of subdimensions specified in the model may depend on theoretical motivation as well as robustness of the item response data being investigated.

To account for the ordinal nature of the item responses, polychoric correlations were used in all analyses. Sampling adequacy of the exploratory data set was confirmed using the Kaiser-Meyer-Olkin (KMO) test, which yielded a sampling adequacy of 0.82 (values over 0.80 are considered adequate) on the polychoric item correlation matrix. A parallel eigenvalue decomposition analysis, which uses resampling techniques to compare the covariance of item responses with identically-sized matrices generated from random Gaussian noise, was conducted in R and indicated that up to 21 factors might be reliably distinguished from noise in the exploratory data set. This large possible number of identifiable factors may have been due to the large sample size and the associated possibility of extreme overfitting. We performed exploratory bifactor analysis specifying one general factor and varying numbers of specific factors, starting with two specific factors and ranging up to seven, in order to compare model fit for both exploratory and confirmatory data sets. Bifactor analysis was conducted in R using the psych package omega

function. The Schmid-Leiman procedure was used to transform an oblique factor analysis of the full item response data into a bifactor structure with orthogonality of the common and distinct factors (Schmid & Leiman, 1957). Resulting factor loadings were thresholded to only include loadings > 0.2.

A comparison of the Bayesian Information Criteria (BIC), which assesses model fit while penalizing for complexity, showed continued improvement in exploratory model fit as additional specific factors (up to the seven total specific factors tested) were included. However, increasing the number of specific factors above five resulted in sub-standard confirmatory fit in the held-out data set, indicating that more complicated models were overfit to the exploratory data. We assessed the fit of each model to the confirmatory data set using both confirmatory fit index (CFI) and Tucker-Lewis measure of fit. Confirmatory fit measures of 0.95 or higher are considered a strong fit. For greater than five specific factors, the confirmatory fit metrics were less than 0.95, indicating that five or fewer specific factors provided the strongest out-of-sample validation fit.

The bifactor model with five subfactors therefore provided the best relative BIC score of exploratory model fit while maintaining a confirmatory fit index (CFI) above 0.95 in the held-out data. Additionally, this model provided cleanly interpretable factors which were not present in models using smaller numbers of specific factors, including unique factors loading primarily on each of apathy and anhedonia, allowing these specific symptoms to be tested while accounting for shared variance across conditions. The bifactor model with five specific factors provided a robust characterization of the total exploratory item response covariance, with Cronbach's alpha = 0.96 and McDonald's $omega_{total}$ = 0.97.

As a check of the bifactor structure against the questionnaire metrics included in the set, we calculated a factor score for each participant in the exploratory data set on each general and specific subfactor. Participant factor scores were computed using the Anderson-Rubin method, which preserves orthogonality of the general and specific scores (Anderson TW, Rubin H., 1956). We then assessed the correlation of these factor scores with participant scores on questionnaire scales. Most scale measures were correlated with the common factor 'G', indicating that this represented shared variance across anxiety and depression measures included in the study. Individual specific factors correlated primarily with measures comprising their major features, as described elsewhere in this paper.

We next used the factor loadings identified in the exploratory analysis to calculate participant factor scores for each in-lab behavioral participant using the same calculation methodology described above. The correlation of these factor scores with participant subscale scores were highly congruent with the correlations found in the exploratory analysis, particularly for the primary subscales associated with each subfactor, confirming that the bifactor structure as applied to the behavioral data was indeed capturing the same subsets of variance (see Figure 2.2).

Behavioral Task: Effortful decision-making

Participants completed two in-lab sessions, each of a maximum two hours duration. The sessions were scheduled to be completed on different days to ameliorate the effects of fatigue, but were

restricted to be within 5 days of each other to mitigate variations in trait-based questionnaires. At the start of each session, after completing informed consent, participants were asked a series of questions to establish eligibility. Participants were informed of restrictions before recruitment, and were not allowed to continue if they reported cannabis use within the previous two weeks, or if they were currently taking antidepressant or anxiolytic medications. Prior to the threat avoidance session, participants were also given a short medical screening to ensure that participants with a history of fainting, vasovagal syncope, or any heart condition or pacemaker-like devices did not participate (out of an abundance of caution). Session order was counterbalanced across participants to perform the reward pursuit or threat avoidance session on the first day.

Calibration of effort

Effort was exerted in the task by applying squeezing pressure to a dynamometer (hand grip) device (BioPac) with the dominant hand. Following task instructions, but before any study trials, the maximum level of force to be used in the experiment was calibrated so as to be challenging but attainable to the individual participant. To perform the calibration, participants were shown an empty tree with a top ladder rung set to a very high difficultly level (no participants reached or exceeded the initial top rung difficulty level during calibration). Participants were given 5 seconds to squeeze with their maximal force while remaining comfortable to see how high they could climb on the tree, and told that they might receive a bonus if they climbed high enough on the tree. Participants were then given two more such trials, with the instruction to see if they could exceed their previous height. All participants received the full bonus regardless of actual effort exerted during calibration.

The maximum force applied by the participant during calibration was set as the Maximum Voluntary Contraction (MVC), following previous work (Bonnelle et al., 2015b). Because our task required participants to briefly sustain the required force for a period of at least 0.5 seconds and participants completed 6 blocks of trials, we set 85% of the MVC as the participant's "sustainable" MVC (SMVC) to mitigate effects of fatigue on ability to exert sufficient effort. The highest level of effort required for the task was set equal to the SMVC, and each lower level of effort was reduced by 10% of the SMVC.

Participants performed the calibration at the beginning of each session to account for possible inter-day differences in grip strength. However, the SMVC on the second session was not allowed to vary by more than 15% of the initial session SVMC in order to mitigate any intentional changes in grip strength once the task structure had been seen.

Calibration of electrical stimulation

Electrical stimulation in the threat avoidance session of the experiment was delivered via a DigiTimer constant current stimulator model DS7A, intended for human subjects research use. Prior to the start of the session, the electrode was sanitized and attached to the inside of the wrist of the participant's non-dominant hand using hypoallergenic medical tape. The experimenter adjusted the voltage to just above zero, and administered single 2ms test pulses while incrementing the voltage until the participant reported that they could feel a sensation similar to that of a small pinprick. This level of voltage was then set as the participant's baseline (level 1).

In order to deliver electrical stimulation of varying magnitudes throughout the task, we then calibrated the maximum amount of stimulation that might be delivered to the participant during the task if sufficient effort was not exerted to avoid it. Levels of stimulation higher than level 1 were administered by delivering varying numbers of rapid level 1 pulses in quick succession (12ms apart).

Starting with a single level 1 pulse, we used the computer to incrementally increase the number of level 1 pulses, asking the participant to rate each stimulation burst on a wish-to-avoid scale of 1 to 10, informing them that a subjective level 7 was the highest level we would ever administer. Participants were instructed that level 7 should be a level that they would wish to avoid, but which they found tolerable and which would not cause them any lasting distress. Once a participant indicated their level 7 had been reached, we randomly sampled 14 additional ratings by repeating various numbers of pulses between but not inclusive of levels 1 through 7 and asking the participant to rate each burst of stimulation. Previous literature has indicated a non-linear (sigmoidal) relationship between increasing numbers of electrical pulses and increased ratings of aversion (Browning et al., 2015). Once ratings were complete, a best-fit sigmoid curve was fit to the ratings data between levels 1 and 7. Evenly spaced points on the curve were defined to represent the intermediate stimulation levels 2-6. In the rare case that a participant had responded very unreliably to the varying pulse numbers (typically by becoming less sensitive to pulses over time, and rating high numbers of pulses as lower subjective levels towards the end of calibration), a sigmoid curve could not be fit and the calibration was repeated beginning with setting a new level 1. Additionally, if a participant reached their subjective level 7 before 8 pulses had been delivered, or did not reach their subjective level 7 by the time 32 pulses were delivered, the calibration was repeated beginning with setting a new level 1.

Once the sigmoid curve had been successfully fit to the participant's final calibration data, the subjective ratings and corresponding number of pulses were used to define the stimulation delivered in the threat avoidance task as detailed below.

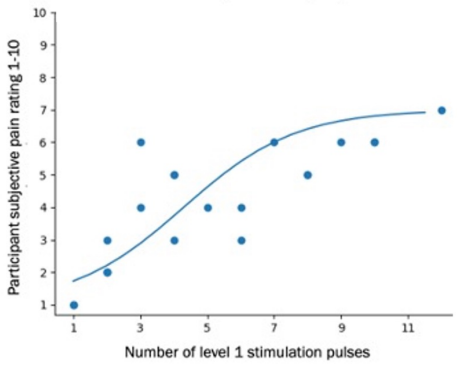


Figure 2.1: Illustrative sigmoidal interpolation in electrical stimulation calibration. Participant level #1 is calibrated to be the single pulse they define as feeling like a small pin-prick. Various numbers of the level #1 pulse are then delivered in train (2ms pulses spaced 12ms apart), and the participant rates the increasing pulses until level #7 is reached. A sigmoidal curve is then fit to the pulse sequences to interpolate between stimulation levels 2-7.

Reward pursuit

After instructions, calibration of force, and two practice trials, each participant completed 6 blocks of 26 trials each, with a 10-second rest period after every 5 trials, and a 30-second rest period after each block. Trials were designed such that each block contained one full set of orthogonally manipulated combinations of the 5 possible effort levels and 5 possible reward levels, plus one additional trial for tracking. Trials were pseudo-randomized and fixed across participants.

On each trial, an image of a tree was displayed. The trunk of the tree contained an empty vertical rectangle, intersected by 5 horizontal lines in the upper half of the rectangle representing "ladder rungs" to which the participant might be asked to climb. The top of the box (i.e., the top ladder rung) represented the participant's sustainable maximum voluntary contraction (SMVC), which was set according to the "calibration of force" section above. Each lower rung then represented a 10% reduction in force, with the rungs visually spaced linearly in the top half of the box to illustrate 60%, 70%, 80%, 90%, and 100% of the SMVC.

For each reward pursuit trial, one of five possible levels of reward was offered by showing 1, 4, 8, 12, or 16 apples displayed in the tree. Each trial also included one highlighted ladder rung indicating the amount of force that would be required to successfully climb the tree and gain the apples for that trial. A participant was given 5 seconds to decide whether or not to climb the tree to obtain the reward. If they declined a trial, they received 0 points, and proceeded to the next trial. If they accepted the trial, they were given 5 seconds to climb to the appropriate ladder rung by applying a sufficient amount of grip force to the dynamometer. Real-time feedback of the grip was shown as a red color filling the rectangle displayed on the tree as the participant squeezed. If the participant held the force at or above the appropriate ladder rung for at least 0.5 seconds, they were awarded the apples in the tree; otherwise, they received zero points and continued to the next trial. When participants successfully completed a trial, they were shown a message that indicated they had added the specified number of apples to their basket. Participants were informed that they would receive a bonus based on how many apples they collected during the task, but were not informed of the conversion rate and were not shown their balance of apples already gathered over time. At the end of the task, the bonus was calculated by prorating the maximum bonus of $7.50 according to the number of apples successfully gathered.

Threat avoidance

The threat avoidance task was analogous to the reward pursuit task, but here participants chose whether to climb trees in order to escape varying numbers of bear traps around the base of the tree (1, 4, 8, 12, or 16 bear traps; no apples were shown) which represented various levels of electrical stimulation. Trial structure of effort levels and stake levels was identical to the reward pursuit session. Participants were informed that their subjective calibration of electrical stimulation levels would be translated into the number of bear traps shown, and that higher numbers of bear traps indicated higher levels of stimulation, but were not specifically informed of the conversion mapping. This mapping was set such that the 5 distinct numbers of bear traps corresponded to subjectively rated stimulation levels 2, 4, 5, 6, and 7 for the given participant.

Timing of trial offer, force requirements, and success requirement was identical to reward pursuit session. If participants did not succeed at a given trial or elected not to exert effort on that trial, they immediately received the level of electrical stimulation associated with that trial.

Computational modeling of behavior

We modeled the binary decision data of participants in reward pursuit and threat avoidance using a class of subjective value (SV) models that traded off between stake (level of reward or threat) and effort. We began with a model which considered only stake level, effort level, and baseline subjective bias against trial acceptance, and then incrementally augmented the model with parameters of known behavioral/theoretical interest. At each stage of augmentation, we assessed model fit using the Watanabe-Akaike Information Criteria (WAIC) (Watanabe S, Opper M, 2010), a method which assesses comparative model fit while penalizing for model complexity in a hierarchical Bayesian fitting framework. We retained parameters that improved model fit in the next model version at each stage.

Hierarchical Bayesian parameter estimation

We used a hierarchical Bayesian sampling approach to infer posterior distributions for model parameters at the individual and group levels as well as to assess the effects of individual differences in trait subdimension scores. Separate parameters for reward pursuit and threat avoidance trials were inferred jointly, and for all threat choice parameters we controlled for the effect of individual differences in baseline tolerance to electrical stimulation.

Models were fit using MCMC sampling using the rstan package (Stan Development Team., 2022) in R (R Core Team, 2018). For each model four chains of posterior sampling were run, each with 5500 total sampling iterations and 3500 "burn-in" iterations. Convergence of the sampling was determined by visual inspection of MCMC sampling chains and inspection of the Gelman-Rubin statistics for parameters. Gelman-Rubin statistics provide a measure of the consistency of parameter variance within sampling chains vs. between sampling chains. Values very close to 1.0 are generally considered ideal as they indicate high reliability of variance sampling across independent Markov chains. 98.7% of Gelman-Rubin statistics across 936 measures of convergence were between 1.00 and 1.01 in our final model. 1.3% of the statistics were between 1.01 and 1.10, and none were above 1.10. All of the group-level parameters quantifying the influence of traits on cognitive processes (which were used for hypothesis inference) had Gelman-Rubin statistics between 1.00 and 1.005, indicating that all had superior convergence reliability.

Each participant parameter was assumed to be drawn from a Gaussian group-level distribution for that parameter described as $\sim Normal(\mu_{group_parameter}, \sigma_{group_parameter}^{2})$. Latent dimension factor scores were parameterized as offsets to each group distribution mean to assess the impact of individual differences in factor scores on cognitive model parameters. Group level hyperparameters were given uninformative priors of mu ~ Normal(0,10) and sigma ~ Exponential(0.5). Latent dimension offset parameters were given uninformative priors of ~Normal(0,5).

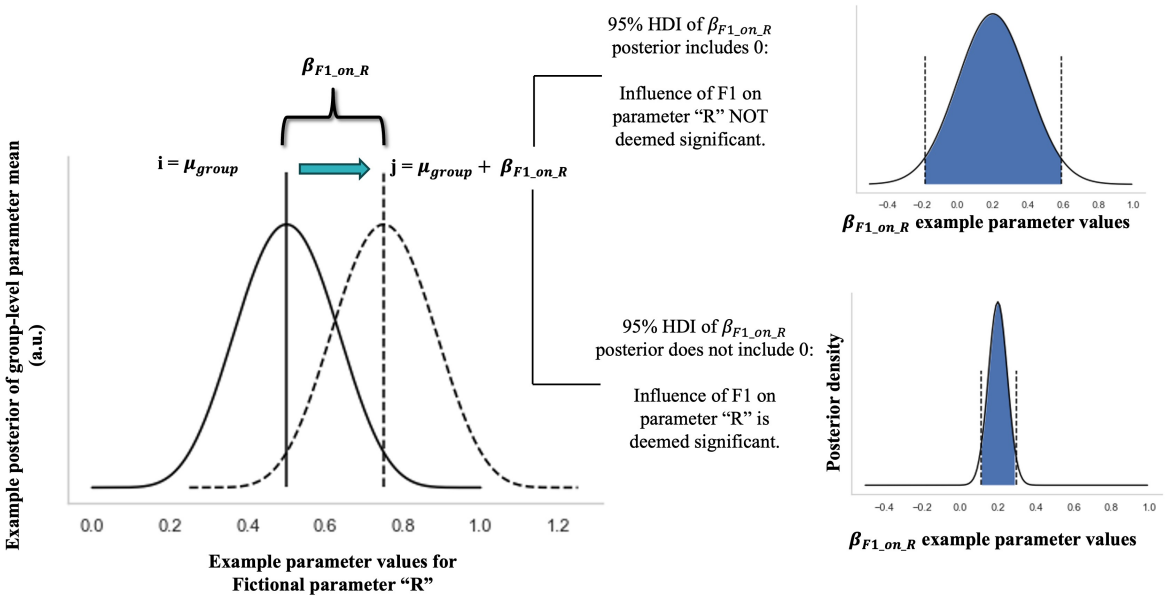


Figure 2.2: Illustration of group-level trait offset to posterior parameter distribution in hierarchical Bayesian framework. A fictional parameter denoted as "R" has its group-level mean μ shifted by group-level adjustment parameter $\beta_{F1_on_R}$, which quantifies the group-level distributional shift due to participant scores on hypothetical F1. Significance of trait-based shift on R due to F1 scores is treated as significant if the 95% highest density posterior distribution interval (HDI) of the adjustment parameter excludes 0; otherwise, the trait-based shift is not treated as significant.

As an example, shown below is the fully expanded assumed distribution for the linear stakescale parameter in all models for reward and threat trials. This example shows the separation of parameters between reward and threat trials and the individual difference offsets of each factor score on the parameters. In addition, for each threat parameter, we modeled a group-level offset parameter for baseline level of shock defined by each participant as their "level 1" (denoted L1mV) to control for potential differences in baseline shock sensitivity. Analogous expansions for all model parameters were fit but are notationally compressed for clarity throughout the paper.

$$\gamma_{_rew_ptp} \sim \text{Normal}\ (\mu_{\gamma_rew} + \sum_i \beta_{\gamma_rew_Fi} * F_{i_ptp},\ \sigma_{\gamma_rew}{}^2) \quad 1(a)$$

$$\gamma_{_threat_ptp} \sim \text{Normal}\ (\mu_{\gamma_threat} + \sum_i \beta_{\gamma_threat_Fi} * F_{i_ptp} + \beta_{\gamma_L1mV} * L1mV_{ptp},\ \sigma_{\gamma_threat}{}^2) \quad 1(b)$$

Model #1: SV Model

For the simplest model (Model #1) we assumed that at each trial the participant would compare the level of stake (reward or threat; γ_{reward}, γ_{threat}) on offer to the effort required for that trial (κ_{reward}, κ_{threat}). We also allowed for an individual baseline bias toward refusing trials in each paradigm (b_{reward}, b_{threat}) to offset the subjective net value. A logistic function was used to map the

participant's valuation of the trial to a probability of trial acceptance, and choice to engage on each trial was assumed to follow a Bernoulli distribution at that probability.

The subjective net value (SNV) for trial t was calculated as:

$$SNV_t = \gamma(\ Stake_t\ ) - \kappa(\ Effort_t\ ) \qquad 2(a)$$

The probability that a participant chose to engage on the trial was then given by:

$$P(Yes)_t = (1 + \exp(-[\ SNV_t - b\ ]))^{-1} \qquad 2(b)$$

We then enhanced the model incrementally to parameterize additional features and performed model comparison to ascertain the model which best characterized the behavioral data. Model #1 had WAIC=7,237.58

Model #2: SVG Model

Previous research has indicated that there may be gender-based differences in willingness to exert effort in reward pursuit paradigms (Treadway et al., 2009). We considered whether parameterizing gender difference for reward pursuit and threat avoidance (g_{reward}, g_{threat}) would provide an improved model fit to our data by adding a parameter for gender.

The SNV for trial t was still calculated as:

$$SNV_t = \gamma(\ Stake_t\ ) - \kappa(\ Effort_t\ ) \qquad 3(a)$$

The probability that a participant chose to engage on the trial was then given by:

$$P(Yes)_t = (1 + \exp(-[\ SNV_t - g \cdot Gender - b\ ]))^{-1} \qquad 3(b)$$

This model did not fit the data as well as Model #1 (Model #2 WAIC=7,238.26), so we did not include the gender parameter in further incremental changes to the model.

Model #3: SVF Model

We next added a parameter to account for individual differences in the influence of fatigue over the course of the experiment (f_{reward}, f_{threat}). This parameter was tied to the (scaled) current trial number. The SNV for this model was calculated as:

$$SNV_t = \gamma(\ Stake_t\ ) - \kappa(\ Effort_t\ ) \qquad 4(a)$$

The probability that a participant chose to engage on the trial was then given by:

$$P(Yes)_t = (1 + \exp(-[\ SNV_t - f \cdot Trial_t - b\ ]))^{-1} \qquad 4(b)$$

This model provided a substantially better fit to the data over Model #1 (Model #3 WAIC=5,836.48), so the fatigue parameter was carried forward to future models.

Model #4: SVFP Model

We next explored whether the data would be better fit by allowing individual differences in the valuation of stake to vary both linearly and nonlinearly. Utility theory in economics has established that increases in subjective value of reward tend to depend not just on the difference in reward offered, but in the relative difference of the rewards. This captures the idea that an increase from a hypothetical reward of $1 to $1000 might be valued more highly than an increase in hypothetical reward from $500,000 to $501,000. For this model we therefore added a participant-level nonlinearity parameter to stake in reward pursuit and threat avoidance trials (θ_{reward}, θ_{threat}).

We parameterized this in our model by calculating SNV at trial *t* as follows:

$$SNV_t = \gamma(\ Stake_t\ )^{\boldsymbol{\theta}} - \kappa(\ Effort_t\ ) \qquad 5(a)$$

The probability that a participant chose to engage on the trial was then given by:

$$P(Yes)_t = (1 + \exp(-[\ SNV_t - f \cdot Trial_t\ - b]))^{-1} \qquad 5(b)$$

The addition of a nonlinear reward parameter improved model fit over Model #3 (Model #4 WAIC=5,741.09), so this parameter was included in further models.

Model #5: SVFP SperE

Next, we evaluated a participant-level interaction parameter on stake per effort (δ_{reward}, δ_{threat}). This quantified the sensitivity of each participant to increases in stake relative to effort as follows.

First, SNV at t was calculated as:

$$SNV_t = \gamma(\ Stake_t\ )^{\boldsymbol{\theta}} - \kappa(\ Effort_t\ )\ + \delta(\ Stake_t/Effort_t\ ) \qquad 6(a)$$

The probability that a participant chose to engage on the trial was then given by:

$$P(Yes)_t = (1 + \exp(-[\ SNV_t - f \cdot Trial_t\ - b]))^{-1} \qquad 6(b)$$

This parameter improved model fit over Model #4 (Model #5 WAIC=5,732.81), so the interaction term was included in further models.

Model #6: SVFP SperE Persev

We next assessed whether participants showed a tendency to perseverate in their choices (also referred to as "stickiness of choice"; η_{reward}, η_{threat}). This model included a parameter to evaluate the repetition of the previous choice made by the participant at each trial as follows.

First, SNV at t was calculated as:

$$SNV_t = \gamma(\ Stake_t\ )^{\theta} - \kappa(\ Effort_t\ )\ + \delta(\ Stake_t/Effort_t\ ) \qquad 7(a)$$

The probability that a participant chose to engage on the trial was then given by:

$$P(Yes)_t = (1 + \exp(-[\ SNV_t - f \cdot Trial_t\ + \rho \cdot Psuccess_{t,e}\ - b]))^{-1} \qquad 7(b)$$

This model improved fit to the data compared with Model #5 (Model #6 WAIC=5,717.80), so the parameter for perseveration was included in subsequent models.

Model #7: SVFP SperE Persev Psuccess

Our experimental paradigm was designed so that all levels of effort should be achievable by each individual. We calibrated each participant's highest possible required level of effort on any trial to be 85% of their maximum force demonstrated during a pre-task calibration period. However, possibly due to extreme fatigue or spontaneous over-exertion during the calibration period, some participants accepted some trials but still failed to achieve the level of effort required for success. We investigated whether an individual's cumulative (within-session) probability of success on trials at a given effort level influenced their choices at that effort level (ρ_{reward}, ρ_{threat}). This was parameterized in the model as follows:

First, SNV at t was calculated as:

$$SNV_t = \gamma(\ Stake_t\ )^{\theta} - \kappa(\ Effort_t\ )\ + \delta(\ Stake_t/Effort_t\ ) \qquad 8(a)$$

The probability that a participant chose to engage on the trial was then given by:

$$P(Yes)_t = (1 + \exp(-[\ SNV_t - f \cdot Trial_t\ + \rho \cdot Psuccess_{t,e}\ + \eta \cdot Choice_{t-1} - b]))^{-1} \qquad 8(b)$$

This model provided an improved fit to the data over Model #6 (Model #7 WAIC=5,709.68), and therefore we included the parameterization of probability of success in the final model. Model #7 was the best-fitting model overall, and was used for all further hypothesis testing and inference.

Model validation: simulated data

To validate the final behavioral model, we simulated data based on the generative final model and compared actual participant behavior to the final simulated data. This provides a test of whether the best-fit model considered is capable of having produced the pattern of choices by actual

participants. The model substantively qualitatively reproduced the participant behavior and is illustrated in Figure 2.2 of the main text.

Model validation: parameter recovery

As a further test of validation, we tested our ability to accurately recover parameter values in the final model by generating simulated data and then fitting the final model uninformatively to the simulated data. We then compared the true (known) parameter values to the model-fit (recovered) parameter values. In hierarchical Bayesian inference, a general goal for parameter recovery is that 95% of the recovered posterior parameter distributions include the true parameter value in their 95% highest density posterior interval. All parameters met or exceeded this criteria except the fatigue parameter in reward pursuit, for which 94% of the recovered posterior distributions contained the true parameter value. See Supplemental Figure 2.S2 for details.

Model validation: model recovery

As a test of model selection validation, we simulated four separate data sets for each of the following models: SVF, SVFP, SVFP SperE, and the final model SVFP SperE Persev Psuccess. Each of these 16 simulated data sets was then hierarchically fit by all four models, and we assessed which model fit each data best based on the lowest WAIC measure across models for that data. This tests the ability of the fitting procedure to correctly identify the model which generated the data as the "best fit" for that data. As illustrated in Supplemental Figure 2.S3, the hierarchical sampling method was able to successfully recover each model well.

Results

Bifactor modeling of latent subdimensions of anxiety and depression

In order to tease apart variance common to subclinical levels of anxiety and depression from lower-level specific latent dimensions capturing different unique aspects of anxious and depressive symptom profiles, we performed dimensionality reduction using bifactor modeling across a number of self-report questionnaires. Previous research has shown that bifactor modeling of item level responses provides a valuable characterization of shared versus unique variance in highly comorbid disorders of internalizing psychopathology (Brodbeck et al., 2011; Clark et al., 1994; Reise, 2012; Reise et al., 2010; Simms et al., 2008; Steer et al., 1998; Steer and Clark, 1995). In our own prior work on learning in contingency volatility paradigms, we applied a bifactor model with 2 specific factors to items from a battery of anxiety and depression questionnaires to examine whether deficits in learning under volatility were linked to symptom variance common to anxiety and depression versus specific to anxiety or specific to depression (Gagne et al., 2020). In that study the hypotheses were specific to characterization of the shared versus distinct variance between measures of anxiety and depression. Crucially, the higher-order factor in a bifactor modeling paradigm provides a quantitative characterization of the shared transdiagnostic variance which may be attributable to condition comorbidity rather than unique diagnosis or syndrome.

In the current study we broadened the questionnaire set previously used by Gagne et al. and explored bifactor models with higher numbers of specific factors with the aim of differentiating latent sub-dimensions of depressive and anxious affect. The measures used were as follows: Beck Depression Inventory (BDI; (Beck, 1961), Penn State Worry Questionnaire (PSWQ; (Meyer et al., 1990), Mood and Anxiety Symptom Questionnaire (MASQ; anhedonic depression and anxious arousal subscales; Clark and Watson, 1991), the State Trait Anxiety Inventory trait (STAI-Trait) questionnaire (which itself comprises both depressed and anxious subscales; Spielberger et al., 1983, Spielberger 1989), the Apathy and Motivation Index (AMI; (Ang et al., 2017) behavioral subscale, and the Temporal Experience of Pleasure (TEPS; (Gard et al., 2006) anticipatory anhedonia subscale. Items evaluating suicidality were removed for online administration; total number of response items = 110.

453 participants completed the item response set online for use in exploratory bifactor analyses (see Methods for details); an additional 252 participants completed the same set of questionnaires online for confirmatory factor analyses. The winning bifactor model was then applied to the same item responses for the 50 in-lab behavioral participants to confirm construct consistency.

Multiple exploratory bifactor analyses were performed specifying one general factor and varying the numbers of specific factors from two through seven. The Schmid-Leiman orthogonalization method was used to compute bifactor item loadings (Schmid & Leiman, 1957). Given the ordinal nature of item responses, polychoric correlations were used in all factor analyses. A comparison of the Bayesian Information Criterion (BIC) (Schwarz, 1978), which assesses model fit while attempting to penalize for complexity, showed continued improvement in exploratory model fit as additional specific factors (up to the maximum seven specific factors considered) were included (see Figure 2.1a). We assessed the fit of each resultant model to the confirmatory data set using comparative fit index (CFI) and Tucker-Lewis measure of fit (see Figure 2.1b and Methods for details). Increasing the number of specific factors above five resulted in suboptimal confirmatory fit in the held-out data set, indicating that more complicated models were potentially overfit to the exploratory data.

The bifactor model with five subfactors (BF 5 model) provided the best combination of exploratory model fit (BIC=4468.38) while maintaining a CFI above 0.95 in the held-out data (CFI = 0.953, RMSE = 0.101; see Figure 2.1). This model also provided broadly interpretable specific factors. In the BF 5 model the general (higher-order) factor loaded highly across most depression and anxiety items, as anticipated. The first specific factor (F1) loaded strongly on items associated with anxious physiological arousal and, to a lesser extent, both somatic and cognitive symptoms of depression. Specific factor 2 (F2) loaded primarily on items pertaining to anhedonic depression. Specific factor 3 (F3) loaded primarily on items related to worry and cognitive aspects of anxiety. Specific factor 4 (F4) loaded primarily onto items associated with behavioral apathy, and specific factor 5 (F5) loaded primarily onto items associated with anticipatory hedonia. (see Figure 2.1c for summary).

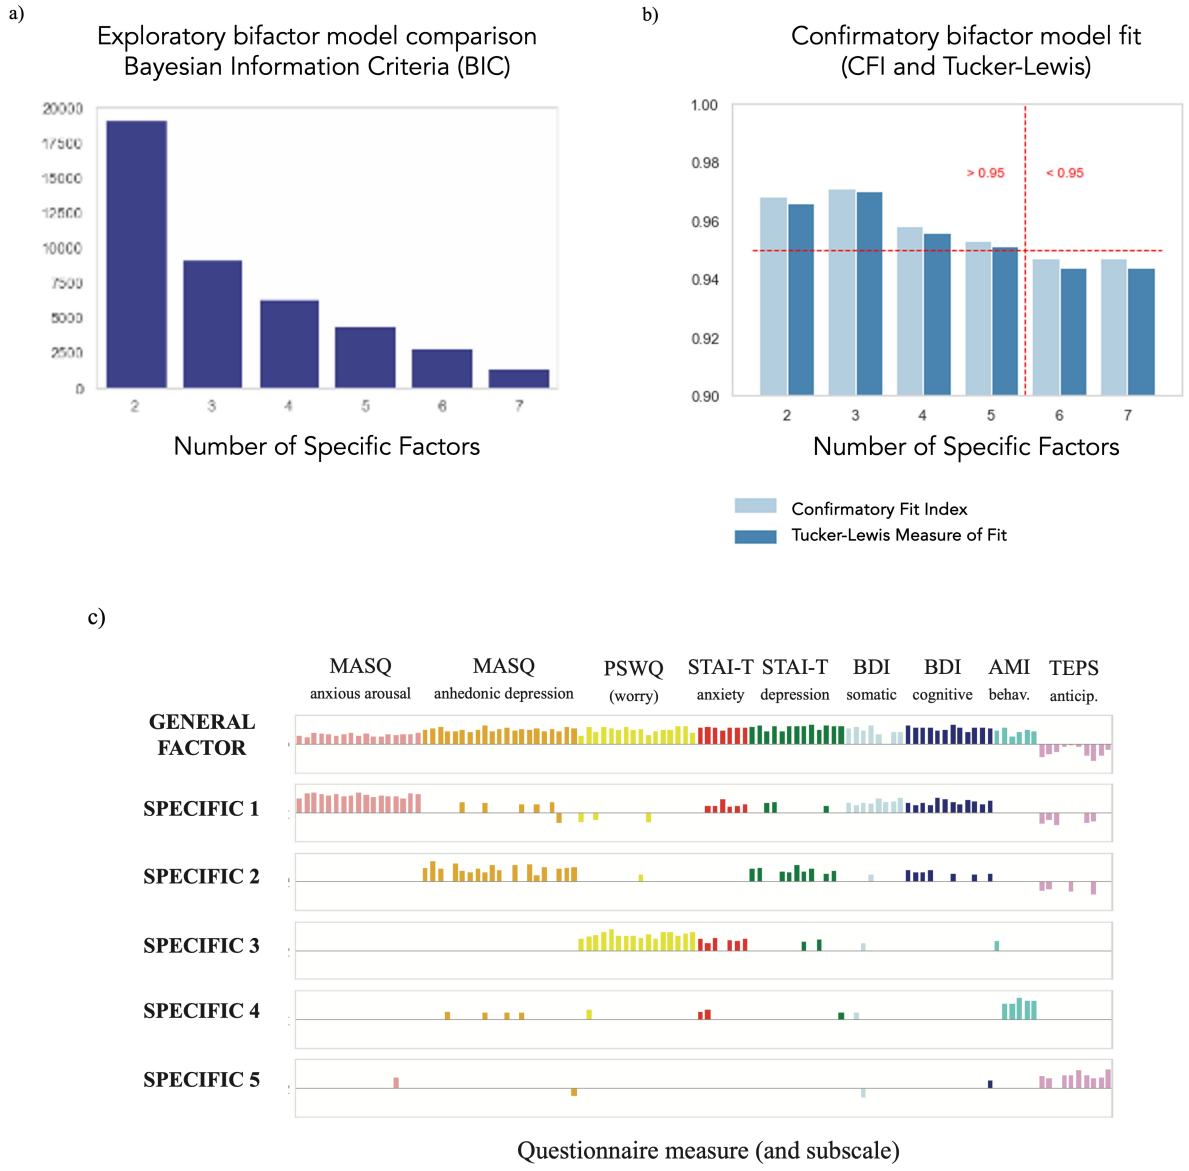


Figure 2.3 Bifactor analysis of item response data. a) Bayesian Information Criterion (BIC) showed improved model fit to exploratory data set as number of specific factors in bifactor model increased. b) Comparative fit of each bifactor model in confirmatory data showed highly credible fit > 0.95 for up to 5 specific factors c) Item loadings of final bifactor model with one general and 5 specific factors. Questionnaires from left to right: Mood and Anxiety Symptom Questionnaire (Anxious Arousal and Anhedonic Depression subscales); Penn State Worry Questionnaire; State Trait Anxiety Inventory- Trait (Anxiety and Depression subscales); Beck Depression Inventory (Somatic and Cognitive subscales); Apathy Motivation Index (Behavioral subscale); Temporal Experience of Pleasure (Anticipatory subscale).

The item loadings of the BF 5 bifactor model were used to calculate factor scores for each participant in the exploratory data set using the Anderson-Rubin method (Anderson and Rubin, 1956; see Methods for details). We then examined the correlation of participant factor scores from the BF 5 model with participant scores on associated questionnaire subscales (see Figure 2.4) to

provide additional insight into factor interpretability. As anticipated, we observed correlations of the general factor with most of the subscales, and high correlations of each specific factor with one or more questionnaire subscale scores.

We performed the same correlational analysis between subscale scores and BF 5 model participant factor scores for the in-lab behavioral participants to ensure consistency in factor interpretation. The correlations in the behavioral group were highly similar with the analogous correlation structure of the exploratory data group. Since all model-based inference was performed on the behavioral participants, we used the correlational analysis from this group to characterize the specific factor descriptions used throughout the analysis.

F1 showed the strongest correlation with scores on the MASQ anxious arousal (AA) subscale in the behavioral group (Pearson correlation $r=0.86$, $p=3.53\text{x}10^{-16}$; coefficient of determination = 0.74, indicating over 70% common variance between Factor 1 and the MASQ AA subscale). F1 also showed substantially weaker relationships with the BDI somatic ($r=0.42$, $p=2.15\text{x}10^{-3}$) and cognitive ($r=0.40$, $p=3.66\text{x}10^{-3}$) subscales. We therefore characterized F1 as capturing anxious arousal/depression.

F2 was highly related to anhedonic depression, with a strong correlation with the MASQ AD subscale ($r=0.68$, $p=3.77\text{x}10^{-8}$), and lesser correlations with the STAI depression subscale ($r=0.36$, $p=9.99\text{x}10^{-3}$) and the TEPS hedonia scale ($r = -0.25$, $p=0.076$; note that TEPS measures positive hedonia, so a negative correlation with this scale indicates greater anhedonia in the factor). This factor was characterized as capturing anhedonic depression.

F3 primarily captured worry and cognitive aspects of anxiety, mainly as assessed by the PSWQ ($r=0.74$, $p=4.20\text{x}10^{-10}$) and to a lesser extent by STAI anxiety items ($r=0.40$, $p=3.72\text{x}10^{-3}$), and was characterized as capturing worry.

Participant scores on F4 were strongly related to scores on the AMI ($r=0.81$, $p=8.22\text{x}10^{-13}$), and participant scores on F5 to scores on the TEPS ($r=0.90$, $p=3.70\text{x}10^{-19}$). These factors were characterized as capturing behavioral apathy (F4) and anticipatory anhedonia (F5_neg). Since the TEPS scale is designed such that higher TEPS relates to higher hedonic response, and all other factors positively measured higher levels of the symptoms of interest, we used negative F5 scores in our analyses so that increases in the factor score F5_neg were related to increases in anhedonia. These factor labels are subjective characterizations intended to reflect the predominant dimensional aspects of each specific factor after shared general variance has been removed.

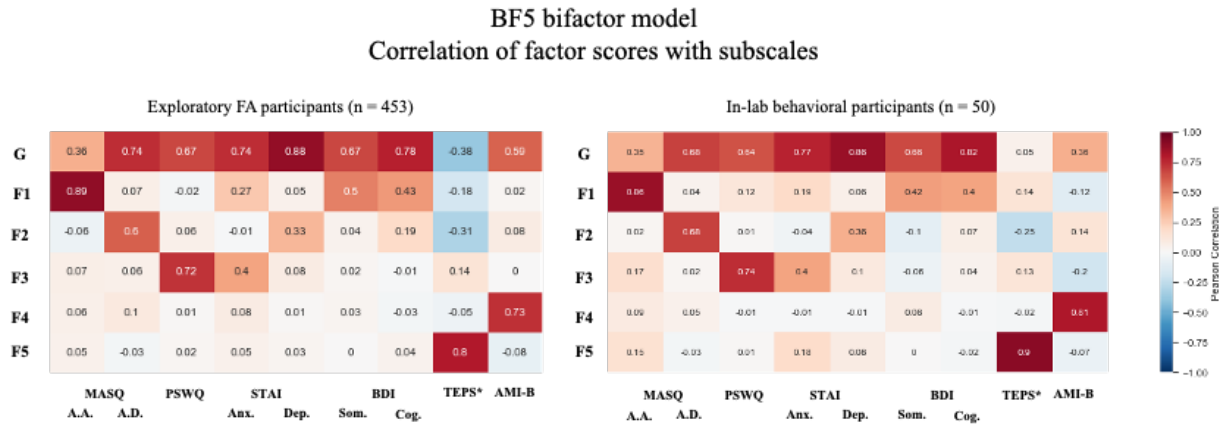


Figure 2.4 Correlation of questionnaire subscale scores with participant-level bifactor general and specific factor scores in exploratory (online) data participants and in-lab behavioral task participants. Questionnaires from left to right: Mood and Anxiety Symptom Questionnaire (Anxious Arousal and Anhedonic Depression subscales); Penn State Worry Questionnaire; State Trait Anxiety Inventory- Trait (Anxiety and Depression subscales); Beck Depression Inventory (Somatic and Cognitive subscales); Apathy Motivation Index (Behavioral subscale); Temporal Experience of Pleasure (Anticipatory subscale).

*TEPS scale is designed such that higher scores indicate higher levels of hedonic response. Unadjusted factor loading correlations are shown here. For modeling and analysis, participant factor scores on F5 were reversed so that higher scores on F5_neg corresponded to higher anhedonia in order to simplify interpretability of results.

Behavioral task: Effort-based decision-making task

To explore individual differences in effortful decision-making, we used an experiment in which participants were asked to choose whether to engage in a series of effortful trials either to pursue reward or to avoid threat. We adapted a previously used "apple-picking" task (Bonnelle et al., 2015a, 2015b) which translates effort exerted through squeezing a dynamometer (handgrip) device into visually "climbing" the trunk of a tree in real time as force is exerted (see Figure 2.3). In addition to the reward version of the task, in which subjects pursue reward by climbing the tree to pick varying numbers of apples (converted to monetary reward), we created a structurally equivalent threat version of the task, in which subjects avoid threat by climbing the tree to avoid varying numbers of bear traps (corresponding to levels of electrical stimulation).

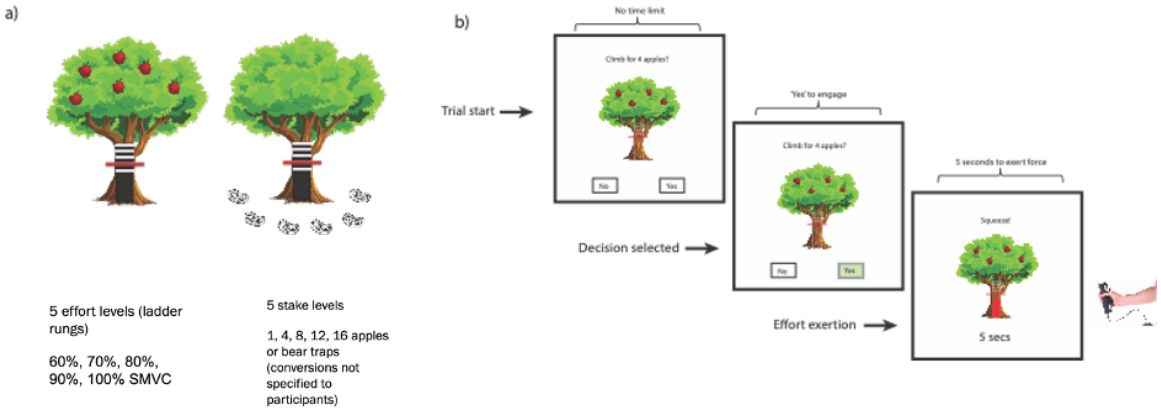


Figure 2.5 Effortful decision-making behavioral task. Task presented participants with 156 trials each of effortful reward pursuit (apples, converted to monetary reward) and effortful threat avoidance (bear traps indicating intensity of electrical stimulation). a) 5 possible stake factors were used in each paradigm illustrated by 1, 4, 8, 12, or 16 apples or bear traps. 5 effort levels were represented by ladder rungs, with the highest rung calibrated to 85% of each individual's maximum force grip (sustainable maximal voluntary contraction, or SMVC). b) Each trial presented participants with a combination of stake level and effort level. Participants could participate in the trial and exert effort to attempt to win the trial (gain the apples in reward paradigm, escape the bear traps in threat paradigm). If they declined the trial, they either received 0 reward or received the electrical stimulation indicated according to paradigm.

The number of apples or bear traps, respectively, indicated the relative magnitude of monetary gain or electrical stimulation for the trial, while the red highlighted ladder rung on each trial indicated the effort required to succeed. There were 5 levels of possible stake and 5 levels of possible effort requirement in each paradigm (reward pursuit and threat avoidance). Required effort levels were calibrated such that the highest effort level of the task represented 85% of each individual participant's maximum grip strength. Levels of electrical stimulation were calibrated such that the highest possible stimulation level during the task was rated 7 out of 10 in aversiveness by participants prior to the start of the task (see Methods). Participants completed the reward pursuit and threat avoidance tasks on different days and were counterbalanced for session valence order.

In both reward pursuit and threat avoidance, a logistic regression of participant choice behavior showed choices were significantly modulated by trial-wise stake and effort level, with increasing stake levels increasing the probability of engaging in trials (reward pursuit: stake level coefficient=1.13, 95% confidence interval (CI)= [1.075,1.186]; threat avoidance: stake level coefficient =0.774, 95% CI=[0.733,0.815]) and increasing effort levels decreasing the probability of engaging in trials (reward pursuit: effort level coefficient= -0.526, 95% CI=[-0.571,-0.480]; threat avoidance: effort level coefficient= -0.430, 95% CI=[-0.468,-0.391]). See Figure 2.6. This choice pattern is consistent with that reported by prior studies of effortful reward pursuit using the original task version (Bonnelle et al., 2015a).

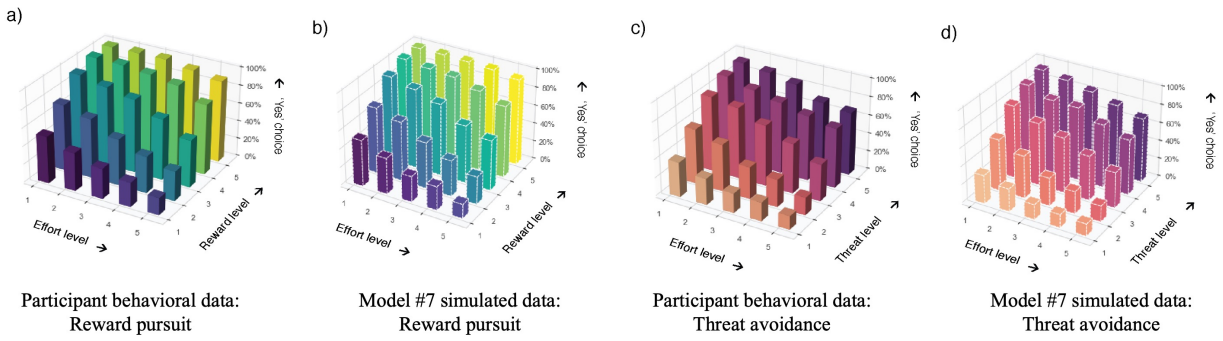


Figure 2.6 Behavioral choice results from participant data and winning computational model. a) Actual mean participant choice patterns (probability of trial engagement) in threat avoidance as a function of effort level and stake (threat) level); b) simulated threat avoidance choice data from winning computational Model #7. c) Actual mean participant choice patterns (probability of trial engagement) in reward pursuit as a function of effort level and stake (reward) level); d) simulated reward pursuit choice data from winning computational Model #7.

Computational modeling of effort-based decision-making behavior

We next sought to model task performance computationally in order to tease apart the influence of different specific decision-making components on participant choice. For this, we fitted participant choice data with various subjective value models as detailed below. For each model, we used a hierarchical Bayesian sampling approach to infer distributions over parameters at the group level and the individual level (see Methods for details).

Since our fundamental hypotheses were regarding the effect of individual differences in latent dimensions of negative affect on effortful choice behavior, we included parameters to assess the influence of participant scores on each factor (including the general factor describing shared variance) in all models. We also similarly parameterized the influence of the level of mV stimulation rated as "Level 1" (L1mV) by each participant in all threat paradigm parameters to control for any variance associated with baseline physiological differences in pain sensitivity, which (although itself of potential interest in relation to dimensions of negative affect) was not the focus of our analysis. Separate reward and threat parameters were jointly inferred during sampling.

We assessed the fit of each model to the data using the Watanabe-Akaike Information Criteria (WAIC) (Watanabe S, Opper M, 2010), a method which assesses comparative model fit while penalizing for model complexity in a hierarchical Bayesian framework (see Methods).

We examined the 95% highest density interval (HDI) of the posterior distribution for each factor score weight (each denoted as a specific "β" parameter) in the winning model. Posterior 95% HDI intervals for factor score weights β which excluded 0 were interpreted as being indicative of significant modulation by the given factor score on the given cognitive parameter (see Methods for further details).

The simplest subjective value model (the SV baseline model, Model #1) included a parameter which linearly scaled effort magnitude, a parameter which linearly scaled outcome magnitude, and a parameter which quantified an individual's overall baseline bias toward avoiding engagement

across trials, as shown in equations 1a and 1b. For notational clarity, the key model formulas shown in 1a and 1b below exclude the extended notation for separate reward and threat indicators and for β-weight group-level offsets (see Methods for full notational explication).

The subjective net value (SNV) of the offer on trial *t* was calculated as a function of the tradeoff between stake and effort:

$$SNV_t = \gamma \cdot Stake_t - \kappa \cdot Effort_t \quad 9(a)$$

Additional model parameters not directly involving stake or effort level were included as adjustments to the value in the calculation of the probability that a participant chose to engage on the trial:

$$P(Yes)_t = (1 + \exp(-[\, SNV_t - b\,]))^{-1} \quad 9(b)$$

Note that the placement of the parameters is a matter of chosen convention; for example, the baseline offset parameter *b* could be moved from the choice probability formula in 1b to the SNV calculation in 1a without changing the model in any way. We chose to characterize SNV as relating only to stake level and effort level for conceptual clarity.

The inferred participant-level parameters for Model #1 were $\gamma_{_reward}$ and $\gamma_{_threat}$ (describing the linear weighting of stake in reward and threat choices, respectively); $\kappa_{_reward}$ and $\kappa_{_threat}$ (describing the linear weighting of effort in reward and threat choices, respectively); and $b_{_reward}$ and $b_{_threat}$ (describing the baseline bias against engaging in trials for reward and threat choices, respectively). Model #1 therefore included inference over 6 participant-level parameters.

The group-level mean and standard deviation for each of these parameters was inferred jointly with the participant-level parameters, as were the adjustment weights β for factor scores and baseline level of electrical stimulation (L1mV; threat valence only) (see Methods). Model #1 therefore additionally included inference over 51 group-level hyperparameters as follows: group-level μ and σ for each of the 6 participant-level parameters; group-level offsets β for each BF5 factor score (G, F1, F2, F3, F4, and F5_neg) applied to each group-level distributional mean μ; and group level offsets β for baseline electrical stimulation level (L1mV) applied to each threat avoidance parameter group-level distributional mean μ. Non-informative Gaussian or exponential priors were set for all hyperparameters (see Methods for details).

The WAIC for Model #1 was 7,237.58 with standard error (SE) = 127.64. This was used as a basis for comparison of model fit for subsequent models, with lower WAIC indicating better fit to the data.

Following Model #1, we considered additional variables that might affect the subjective value computation, and fitted a number of additional models, increasing the complexity by including one additional participant-level parameter at each stage (see Methods for additional modeling details).

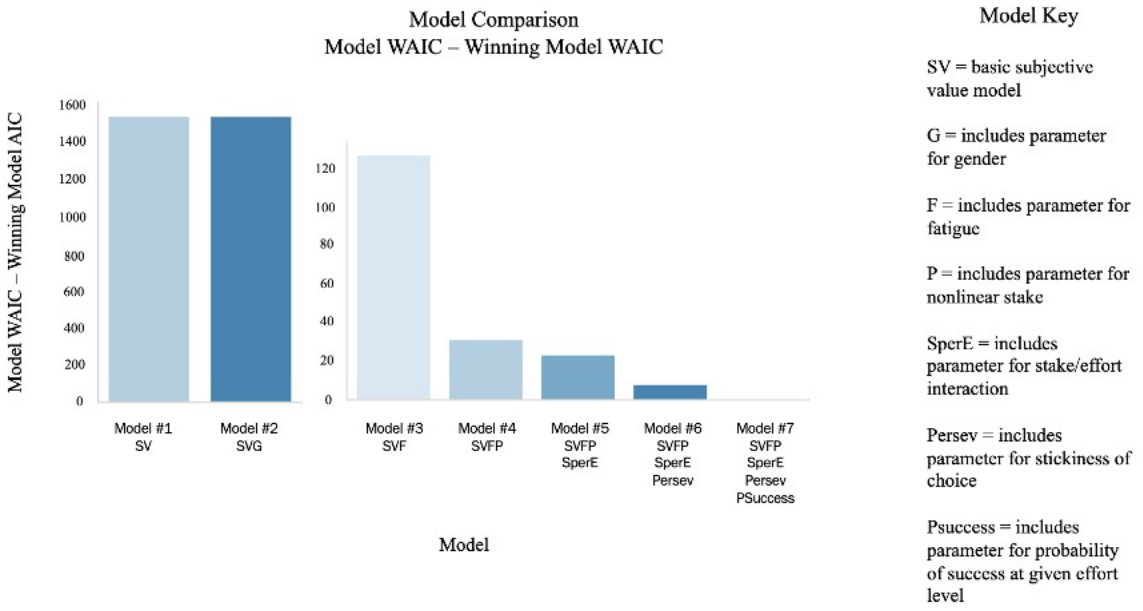


Figure 2.7: Model fit comparison WAIC. Model fit comparison of Watanabe-Akaike Information Criterion (WAIC) for each model relative to winning model (lowest WAIC). Parameter space and model descriptors for each model outlined. See Methods for modeling details.

The winning model (Model #7), given in equations 10(a) and 10(b) below, included the following components in addition to those of Model #1: a nonlinear scaling parameter on stake (θ) as suggested by utility theory (Von Neumann, J., & Morgenstern, O., 1944); a parameter to quantify fatigue across trials (f); a stake per effort interaction parameter (δ); a parameter to account for a participant's previous aggregate probability of success on 'yes' responses at a given level of effort (ρ); and a parameter to characterize a participant's tendency to perseverate (i.e., to repeat the previous response regardless of changes in trial offer, sometimes referred to as 'sticky choice') (η).

The winning model is described by the following equations. As before, for notational clarity the key model formulas shown in 10(a) and 10(b) exclude the notation for separate reward and threat indicators and for β-weight group-level offsets, and all parameters not relating directly to stake or effort level are included in formula 10(b) by convention.

$$SNV_t = \gamma(\text{Stake}_t)^{\theta} - \kappa(\text{Effort}_t) + \delta(\text{Stake}_t/\text{Effort}_t) \quad 10(a)$$

$$P(\text{Yes})_t = (1 + \exp(-[SNV_t - f(\text{Trial}_t) + \rho(\text{Psuccess}_{t,e}) + \eta(\text{Choice}_{t-1}) - b]))^{-1} \quad 10(b)$$

Model #7 had the lowest WAIC of all models analyzed; WAIC = 5709.68, se = 119.71.

Multiple influences of negative affect upon effortful reward pursuit

We first tested the relationship of individual differences in latent specific negative affect factor scores with parameters of the winning computational model describing effortful reward pursuit.

Population-level weights β captured the influence of participant factor scores on each parameter of the winning model (Model #7). We examined the posterior distributions of these population-level weights to investigate whether factor scores led to a significant shift in the group distributional mean of participant-level parameters (Figure 2.2).

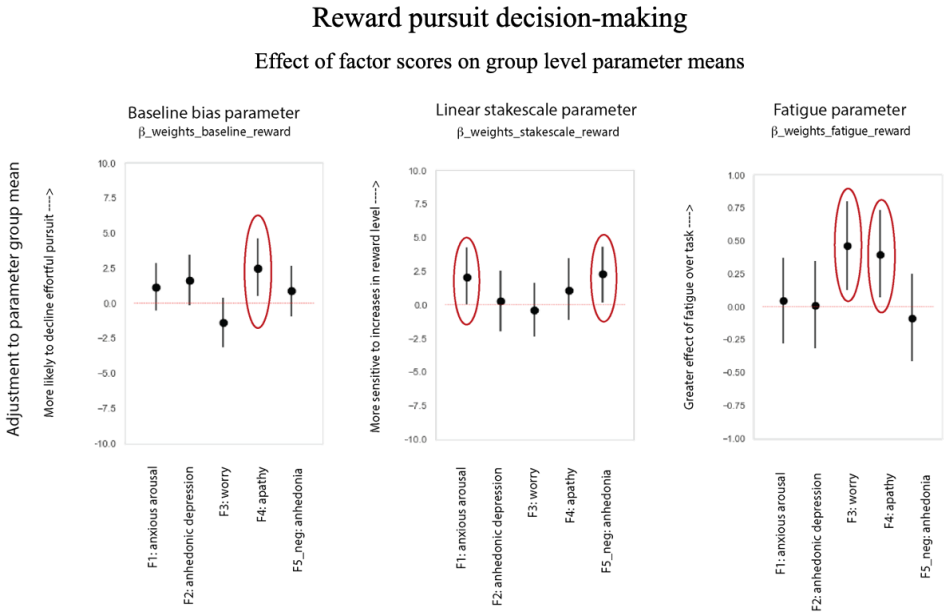


Figure 2.8: Effect of factor scores on decision-making parameters in reward pursuit for winning computational model of choice. Vertical lines show 95% posterior highest density intervals (HDI) for selected group-level trait effect parameters. Increases in F4 (behavioral apathy) are associated with an increased baseline bias to refuse effortful reward pursuit across all effort and reward levels. Increases in F5_neg (anticipatory anhedonia) are associated with a greater sensitivity to increases in stake (higher levels of reward are generally required for this group to exert effort). A similar effect is shown for F1 (anxious arousal). Increases in F4 (behavioral apathy) are associated with an increased impact of fatigue on choice refusal over time. A similar effect is shown for F3 (worry)

We found that elevated scores on factors capturing heightened behavioral apathy (F4) and anticipatory anhedonia (F5_neg) were significantly associated with reduced effortful reward pursuit, in line with previous findings. However, each operated on distinct mechanistic components of decision-making.

Specifically, we found that higher levels of behavioral apathy (F4) were associated with a significantly greater bias against engaging in effortful reward pursuit ($\beta_{_F4_baseline_reward} = 2.48$, 95% posterior HDI=[0.572, 4.681]; see Figure 2.8). Higher levels of behavioral apathy (F4) were also associated with a significantly larger drop in effortful engagement over time as indexed by trial number ($\beta_{_F4_fatigue_reward} = 0.398$; 95% posterior HDI = [0.073, 0.735]). In line with some previous findings, this might be interpreted as individuals with higher levels of behavioral apathy showing greater fatigue over time during effortful reward pursuit (Bonnelle, 2016), but we note that a model parameter indexed by trial number could also reflect other sources of decreased task engagement over time. Interestingly, variance in apathy (F4) was not significantly associated with a differential response to changes in the level of effort required or to the level of reward that could be obtained.

In contrast, anticipatory anhedonia as measured by F5_neg was associated with elevated sensitivity to increases in reward level on offer ($\beta_{_F5neg_stakescale_reward}$ = 2.28, 95% posterior HDI=[0.158, 4.293]) but showed no significant relationship with baseline, effort level, or fatigue parameters. This elevated sensitivity to increases in reward level corresponded behaviorally to individuals with higher levels of anhedonia (higher F5_neg) generally requiring higher levels of reward to begin engaging in effortful pursuit relative to controls. Jointly, these results indicate distinct mechanisms of operation for different latent dimensions of depressive symptoms in effortful reward pursuit.

Like higher levels of behavioral apathy (F4), higher levels of worry (F3) were also associated with a significantly larger drop in effortful engagement over time as indexed by trial number ($\beta_{_F3_fatigue_reward}$ = 0.464, 95% posterior HDI=[0.123, 0.801]; see Figure 2.8). This may indicate that latent dimensions of anxiety (in this case, cognitive anxiety) are also related to mechanistic differences in effortful reward pursuit decisions.

We also found that high scores on F1, which taps physiological symptoms of anxiety and to a lesser extent depression, were significantly associated with heightened sensitivity to increases in reward level ($\beta_{_F1_stakescale_reward}$ = 2.06, 95% posterior HDI=[0.038, 4.267]).

Effortful avoidance of threat: heightened bias but reduced sensitivity in worriers

We next examined the factors influencing effortful threat avoidance. In line with the hypothesis that anxiety-related components of negative affect might be associated with increased willingness to exert effort to avoid threat, elevated scores on the dimension associated with worry (F3) were significantly associated with an increased baseline willingness to engage in effortful avoidance of shock receipt ($\beta_{_F3_baseline_threat}$ = -2.304, 95% posterior HDI = [-4.585, -0.198]; note parameter describes baseline avoidance of effort and the adjustment here is negative, indicating increased baseline willingness to exert effort). Interestingly, we found that heightened worry (F3) was simultaneously significantly associated with a reduced linear scaling of threat level ($\beta_{_F3_stakescale_threat}$ = -3.112, 95% posterior HDI=[-5.948, -0.284]), indicating a flattened relationship between increasing threat levels and increasing probability of willingness to exert effort in individuals with higher worry (see Figure 2.9).

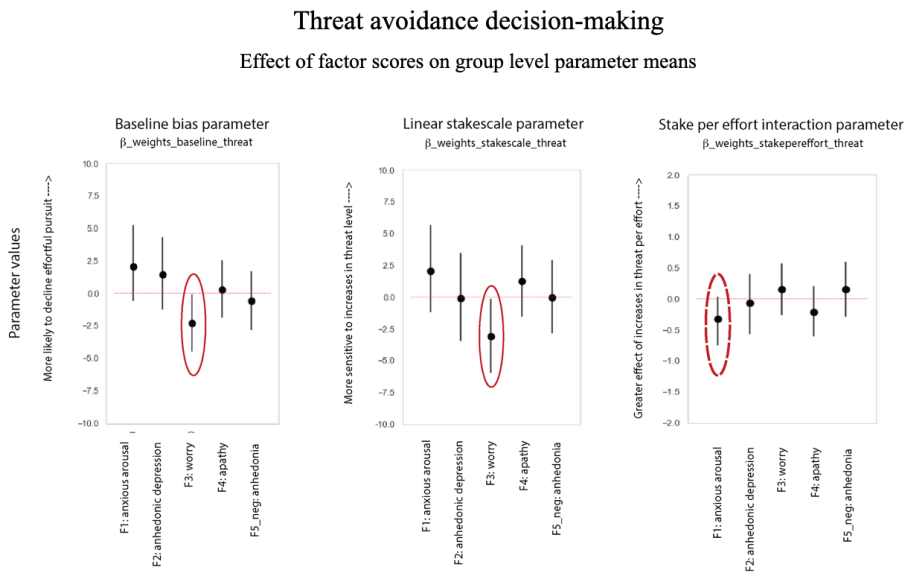


Figure 2.9: Effect of factor scores on decision-making parameters in reward pursuit for winning computational model of choice. Vertical lines show 95% posterior highest density intervals (HDI) for selected group-level trait effect parameters. Increases in F3 (worry) are associated with a reduced baseline bias against effortful threat avoidance for all effort and reward levels (i.e., higher F3 is associated with higher broad willingness to exert effort to avoid threat). Increases in F3 (worry) are also associated with a decreased sensitivity to increases in stake (these individuals are less influenced by increases in threat magnitude). We find a trending relationship between increases in F1 (anxious arousal) and reduced influence of threat level per effort level

No other latent dimensions of negative affect modulated any of the computations involved in effortful threat avoidance as assessed in this paradigm.

Discussion

Decisions to effortfully pursue reward or avoid threat can be broken down into distinct component computational mechanisms, and latent dimensions of psychopathology may differentially impact these distinct mechanisms. In the current study we sought to investigate the relationships of key latent subdimensions of anxiety and depression with the mechanistic components of effortful decision-making processes.

We used bifactor modeling of item level scores from a battery of standardized questionnaire measures of anxiety and depression administered to a large on-line participant group to characterize a general latent factor that captured variance common to anxious and depressive affect and five specific latent factors that captured unique orthogonal components of variance in anxious and depressive affect. Factor loadings were used to relate in-lab participant scores on these latent factors to performance on effortful reward pursuit and threat avoidance tasks. We focused on the five specific factors and exploration of three key questions. First, we sought to test whether different latent subdimensions of depressive affect impacted the same or different component mechanisms of decision-making in effortful reward pursuit. Second, we sought to investigate whether cognitive or physiological latent subdimensions of anxious affect were associated with differences in decision-making in either effortful reward pursuit or effortful threat avoidance, and

if so whether these associations varied between decision-making in reward versus threat paradigms. Finally, we more broadly sought to assess whether anxious versus depressive affect differentially modulated effortful decision-making in reward pursuit or threat avoidance.

Our first main question sought to address the relationship of specific dimensions of depressive affect with the component mechanisms underlying effortful reward-based decision-making. A commonly used task in human effortful reward pursuit studies is the Expenditure of Effort for Reward Task (EefRT) (Treadway 2009, 2012), which compares choice preferences between high effort/high reward options and low effort/low reward options, while simultaneously varying the levels of risk associated with decision option outcomes. This task has the advantage of allowing translational comparisons from rodents to humans, but one disadvantage is that it cannot give us precise information as to whether reduced willingness to choose high effort/ high reward options is linked specifically to altered valuation of effort or of reward. To address this issue in the current study, we orthogonally titrated 5 levels of effort and 5 levels of stake (level of financial reward or electrical stimulation) in reward pursuit and threat avoidance. We used hierarchical Bayesian modeling of subjective value decision-making to model task behavior, enabling us to dissociate altered valuation of effort from that of stake, while also assessing baseline biases, changes in willingness to engage across trials, and other influences on choice. The two latent factors from the bifactor model of particular interest here were the fourth and fifth factors which showed high loadings on items from the behavioral apathy and anticipatory hedonia scales, respectively.

While scores on both these factors impacted willingness to exert physical effort in reward pursuit, our findings indicated that they acted upon distinct computational mechanisms. Specifically, higher scores on latent factor F4, associated with increased behavioral apathy, were linked to reduced across-the-board baseline willingness to exert effort in pursuit of reward, as well as a greater diminishment of effort exertion in reward pursuit over the course of the task. Meanwhile, higher scores on latent factor F5_neg, associated with elevated anticipatory anhedonia, were linked to heightened sensitivity to increases in reward magnitude, which behaviorally reflected a reduced willingness to exert effort for rewards of low magnitude.

These findings illustrate how distinct latent dimensions of depressive affect can influence effortful reward pursuit through different computational mechanisms. Given that anhedonia and apathy often occur together in the same individuals in disorders such as depression, we might expect to see a compound effect in such individuals of both a reluctance to engage in effortful reward pursuit in general that increases with the time spent on the pursuit in hand and reduced willingness to pursue rewarding activities unless the reward magnitude is perceived to be of especially high subjective value. Importantly, identifying the specific component mechanism impacted may help in the development of personalized interventions. For example, the current findings might suggest that an individual with high anticipatory anhedonia might benefit most from encouragement to revalue activities seen as of low reward value or to test how they feel after engaging in such activities. In addition, establishing the extent to which pharmacological interventions increase baseline willingness to engage, decrease the drop off in willingness across time or alter the change in weight given to activities seen as mildly, moderately or extremely pleasurable would also be of value.

Our second main aim was to investigate the role of latent dimensions of anxious phenotypes in effortful reward pursuit and threat avoidance. Few studies have explored the relationship of anxiety with effortful reward pursuit. Our analysis found that scores on latent factor F3, which showed high loadings for worry-related symptoms, were linked to a significantly greater reduction in effortful reward pursuit over time. This relationship was similar in strength and direction to the effect of heightened behavioral apathy, presenting the intriguing possibility that cognitive dimensions of anxious phenotypes may have a similar effect as certain dimensions of depressive phenotypes on effortful reward pursuit. However, this similarity also raises a more complex question which our present analysis cannot answer: while worry (F3) and behavioral apathy (F4) acted similarly on the model parameterization of decreasing willingness to effortfully pursue reward over time (as indexed by trial number), it is possible that this parameter captured not only the effects of fatigue from effort exertion (as was intended), but also effects of various other sources of decreasing task engagement over time, such as compromised sustained attentional abilities. An effortful reward pursuit investigation which additionally measures ongoing attention and motivation throughout the task could be used to more fully explicate the underlying causes of the similar mechanistic impairment seen here for orthogonal dimensions of negative affect strongly linked to worry (F3) and behavioral apathy (F4).

In the effortful threat avoidance task, the latent factor with high loadings on worry-related symptoms (F3), was the only latent factor to significantly modulate parameter values. Scores on this factor were associated both with an increased baseline willingness to effortfully avoid threat, and decreased sensitivity to increases in threat level.

These findings for threat avoidance are consistent with work from animal research revealing that in rodents, heightened anxious-like states are associated with greater willingness to exert effort to avoid threat (as indexed by longer swim times on the forced swim test) (Cryan et al., 2005; Lee et al., 2017; Porsolt et al., 1977; Ramboz et al., 1998). Studies using instrumental threat avoidance paradigms in humans have found that individuals with higher levels of anxiety engage in greater avoidance of threat (Davis et al., 2010; Grillon et al., 2008), but these studies have typically not experimentally modulated the effort level required or attempted to differentiate the correlates of cognitive versus physiological anxiety. We note that elevated scores on the factor tapping both anxious arousal and depression (F1, see Figure 2.9) did not significantly modulate effortful avoidance of threat. To check that the absence of a finding for physiological anxiety was not merely due to the mixed loadings on this factor, we replaced scores on this latent factor with scores on the MASQ-AA in our computational behavioral model. There was no significant impact of MASQ-AA scores on any of the parameters of interest (see Supplements), suggesting that it is indeed the cognitive aspect of anxiety that most modulates willingness to exert effort in pursuit of threat.

Our third main aim was to determine broadly whether anxious versus depressive affect differentially modulated effortful decision-making in reward pursuit and threat avoidance.
When we consider the full picture across both effortful pursuit of reward and avoidance of threat, we see clear differential influences of depressive versus anxious phenotype-related symptom dimensions. High scores on latent dimensions tapping symptoms of apathy and anticipatory anhedonia were associated with different facets of reduced willingness to exert effort in pursuit of reward but did not modulate effortful avoidance of threat. In contrast, elevated scores on the latent

dimension tapping worry were linked with reduced effort in pursuit of reward as a function of time on task but increased effort in avoidance of threat.

Our initial hypotheses regarding effortful threat avoidance decision-making were informed by the rodent literature wherein anxious phenotypes have showed reduced inertia on the forced swim task in contrast to depressive phenotypes which showed increased inertia. While we found worry to be linked to increased willingness to exert effort in avoidance of threat we found no significant relationship between depressive symptoms and decreased willingness to exert effort in avoidance of threat. One possibility is that due to the ethical constraints in the levels of electrical stimulation used in our task with humans (where participants select the highest level they are willing to receive), our manipulation of threat was not as extreme as that used in the rodent paradigm and does not induce feelings of helplessness/hopelessness which might be present in situations of true mortal threat or that more closely parallel the rodent experience of the forced swim test.

Here it is also of interest to note that scores on the factor that tapped both anxious arousal and depressive affect (F1) showed a trend level negative relationship with parameter scores for the influence of stake per threat interaction in effortful threat avoidance. Since F1 loaded on both anxious arousal and depression items (largely from the MASQ anxious arousal subscale and the Beck Depression Inventory), we investigated the effect of replacing F1 with either standardized MASQ AA scores or standardized BDI scores. When F1 was replaced with BDI scores, we observed a significant relationship between BDI and the stake per effort interaction parameter, with elevated BDI scores being linked to a reduced influence of threat level increases relative to effort level in decision-making (see Supplements for analyses). No significant effect was found when F1 scores were replaced with MASQ AA scores. While we do not wish to give too much weight to this post-hoc analysis, it does align with the hypothesized opposite effects of anxious versus depressed symptom dimensions on effortful avoidance of threat.

In summary, our findings indicate that different subdimensions of depressed and anxious symptoms do indeed have distinct mechanistic influences on effortful decision-making in humans. Among individuals with heightened depressive affect, the symptom of apathy and the symptom of anhedonia appear to impact separate mechanisms of the decision-making process which each lead to a similar high-level behavioral outcome (reduced effortful reward pursuit). This may also hint at a potential compound effect when these symptoms co-occur in depression. We also find that cognitive worry is associated with a similar mechanistic impairment in effortful reward pursuit as apathy. While this suggests that dimensions of anxiety should also be investigated with respect to differences in effortful reward pursuit, it is also interesting to note that the dimension of cognitive anxiety captured by worry may also share high co-occurrence with depressive symptoms, and should therefore be accounted for during exploration of subdimensions of depressive symptoms in order to uniquely parse variance across other symptoms. Within the domain of effortful threat avoidance, we do not find a significant difference in behavior for any dimensional characterization of depression specifically. Interestingly, we find that the cognitive rather than physiological dimension of anxiety is associated with greater bias to effortfully avoid threat and less sensitivity to the actual magnitude of threat itself.

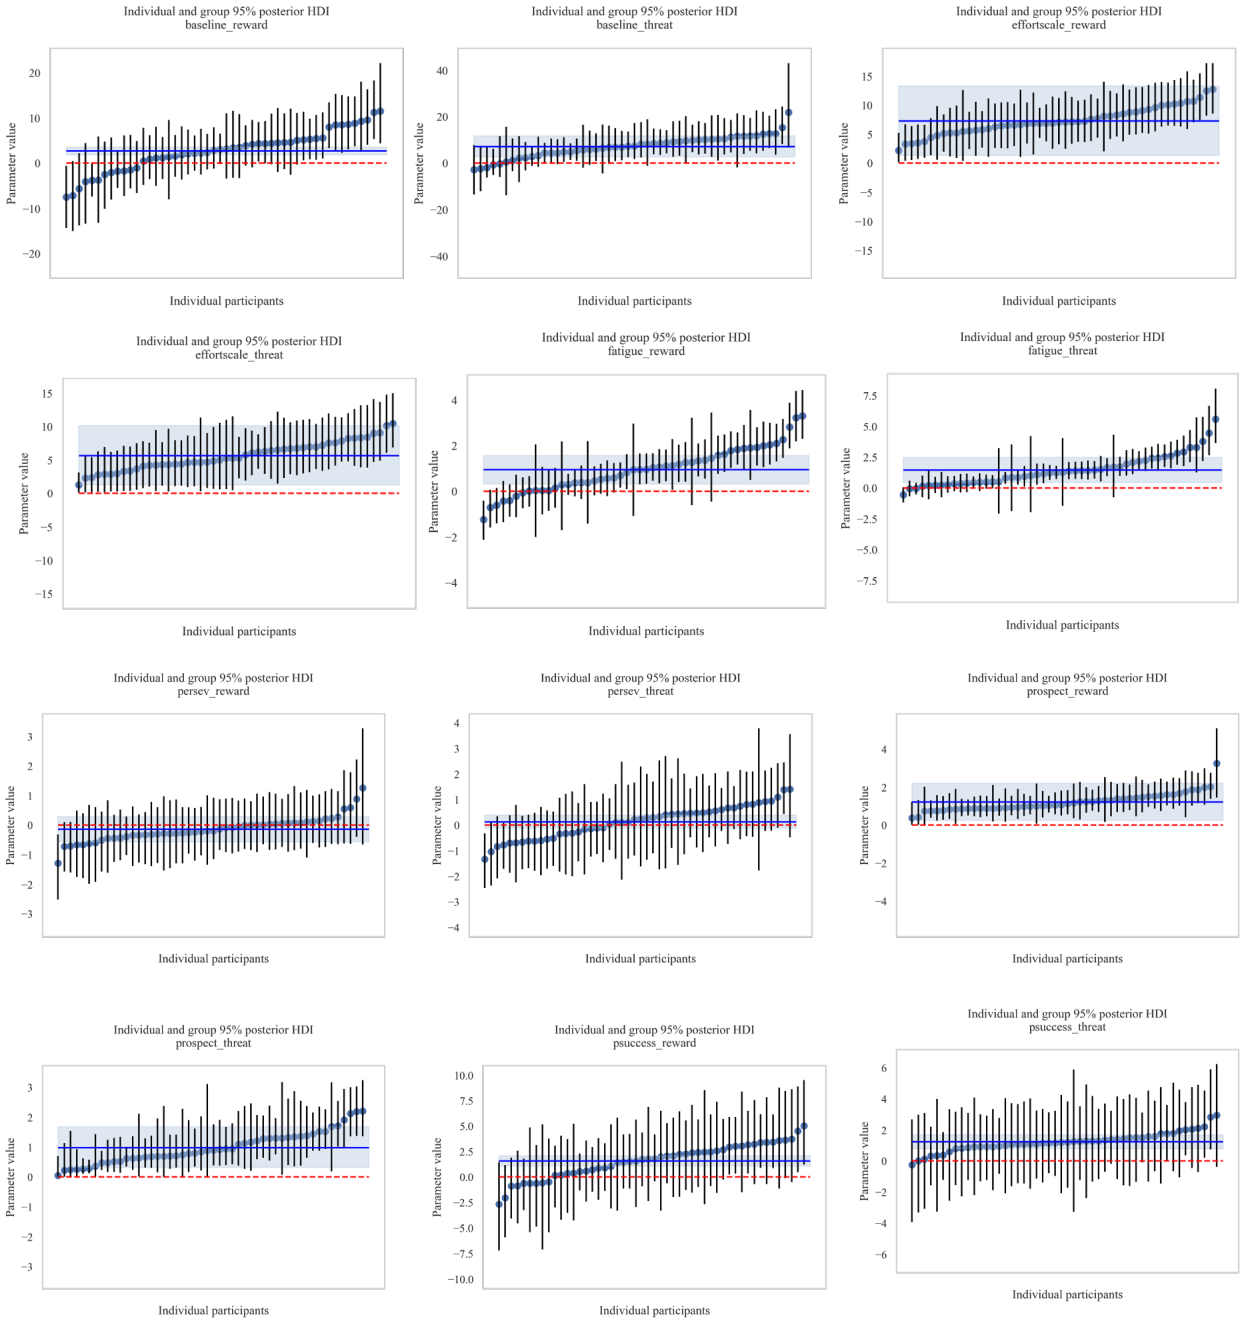
We hope that this delineation of the influences of different subdimensions of depressed and anxious affect on the component computational mechanisms involved in effortful decision-making

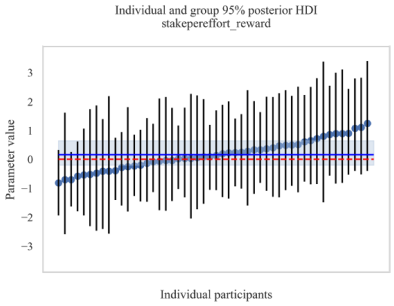
might valuably inform more personalized treatments for psychopathology. Currently, many heterogeneously presenting or comorbidly occurring forms of depression or anxiety are treated through a trial-and-error approach with a single set of psychological and pharmacological tools. More nuanced information about the relationship of specific symptom profiles with specific deficits in effortful decision-making may allow for treatments to be developed that are more precisely targeted to the manifestation of internalizing psychopathology in individuals.

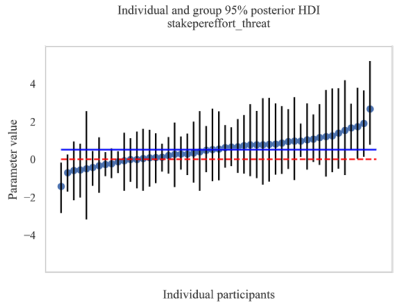
Chapter 2: Supplemental materials

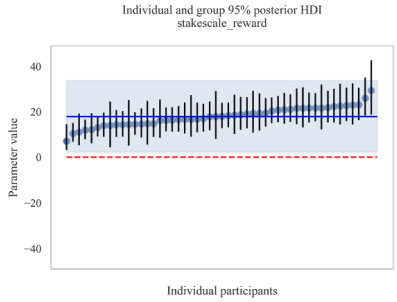
Figure S.2.1: Participant parameter posterior distributions for winning model

Winning model (Model #7) parameters. Vertical lines with dots show 95% posterior HDI and mean for each participant. Horizontal blue lines and shading show 95% posterior HDI for group-level distribution




Individual and group 95% posterior HDI
stakepereffort_reward
Parameter value
Individual participants


Individual and group 95% posterior HDI
stakepereffort_threat
Parameter value
Individual participants


Individual and group 95% posterior HDI
stakescale_reward
Parameter value
Individual participants

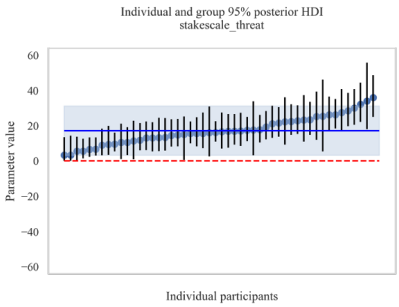
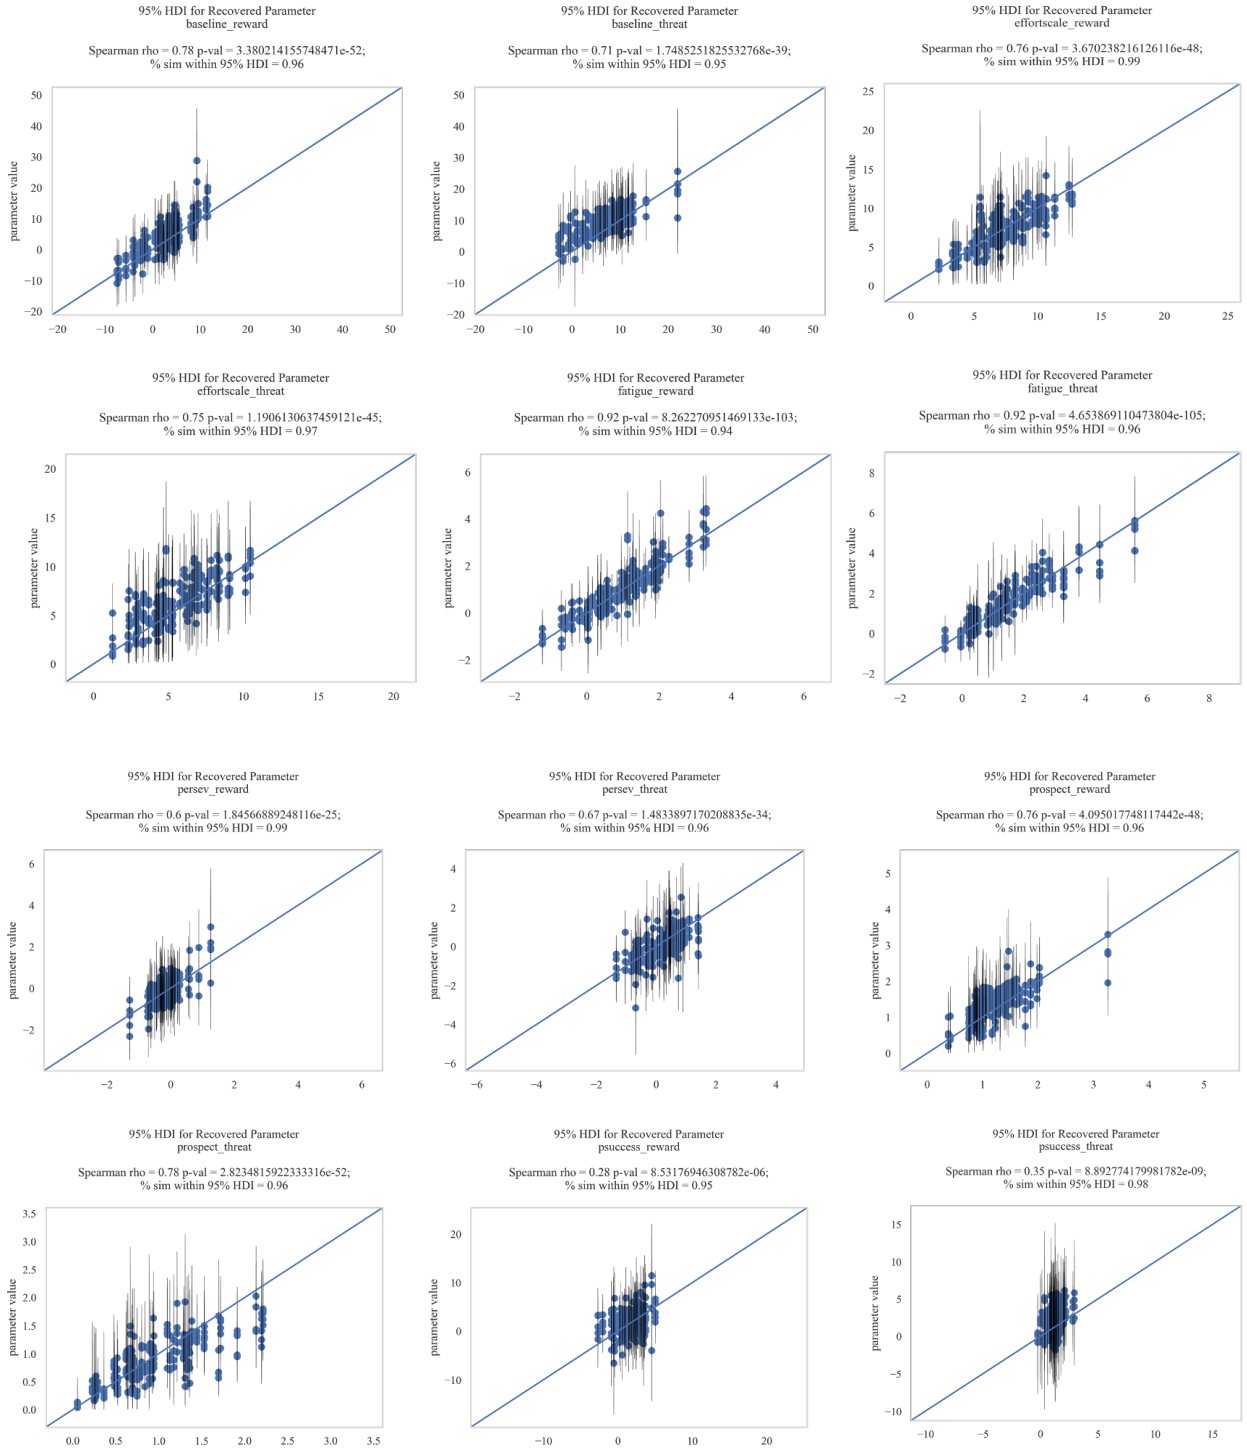
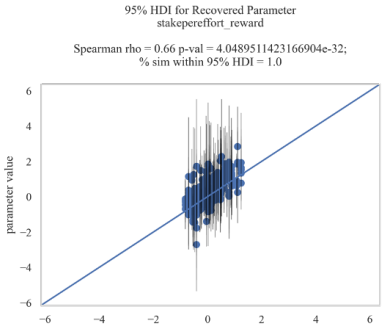
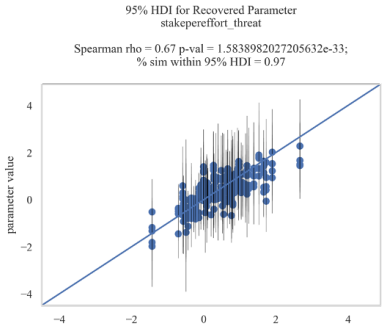
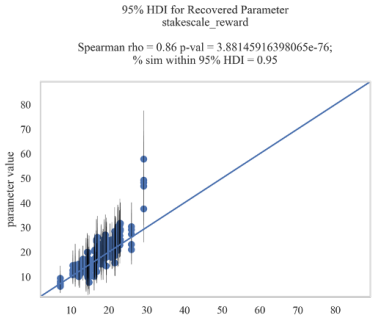

Individual and group 95% posterior HDI
stakescale_threat
Parameter value
Individual participants

Figure S.2.2: Parameter recovery for winning model




95% HDI for Recovered Parameter
stakepereffort_reward
Spearman rho = 0.66 p-val = 4.048951142316690 4e-32;
% sim within 95% HDI = 1.0
parameter value


95% HDI for Recovered Parameter
stakepereffort_threat
Spearman rho = 0.67 p-val = 1.583898202720563 2e-33;
% sim within 95% HDI = 0.97
parameter value


95% HDI for Recovered Parameter
stakescale_reward
Spearman rho = 0.86 p-val = 3.88145916398065e-76;
% sim within 95% HDI = 0.95
parameter value

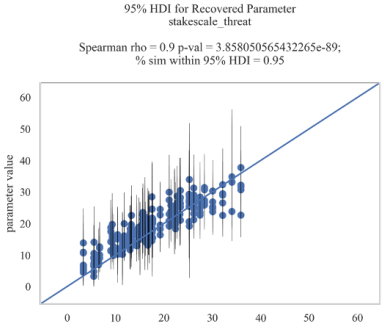

95% HDI for Recovered Parameter
stakescale_threat
Spearman rho = 0.9 p-val = 3.858050565432265e-89;
% sim within 95% HDI = 0.95
parameter value

Figure S.2.3: Model recovery

Left: Recovering Model WAIC – best recovering Model WAIC for each simulation and recovery, averaged (mean) over four separate sets of simulations. Darker blue shades on the diagonal indicate more accurate recovery of the generating models.

Right: Percentage of model winning best recovery across four iterations of simulations and recovery. Darker blue shades on the diagonal indicate closer to 100% accurate recovery.

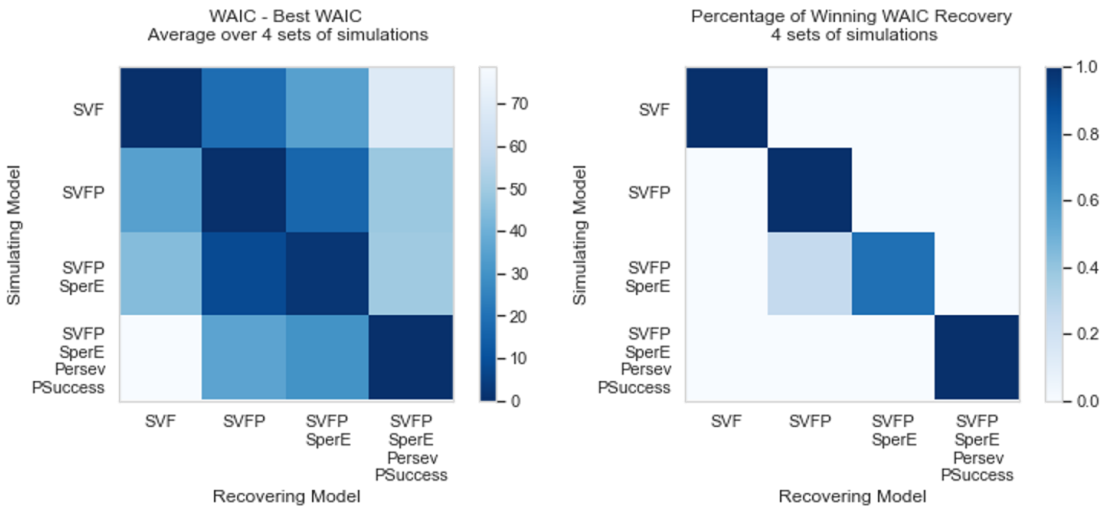


Figure S.2.4: Full list of item loadings for specific factors

Full specific factor item loadings BF5 model

Factor 1: Anxious arousal/depression

F1	Item	Text
0.59	MASQ_25	Felt numbness or tingling in my body
0.59	MASQ_61	Felt like I was choking
0.57	MASQ_19	Felt faint
0.56	MASQ_87	Had trouble swallowing
0.53	MASQ_88	Hands were cold or sweaty
0.53	MASQ_45	Had pain in my chest
0.53	MASQ_52	Felt dizzy or lightheaded
0.51	MASQ_67	Had a very dry mouth
0.50	MASQ_73	Was afraid I was going to die
0.50	MASQ_57	Hands were shaky
0.49	MASQ_48	Had hot or cold spells
0.48	MASQ_55	Was short of breath
0.47	MASQ_79	Was trembling or shaking
0.47	MASQ_75	Heart was racing or pounding
0.45	MASQ_69	Muscles twitched or trembled
0.44	BDI_21	Lost interest in sex
0.43	BDI_5	Feel guilty
0.42	MASQ_85	Had to urinate frequently
0.41	MASQ_3	Startled easily
0.41	BDI_18	Low appetite
0.40	BDI_6	Feel I am being punished
0.39	STAI_T_11	I have disturbing thoughts
0.37	BDI_10	Crying more than usual/can’t cry
0.35	BDI_14	I feel ugly/unattractive
0.34	BDI_12	Lost interest in other people
0.34	BDI_20	Worried about my health
0.32	BDI_7	Disappointed with/hate myself
0.32	MASQ_26	Felt withdrawn from other people
0.31	MASQ_66	Felt really slowed down
0.31	BDI_19	Lost weight
0.31	STAI_T_5	I feel like a failure
0.30	BDI_3	Feel like a failure
0.29	BDI_11	Irritable
0.29	MASQ_33	Felt like nothing was very enjoyable

0.29	BDI_1	Feel sad
0.29	BDI_16	Not sleeping well
0.29	STAI_T_4	I wish I could be as happy as others seem to be
0.28	BDI_17	Tired
0.26	BDI_13	Decision-making difficulty
0.26	STAI_T_20	I get in a state of tension or turmoil as I think over my recent concerns and interest
0.25	BDI_8	I blame myself
0.25	MASQ_44	Felt like there wasn't anything interesting or fun to do
0.25	MASQ_53	Felt unattractive
0.24	BDI_15	Can't get work done
0.24	BDI_2	Discouraged about future
0.23	BDI_4	As much satisfaction out of things
0.21	STAI_T_18	I take disappointments so keenly that I can't put them out of my mind
0.21	STAI_T_15	I feel inadequate
0.21	STAI_T_9	I worry too much over something that really doesn't matter
0.21	STAI_T_8	I feel that difficulties are piling up so that I cannot overcome them
0.20	STAI_T_17	Some unimportant thought runs through my mind and bothers me

Factor 2: Anhedonic depression

F2	Item	Text
0.58	MASQ_14	Felt really happy
0.52	MASQ_23	Felt like I was having a lot of fun
0.50	MASQ_49	Was proud of myself
0.49	STAI_T_10	I am happy
0.49	MASQ_40	Felt like I had a lot to look forward to
0.49	MASQ_35	Felt like I had accomplished a lot
0.42	MASQ_86	Felt really good about myself
0.42	MASQ_58	Felt really "up" or lively
0.41	STAI_T_3	I feel satisfied with myself
0.40	MASQ_1	Felt cheerful
0.39	MASQ_78	Felt hopeful about the future
0.38	STAI_T_1	I feel pleasant
0.38	STAI_T_13	I feel secure
0.37	MASQ_30	Looked forward to things with enjoyment
0.37	MASQ_18	Felt optimistic
0.37	MASQ_72	Felt like I had a lot of energy
0.34	MASQ_36	Felt like I had a lot of interesting things to do
0.34	BDI_1	Feel sad
0.33	BDI_4	As much satisfaction out of things

0.32	MASQ_26	Felt withdrawn from other people
0.32	STAI_T_16	I am content
0.29	STAI_T_12	I lack self-confidence
0.29	STAI_T_6	I feel rested
0.28	BDI_2	Discouraged about future
0.28	STAI_T_7	I am “calm, cool, and collected”
0.27	BDI_3	Feel like a failure
0.27	MASQ_27	Seemed to move quickly and easily
0.27	MASQ_33	Felt like nothing was very enjoyable
0.24	BDI_7	Disappointed with/hate myself
0.24	STAI_T_15	I feel inadequate
0.24	BDI_14	I feel ugly/unattractive
0.22	BDI_17	Tired
0.22	BDI_12	Lost interest in other people
0.21	PSWQ_9	As soon as I finish one task, I start to worry about everything else I have to do.
0.20	MASQ_53	Felt unattractive

Factor 3: Cognitive anxiety/worry

F3	Item	Text
0.61	PSWQ_5	I know I should not worry about things, but I just cannot help it.
0.53	PSWQ_12	I have been a worrier all my life.
0.52	PSWQ_16	I worry about projects until they are all done.
0.52	PSWQ_13	I notice that I have been worrying about things.
0.52	PSWQ_4	Many situations make me worry.
0.48	PSWQ_15	I worry all the time.
0.47	PSWQ_10	I never worry about anything.
0.44	PSWQ_3	I do not tend to worry about things.
0.43	PSWQ_14	Once I start worrying, I cannot stop.
0.43	PSWQ_8	I find it easy to dismiss worrisome thoughts.
0.43	PSWQ_6	When I am under pressure I worry a lot.
0.42	PSWQ_7	I am always worrying about something.
0.41	PSWQ_2	My worries overwhelm me.
0.37	STAI_T_9	I worry too much over something that really doesn’t matter
0.36	PSWQ_9	As soon as I finish one task, I start to worry about everything else I have to do.
0.35	STAI_T_2	I feel nervous and restless
0.34	PSWQ_1	If I do not have enough time to do everything, I do not worry about it.
0.34	PSWQ_11	When there is nothing more I can do about a concern, I do not worry about it any more.
0.34	STAI_T_20	I get in a state of tension or turmoil as I think over my recent concerns and interest
0.32	STAI_T_14	I make decisions easily

0.31	STAI_T_17	Some unimportant thought runs through my mind and bothers me
0.28	STAI_T_18	I take disappointments so keenly that I can't put them out of my mind
0.28	AMI_5	I make decisions firmly and without hesitation.
0.25	STAI_T_12	I lack self-confidence
0.21	STAI_T_8	I feel that difficulties are piling up so that I cannot overcome them
0.21	BDI_16	Not sleeping well

Factor 4: Behavioral Apathy

F4	Item	Text
0.61	AMI_11	I get things done when they need to be done, without requiring reminders from others.
0.53	AMI_12	When I decide to do something, I am motivated to see it through to the end.
0.53	AMI_15	When I have something I need to do, I do it straightaway so it is out of the way.
0.45	AMI_10	I don't like to laze around.
0.44	AMI_9	When I decide to do something, I am able to make an effort easily.
0.29	STAI_T_8	I feel that difficulties are piling up so that I cannot overcome them
0.28	PSWQ_2	My worries overwhelm me.
0.23	MASQ_21	Felt really bored
0.22	STAI_T_2	I feel nervous and restless
0.21	STAI_T_19	I am a steady person
0.21	BDI_15	Can't get work done
0.21	MASQ_33	Felt like nothing was very enjoyable
0.21	MASQ_39	Felt like it took extra effort to get started
0.20	MASQ_44	Felt like there wasn't anything interesting or fun to do

Factor 5 (unadjusted): Hedonia

F5	Item	Text
0.53	TEPS_18	When I hear about a new movie starring my favorite actor, I can't wait to see it.
0.51	TEPS_14	When I think of something tasty, like a chocolate chip cookie, I have to have one.
0.37	TEPS_15	Looking forward to a pleasurable experience is in itself pleasurable.
0.36	TEPS_11	When I'm on my way to an amusement park, I can hardly wait to ride the roller coasters.
0.36	TEPS_12	I get so excited the night before a major holiday I can hardly sleep.
0.34	TEPS_1	When something exciting is coming up in my life, I really look forward to it.
0.31	MASQ_79	Was trembling or shaking
0.30	TEPS_17	When ordering something off the menu, I imagine how good it will taste.
0.29	TEPS_16	I look forward to a lot of things in my life.
0.28	TEPS_3	When I think about eating my favorite food, I can almost taste how good it is.
0.22	BDI_14	I feel ugly/unattractive

Chapter 3

Ambiguous decision-making and valence-specific framing in dimensional characterization of anxiety, depression, and mania

Abstract

Individuals vary widely in their tolerance of uncertainty while making decisions. One such form of uncertainty occurs when we have incomplete information about the decisions we must make, also referred to as decision-making under ambiguity. Here, research has found that individuals tend to be biased towards risk and ambiguity avoidance in paradigms of reward gain (Cohen et al., 1985; Ellsberg, 1961), but conversely with risk and ambiguity seeking in some studies of reward loss (Machina et al., 2011). Higher levels of trait anxiety have been associated with an increased sensitivity to levels of missing information in threat avoidance decision-making under ambiguity (Lawrance et al., 2022), while depression and mania have been associated with altered valuation of rewarding outcomes in decision-making (Johnson et al., 2012; Whitton et al., 2015). In the present study we used hierarchical Bayesian modeling of decision-making under ambiguity to test the relationship of dimensions of internalizing psychopathology with ambiguous decision-making across reward pursuit, loss avoidance, and threat avoidance paradigms. We find that anhedonic depression relates to less risky loss avoidance decision-making under ambiguity, while physiological anxiety relates to more risky decisions in the same paradigm. In threat avoidance, we find a significant dissociation between levels of mania and anhedonic depression, which heightened levels of mania showing less risk sensitivity than anhedonic depression in this paradigm. Also in threat avoidance, we find suggestive (although not conclusive) evidence that sensitivity to magnitude of threat may relate to shared variance across depression and anxiety.

Introduction

Every day we must make decisions without complete information about the outcome of our actions. When the outcome of a decision is probabilistic rather than deterministic, the associated uncertainty is referred to as "risk." Individual differences in decision-making under risk have been widely studied, particularly in the field of behavioral economics. Classic research findings have shown that it is typical for individuals to be avoidant of risk when outcomes involve reward (often operationalized as monetary gains in research), but to be risk-seeking in decisions when outcomes involve avoiding loss (Fox et al., 2015; Kahneman and Tversky, 1979). This has often been referred to as the "gain-loss framing" attribute of individual risk preference, and hints at a broader pattern of both within- and across-individual variability in valence-specific decision-making

preferences which may also extend to decisions involving the avoidance of threat (Lawrance et al., 2022).

Human and animal research has implicated the prefrontal cortex and its dopaminergic afferents in propensity for risk-taking (Orsini et al., 2015). Individuals with damage to ventromedial and orbital frontal cortex (OFC) show an increase in risk-taking behavior relative to controls as studied by the well-known Iowa Gambling Task, in which participants select between card decks which yield rewards and penalties at various (hidden) schedules; over time, more risk-seeking behavior manifests in a greater number of draws from the high reward/high penalty card decks, even though the ultimate net gain is lower from those decks (Bechara et al., 1999, 1994). In addition, conditions which are behaviorally linked to increased risk-taking behavior such as substance abuse and attention deficit hyperactivity disorder (ADHD) show decreased activity of prefrontal cortical regions associated with increased risky decision-making relative to controls (Crowley et al., 2010; Ernst et al., 2003; Fishbein et al., 2005; Klein et al., 2012).

Various forms of internalizing psychopathology have also been associated with alterations in risk tolerance. In a study using a gambling based decision-making task in which participants chose between risky and safe outcomes of varying magnitudes of financial reward gain and loss, participants clinically diagnosed with generalized anxiety disorder or panic disorder were found to make significantly fewer risky choices than paired controls (Giorgetta et al., 2012). Another study using a similar gambling task involving only financial reward gain found that induced anxious mood in healthy participants was also linked to a bias against risky choices (Raghunathan and Pham, 1999). Interestingly, the latter study also examined the effect of induced sad mood on the same task and found that sadness was related to a bias towards, rather than away from, risky choices, suggesting a differential influence of subtypes of negative affect. Meanwhile, research using the balloon analog risk task (BART), a well-known paradigm to test risk aversion in research, confirmed that individuals with higher levels of subclinical trait anxiety showed increased risk aversion in reward/loss decision-making, while scores on the negative affect subscale of the Positive and Negative Affect Scale (PANAS; Watson et al., 1988), which has been posited to measure symptomatology generally corresponding to shared variance between anxiety and depression, failed to show the same association (Maner et al., 2007).

These and other studies have clearly established a relationship between anxiety and risk avoidance, but such studies often have not fully separated reward gain and loss avoidance decisions, instead allowing decisions to involve both gain and loss, obscuring any framing effect and its role in the relationship between anxiety and risk aversion. One study designed to correct for this used a variant of a gambling task in which some trials involved both gain and loss, and other trials involved gain only (but no trials involved loss only), so that risk aversion and loss aversion could be parametrically distinguished (Charpentier et al., 2017). This study found that individuals with high anxiety scores showed increased risk aversion, but not increased loss aversion, suggesting specificity of effect across reward and loss and emphasizing the importance of differentiating between valence paradigms in decision-making. Meanwhile, symptoms of other forms of psychopathology, such as mania, have been associated with increased tolerance of risk and risk-seeking behavior (Devlin et al., 2015; Edge et al., 2013; Lukacs et al., 2021; Mason et al., 2014; Wong et al., 2021), again hinting at a differentiable and perhaps opposing influence of varying dimensions of psychopathology on risk sensitivity.

In addition to the risk inherent in a probabilistic outcome, which is considered "first-order" uncertainty, decision-makers frequently encounter forms of "second-order" uncertainty in the world. Second-order uncertainty occurs when there is some uncertainty regarding the first-order uncertainty (e.g., probability) itself. This may occur, for example, when the underlying unknown probabilities of outcomes are changing (or volatile) over time. Individuals generally consider previous outcomes in making decisions; periods of heightened outcome volatility require individuals to more heavily weight recent rather than older past experience given the rapidly changing outcome environment. Research has shown that while people generally do this with surprising efficacy (Behrens et al., 2007), higher levels of anxiety are associated with a reduced ability to adjust learning rate between volatile and stable environments in threat avoidance decision-making (Browning et al., 2015). Recent advances on this work have found that here the impairment may be related to the shared variance between depression and anxiety, rather than anxiety alone (Gagne et al., 2020), highlighting the importance of a dimensional approach to evaluation of psychopathological traits or symptoms.

Second-order uncertainty also occurs in decisions involving information ambiguity, when some of the information required to accurately estimate a probabilistic outcome is missing or unavailable. In ambiguous decision-making, individuals must estimate outcome probabilities based on the information that is available. Like risk preference, ambiguity sensitivity may vary across individuals. Individuals have been shown to be broadly ambiguity averse when faced with the presence versus absence of ambiguity (Cohen et al., 1985; Ellsberg, 1961); however, as with decision-making under risk, findings indicate that this may be specific to ambiguous decisions involving potential reward gain, and that ambiguous decisions involving potential reward loss may in fact be associated with more ambiguity-seeking behavior (Machina et al., 2011). Research into decision-making under ambiguity using aversive outcomes (electrical stimulation) has linked higher levels of trait anxiety to elevated sensitivity to the precise level of missing information in the absence of any significant relationship with the categorical presence or absence of ambiguity (Lawrance et al., 2022), suggesting that categorical sensitivity to ambiguity and sensitivity to changes in level of missing information may be experimentally separated and may yield different findings with respect to psychopathology.

In addition to risk and ambiguity sensitivity, various symptoms of psychopathology are also differentially related to valuation of reward, loss, and threat. For example, depression (Treadway et al., 2012) and its associated symptoms of anhedonia (Treadway and Zald, 2011) and apathy (Bonnelle et al., 2015a, 2015b) have been associated with a reduced anticipation of reward and a reduced willingness to exert effort to pursue rewarding outcomes, while anxiety has long been linked to an increased valuation of threat (Chen and Lovibond, 2020; Huys et al., 2015; Levy and Schiller, 2021). These differences motivate the application of computational models to empirical data collected using parallel versions of decision-making tasks that assess risk aversion and ambiguity sensitivity during pursuit of reward versus avoidance of loss or threat to tease apart differences in outcome valuation, risk sensitivity, and ambiguity sensitivity in psychopathology.

In the current study, we aimed to investigate the effect of dimensional measures of depression, anxiety, and mania on individual differences in ambiguity sensitivity across three parallel versions of a decision-making under ambiguity task which differed in outcome type (potential reward gain,

potential reward loss and potential shock receipt). We used hierarchical Bayesian modeling of decision-making under ambiguity to assess individual differences in mechanisms of the decision process including risk preference, categorical ambiguity sensitivity, and sensitivity to level of missing information. In particular, our aims for the present analysis were as follows: 1) We sought to test whether a dimensional characterization of subclinical mania symptomatology was associated with reduced sensitivity to categorical ambiguity, reduced sensitivity to increasing level of missing information, or increased risk preference during decisions involving gain (and secondarily during loss and threat avoidance decisions). 2a) We also sought to build upon previous findings that elevated trait anxiety is linked to increased sensitivity to level of missing information by exploring the dimensional specificity of this relationship. In particular, we tested whether physiological anxiety, cognitive anxiety, or a general dimension of negative affect (representing variance common to both anxious and depressive symptoms; Gagne et al., 2020) were related to ambiguity sensitivity in threat avoidance decisions; 2b) We further sought to test whether effects of anxiety and depression on ambiguous decision-making are outcome specific varying across reward-gain, reward-loss and threat (shock) versions of the decision-making under ambiguity task.

Methods

Behavioral participants and experimental paradigm

Participants for the in-lab behavioral study (final n=71; see Results) were recruited through the UC Berkeley undergraduate research participation pool (RPP), the UC Berkeley experimental social sciences lab (XLAB), fliers or advertisements posted online or within the local community (up to a 10-mile radius), and via e-mail advertisements sent to community e-mail lists. All study procedures were approved by the UC Berkeley Committee for the Protection of Human Subjects (CPHS).

Participants interested in the study were sent an e-mail outlining the study eligibility criteria, the requirement that ID and proof of local residence be shown during the first online session for verification purposes, and the estimated time commitment possible during the study. Participants were then able to schedule a paid initial study screening via online administration. During the online screening session, after informed consent, participants who were not recruited directly through the University were required to show ID and proof of local residence to the experimenter prior to proceeding in order to verify identity. This ID requirement was implemented approximately halfway through the recruiting process due to a high number of identified (and excluded) fraudulent repeat online survey completions. These fraudulent surveys did not affect the final study data since no fraudulent participants completed the in-lab portions of the study.

Following informed consent and proof of ID, the participant was then directed to an online screening session, which included questions about demographics, medical questions, and a number of self-report questionnaires to measure various traits spanning the internalizing psychopathology spectrum. All participants who completed the initial online screening were compensated at the rate of $15 per hour paid electronically via Hyperwallet (hyperwallet.com). Participants recruited via RPP could alternatively elect to be compensated with partial course credit according to the UC Berkeley RPP policies for online psychology study participation.

After completion of the online screening, participant responses were reviewed for eligibility by the experiment team. In particular, participants were excluded following this stage if they reported more than occasional hard drug use, frequent recent soft drug use, diagnosis with neurological disorders or psychiatric disorders characterized by psychotic symptoms, or failure to provide proof of ID and local residence. Participants who were eligible to proceed were contacted via phone or e-mail to schedule a secure Zoom-based structured clinical interview (SCID) with a clinically trained team member. Participants who completed the Zoom-based SCID were compensated $15 per hour for their time via Hyperwallet (or again, via RPP course credit if eligible and if this option was elected).

Following the SCID, eligible participants were next invited to complete a 2-session in-lab experiment in which they would answer questionnaires and complete a decision-making under ambiguity task in reward pursuit, loss avoidance, and threat avoidance. Participants completed informed consent prior to each in-lab session. Following informed consent in the first session, participants completed a health screening questionnaire and were also asked to complete a urine-based drug test to confirm no recent cannabis use. Participants deemed ineligible to participate at this stage for any reason were compensated for 1 hour of their time and informed they did not meet study eligibility criteria. Drug test results or illicit drug use disclosure were not recorded or stored for participants.

Participants then completed a computer-based battery of questionnaire measures pertaining to the trait space of interest (see detail in the Item Response section below). This in-person set of questionnaires included two depression measures with questions pertaining to suicidality. Although the experimenters were not able to see participant item responses in real-time, high-risk participants were indicated to the experimenter via several possible alpha-numeric "completion codes" which were displayed at the end of the questionnaire session. These completion codes were not interpretable by participants as related to responses but were known to experimenters. For completion codes indicating high levels of depression or suicidal intent, we followed a CPHS-approved safety plan which was implemented at the end of the first behavioral session. At the conclusion of the first session, a specially trained clinical experimenter joined the session (if not already the primary experimenter for the session) and discussed the depression or suicidality flag that had been raised; provided a written list of resources for help with depression and other mental health burdens; and offered to escort the participant directly to student counseling or health services. Additionally, a senior clinical team member was available for an immediate phone call counseling session with the participant.

Participants completed both the threat avoidance task and either the reward or loss task (counterbalanced across individuals) in their first in-lab session (with order of within-first-session tasks also being counterbalanced).

Following informed consent prior to the second in-lab session, participants completed another series of questionnaires, followed by their remaining ambiguous decision-making task valence (reward or loss). Participants also completed a sustained attention to response task (SART) and N-back task in the second session, which were not used in the present study analysis.

Half of participants were invited back to repeat this experimental paradigm 4-8 weeks later for test-retest validation of the experimental protocol. Data from these repeat sessions are not included or analyzed in the present write-up.

Participants were compensated for in-lab sessions at the rate of $15 per hour in cash (or via partial course credit according to the UC Berkeley RPP guidelines for RPP participants electing this option). In addition, participants could earn a monetary bonus of up to $10 during the reward pursuit task. Participants who completed both in person sessions were paid a completion bonus of $15 to incentivize study completion.

Data processing and exclusion: in-lab participants

A total of 120 participants completed part or all of the decision-making under ambiguity task. We included only participants who successfully completed all three task valences (reward pursuit, loss avoidance, or threat avoidance) in the present analysis. 8 participants were excluded due to incomplete task data. There were 13 unambiguous trials in the experiment in which both probability of outcome and magnitude of outcome favored the same urn (referred to as "no-brainer" trials because all metrics were in favor of the same urn). Participants who failed to select the more favorable urn on at least 85% of these trials in each valence paradigm were excluded from the data. This resulted in the exclusion of 41 additional participants. There were n=71 in-lab behavioral participants used in the final analysis; mean age = 21.67 years, standard deviation = 4.08 years. Of these, 60 = female; 14 identified as Hispanic or Latino and 57 identified as not Hispanic or Latino; 42 identified as Asian, 20 as White, 6 as unknown race, and 3 as more than one race.

Online participants (questionnaire item responses)

To adequately power exploratory bifactor analyses across self-report questionnaires, we used questionnaire response data from an online set of participants taking part in another behavioral study for our lab. The behavioral data from this experiment is not considered in the present study, but the questionnaire responses for the final online participant group (after all exclusions) were used for exploratory bifactor modeling as detailed below.

An online decision-making task was administered to a large data set (n=1022) recruited through Amazon Mechanical Turk (MTurk). All procedures for the online study were approved by the UC Berkeley CPHS. To ensure high-quality responses, only MTurk workers who had successfully completed over 500 MTurk Human Intelligence Tasks (HITs) with at least a 98% rate of task approval were allowed to enroll for the online study HIT. We additionally specified that our HIT should only be shown to individuals who had identified themselves to Amazon as physically located in the United States and as being between the ages of 18 and 55.

Participants first completed informed consent online and a short screening in which they provided demographic and medical information to determine study eligibility. Participants were not eligible to continue if they endorsed cannabis use within the 2 preceding weeks, or if they were currently taking antidepressant or anxiolytic medication. Eligible participants were next directed to an online series of questionnaires relating to internalizing psychopathology. After completing the first set of questionnaires, participants were directed to an online behavioral task (not analyzed here).

Following the online behavioral task, participants completed another short round of questionnaire responses. They were then directed to a final screen in which they were informed that they had completed the task and their payment was guaranteed, and they were asked to respond honestly to three final questions to aid in the accuracy of our research (following Meade & Craig, 2012). Participants were asked to confirm 1) whether they were physically located in the United States; 2) whether they had used cannabis in the previous two weeks; and 3) whether they thought we should include their data in our study based on their own judgment of their attentiveness and honesty during the task and questionnaires. Following these final questions, participants were paid $20 via the MTurk platform for their completion of the study (which was not allowed to exceed two hours).

Data processing and exclusion: online participants

To ensure the best data quality possible in the questionnaire responses for the online data set, we performed exclusions on this data based on questionnaire attention checks, responses to final questions, and poor performance on the behavioral task that the participants completed.

Participants were excluded from the online data if they missed two or more of the five questionnaire attention checks; if they disclosed that they did not live in the US or that they had used cannabis within the last 2 weeks; if they responded that we should not use their data in our research; or if they did not successfully answer at least one iteration of comprehension check questions during the behavioral task. Participants were also excluded based on high number of timeouts or low performance on "no-brainer" trials during the behavioral task (not considered here). We used the full set of exclusions (both questionnaire and behavioral) for the online questionnaire response data to ensure only high-quality responses were included in factor analysis. The final online participant group comprised n=470 individuals. The questionnaire responses for these individuals were used for exploratory bifactor modeling as detailed below.

In-lab participants (questionnaire item responses)

To adequately power confirmatory bifactor analyses across self-report questionnaires (i.e., a greater number of participants than items, with number of items = 126), we performed CFA in a larger set of in-lab participants across the current behavioral task and another (fMRI) study performed by our lab in which participants also completed the necessary questionnaire responses. Data from that study are not considered in the present analysis except for relevant questionnaire responses. Participants in the fMRI study shared the entire recruitment pipeline described above for the in-lab behavioral study except for the inclusion of additional medical questions to verify MRI eligibility. Participants in the in-lab ambiguous decision-making task were not required to complete all behavioral valence task sessions to have questionnaire response data included for CFA. For uniformity, scores for all questionnaire measures used in the analysis were recorded during the reward task behavioral session, except mania scores which were recorded only at screening. The final participant group for confirmatory factor analysis comprised n=179 individuals.

Item-level responses to measures of anxiety, depression, and mania

Participants for the in-lab and online experimental paradigms completed a number of self-report questionnaires to assess multiple dimensions of psychopathology. The current analysis sought to assess the relationship between subclinical internalizing psychopathology (particularly anxiety and depression) and hypomania with ambiguous decision-making. Measures used in the present analysis evaluated trait levels of symptomatology (i.e., stable over the course of the experiment) rather than state levels (i.e., specific to the moment of inquiry). In-lab participants completed some questionnaire measures multiple times over the course of the screening and multi-session behavioral study. For consistency, responses from the reward session were used for any duplicated questionnaires. The Chapman Hypomanic Personality measure was only completed during online screening, so the screening responses for that measure were used in the present analysis.

The questionnaire measures used comprised the following: State Trait Anxiety Inventory Trait scale (STAI-T; Spielberger, C. D., Gorsuch, R. L., Lushene, R., Vagg, P. R., & Jacobs, G. A., 1983); the Mood and Anxiety Symptoms Questionnaire (MASQ) anxious arousal (AA) and anhedonic depression (AD) subscales (Watson and Clark, 1991); the Penn State Worry Questionnaire (PSWQ) (Meyer et al., 1990); Beck Depression Inventory (BDI-II) (Beck, 1961; Steer et al., 1998); Centers for Epidemiological Studies Depression Scale (CES-D) (Radloff, 1977); the Inventory of Depression and Anxiety Symptoms II (IDAS-II) mania and euphoria subscales (Watson et al., 2012); and the Hypomanic Personality Scale (HPS) (Eckblad and Chapman, 1986). Items addressing suicidality were excluded. Final number of items n=126. See Results section for additional details.

The set of online participants (n = 470) was used for exploratory bifactor analysis. Key main score means and standard deviations for this group were as follows: MASQ AA (23.99 +/- 8.48) (total possible score range on MASQ AA = [17,85]); MASQ AD (58.04 +/- 19.0) (total possible score range on MASQ AD = [22,110]); PSWQ (47.73 +/- 17.83) (total possible score range on PSWQ = [16,80]); Chapman hypomania scale (10.43 +/- 8.19) (total possible score range on Chapman hypomania scale = [0,48]); IDAS mania/euphoria scale items (15.83 +/- 5.91) (total possible score range on IDAS mania/euphoria items = [10,50]). Full score distribution plots for each subscale and correlations across subscales are provided in Supplements 3.S.1a and 3.S.1b.

For confirmatory analysis, the questionnaire responses for a combined set of in-lab participants across 2 studies were used (n=179) in order to adequately power confirmatory bifactor analysis across the number of response items (items=126). For the confirmatory analysis group, key main score means and standard deviations were as follows: MASQ AA (21.89 +/- 5.98); MASQ AD (56.57 +/- 16.16); PSWQ (54.21 +/- 14.61); Chapman hypomania scale (15.32 +/- 8.16); IDAS mania/euphoria scale items (15.01 +/- 5.2). Full score distribution plots for each subscale and correlations across subscales are provided in Supplements 3.S.2a and 3.S.2b.

Following determination of the best fit bifactor model (BI-4), participant factor scores were calculated for the in-lab participants of the decision-making under ambiguity task (n=71). These factor scores were used for hypothesis testing along dimensions of internalizing psychopathology and manic symptomatology. Within this group, key main score means and standard deviations were as follows: MASQ AA (21.85 +/ 4.46); MASQ AD (54.89 +/- 15.82); PSWQ (53.77 +/- 13.5);

Chapman hypomania scale (13.87 +/- 7.21); IDAS mania/euphoria scale items (14.56 +/- 5.54). Full score distribution plots for each subscale and correlations across subscales are provided in Supplements 3.S.3a and 3.S.3b.

Scores across populations were broadly similar, with the exception of somewhat lower mean hypomania scores among online participants. Since exploratory data analysis was performed in the online group, a high confirmatory fit in the in-lab samples was key to validate the structure of the bifactor model (see Bifactor analysis of latent dimensions of psychology section below).

Bifactor analysis of latent dimensions of psychopathology

Bifactor analysis was used to parse common from specific variance across the anxiety, depression, and hypomania trait space. Prior work exploring the latent structure of factors in anxiety and depression has found that the high comorbidity between the two disorders is well-described by a tripartite model in which shared variance is orthogonally removed from specific anxiety and depression elements (Clark et al., 1994; Gagne et al., 2020).

More recent work has found that a bifactor analysis with more than 2 specific factors may be used to parse symptom-level heterogeneity within each condition (see Chapter 2 of this dissertation). It is relatively rare for those with anxiety and depressive symptoms to report manic symptoms; however, two-thirds of individuals diagnosed with bipolar disorder experience anxiety disorders during their lifetime, and a similar number experience depressive episodes during their lifetime (Kotov et al., 2017). Given the hypotheses of the current study, we included measures of hypomania in our bifactor analysis in order to confirm absence of any large dimensional shared variance with internalizing psychopathology in general, and to control for any potential shared variance between symptoms of hypomania and those of physiological anxiety.

To facilitate optimal modeling power and ensure model validation in an independent sample to mitigate any overfitting, we performed exploratory factor analysis using a large online sample of participants who completed the requisite set of item response questionnaires in connection with a separate study run by our lab (n=470). We performed confirmatory bifactor analysis in a set of in-lab participants who had completed one of two decision-making tasks in our lab (ambiguity and volatility tasks) and who had completed the requisite item response questionnaires (n=179).

We tested bifactor model structures in the EFA data set (n=470) specifying 3, 4, 5, and 6 specific factors (with one general factor in each model describing common variance). Exploratory and confirmatory bifactor analyses were performed in R using the ‘psych’ package. Polychoric correlations were used for exploratory analyses given the ordinal nature of item responses (Jöreskog, 1994). Schmid-Leiman (Schmid and Leiman, 1957) bifactor rotations were used to extract one common factor and the specified number of specific factors from the item response covariances with all factors orthogonal to one another in the multidimensional space. Participant factor scores were calculated using the Anderson-Rubin method (Anderson TW, Rubin H., 1956) which preserves orthogonality in the resulting individual scores.

We compared model fit across the four proposed bifactor models using the Bayesian Information Criterion (BIC), a measure of model fit which penalizes for parametric complexity. We found that

BIC improved with each additional specific factor included in the analysis, with the bifactor model specifying 6 specific factors showing the lowest BIC= -26,753.13 in the exploratory data (where lower BIC corresponds to better fit; see Figure 3.1).

We performed confirmatory analysis using each bifactor structure fit above in the CFA data set (n=179) and compared both the confirmatory fit index (CFI) and the Tucker-Lewis measure of fit (TL) across structures. The bifactor model structure with 4 specific factors (BI-4) provided the best fit to the out-of-sample confirmatory data (CFI = 0.989, TL = 0.989, RMSE = 0.036), indicating that the models with 5 and 6 specific factors may have shown better fit in the EFA data set due to overfitting. The bifactor model with 4 specific factors delineated latent dimensions within the data which we characterized based on predominant item loadings. Figure 3.1(c) illustrates the correlations between the exploratory data participant factor scores on the BI-4 model and the questionnaire subscale scores. Consistent with expectations, the general factor was highly correlated with the internalizing scale measures, with the lowest correlation across anxiety/depression metrics seen in the MASQ AA subscale (Pearson r = 0.45, p = 6.15e-25). The two measures of mania included (HPS and IDAS Mania/Euphoria) did not show strong levels of common variance with the general internalizing factor (Pearson r = 0.08, p = 0.09, Pearson r = 0.04, p = 0.36 respectively).

After accounting for shared internalizing variance, one specific factor was highly correlated with measures of depression, particularly with the MASQ anhedonic depression subscale (Pearson r = 0.68, p = 4.12e-64). A second specific factor was most strongly correlated with the MASQ anxious arousal subscale (Pearson r = 0.85, p = 1.84e-129) and the IDAS mania/euphoria items (Pearson r = 0.56, p = 1.91e-39), in addition to moderate but lesser correlations with measures of anxiety and depression (Pearson r < 0.35 all other measures). A third specific factor was very highly correlated with the two measures of mania (HPS Pearson r = 0.93, p = 5.83e-203; IDAS Mania/Euphoria Pearson r = 0.62, p = 2.19e-50) and showed low or no correlations with all other measures. The fourth specific factor correlated highly with PSWQ (Pearson r = 0.70, p = 1.18e-70) and EPQ neuroticism (Pearson r = 0.54, p = 2.09e-36), with additional more moderate correlations to STAI-Trait (Pearson r = 0.35, p = 6.59e-9). Based upon these correlations and the associated item-specific loadings, we characterized the factors as follows: Factor 1 as (anhedonic) depression; Factor 2 as physiological anxiety/arousal; Factor 3 as (hypo)mania; and Factor 4 as cognitive anxiety/worry.

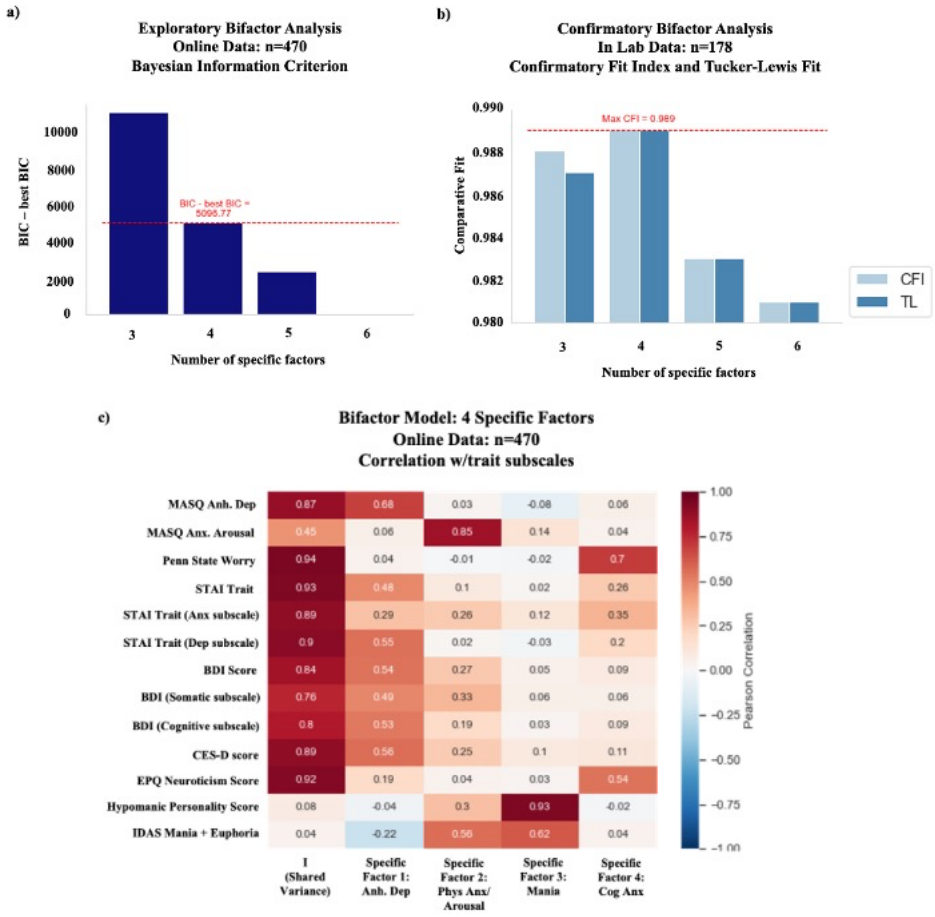


Figure 3.1 Bifactor analysis results. a) Model fit in exploratory data as measured by Bayesian Information Criterion (BIC) for proposed bifactor models with 3, 4, 5, and 6 specific factors. b) Confirmatory fit of bifactor model structures in confirmatory analysis data set, measured by confirmatory fit index (CFI) and Tucker-Lewis measure of fit (TL). c) Correlation of individual factor scores from winning bifactor model (BI-4) with subscale scores on questionnaire measures.

Decision-making under ambiguity behavioral task

The decision-making under ambiguity task used was based upon a modified version of the classic Ellsberg (Ellsberg, 1961) urns task which had previously been adapted by our lab for ambiguous decision-making investigations (Lawrance et al., 2022). Here we extended this experimental paradigm to include reward pursuit, loss avoidance, and threat avoidance.

The task consisted of 156 trials in each valence paradigm. On each trial, participants were presented with two urns, each filled with a mixture of 'X', 'O', and '?' tokens. Each urn contained 50 total tokens. In each paradigm, 'X' corresponded to receiving the outcome of the trial (i.e., a point gain in the reward pursuit task, a point loss in the loss avoidance task, or an electrical stimulation in the threat avoidance task) such that (for unambiguous trials) the percentage of 'X' tokens in the urn gave the probability of receiving the outcome for that urn if selected. Each urn also had an associated magnitude displayed above it which corresponded to the level of outcome associated with that urn. Participants therefore were required to evaluate both the relative chance

of the urn selection leading to an outcome (as indicated probabilistically by the ratio of number of 'X's to total number of both 'X's and 'O's in the urn) and the associated outcome magnitude for that urn. Participants selected one of the two urns shown on each trial, and were then shown the outcome for the chosen urn. Urn outcomes were tracked on the screen for each block of 6 trials, after which one of the preceding 6 outcomes was randomly selected for the participant to receive. The selected outcome for the block was then immediately delivered (point gain, point loss, or electrical stimulation). This manipulation was used to minimize the actual total amount of electrical stimulation received by participants during the threat avoidance version of the task, while still incentivizing them to make a meaningful decision on every trial.

One-half (78) of the trials in each paradigm involved an ambiguity manipulation in which varying numbers of characters in one urn (the "ambiguous" urn) were masked with a '?' character. Participants were never shown the true masked characters, and were informed that the order of the revealed characters could not be used to accurately deduce the masked characters of the urn. Therefore the probability of the outcome for the "ambiguous" urn had to be inferred by the participant based on the tokens revealed. There were 6 different levels of ambiguity throughout the task, in which either 10, 30, 40, 45, 47, or 49 out of 50 total tokens were masked, allowing for the orthogonal manipulation of level of missing information with magnitude and probability of outcomes.

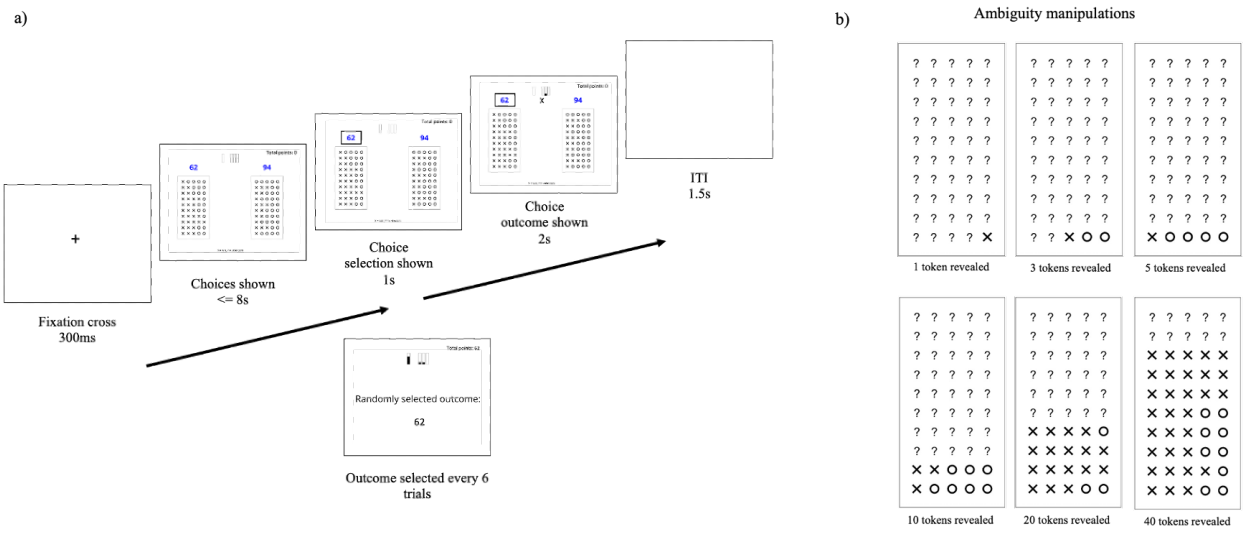


Figure 3.2 Decision-making under ambiguity task. a) On each trial, participants are presented with two urns, and have up to 8 seconds to make a selection between them. A token is then drawn from the selected urn, with an 'X' representing the magnitude outcome shown for that urn (winning points in the gain scenario, or losing points in the loss scenario), and an 'O' representing a 0 outcome. Outcomes are binned by value (0; 1-50; 50-100; 100-150) and tracked visually for each block of 6 trials. After each block, a single outcome is chosen at random and the participant receives the selected outcome. b) Ambiguous urns have 6 possible levels of missing information, with 1, 3, 5, 10, 20, or 40 tokens revealed.

Behavioral analysis

Participant choice behavior was analyzed using Python programming language. Statistics were performed using scipy and numpy packages for Python except where otherwise noted. Pearson correlations were used to evaluate the relationships between subscale scores and bifactor scores across individuals. Other statistical tests were nonparametric when underlying data was non-normal, and parametric otherwise, as specified throughout the analysis. P-values throughout are not corrected for multiple comparison family-wise error unless otherwise noted.

Computational modeling of value-based decision-making under ambiguity

We used a series of computational models of value-based decision-making under ambiguity to investigate cognitive decision-making mechanisms and associated individual differences. Since decision-making behavior for a given individual may be fundamentally different between paradigms of reward pursuit, loss avoidance, and threat avoidance, we modeled each valence paradigm separately and performed model comparison within each valence. We then examined comparative model fit using the Widely Applicable Information Criterion (WAIC; Watanabe S, Opper M, 2010) to determine the best fitting model for each paradigm.

On each task trial, participants were presented with both a magnitude for each urn choice and information which could be used to estimate the probability of outcome for each urn. The information for probability estimation was partially masked at various levels on half of trials, manipulating the level of ambiguity under which probability must be estimated (see Task description above).

The actual decision-making computation admits of various possibilities in both subjective estimation of relevant quantities and in relative value comparison. In the present analysis we considered a model space including representations of each foundational algorithmic possibility as explicated below.

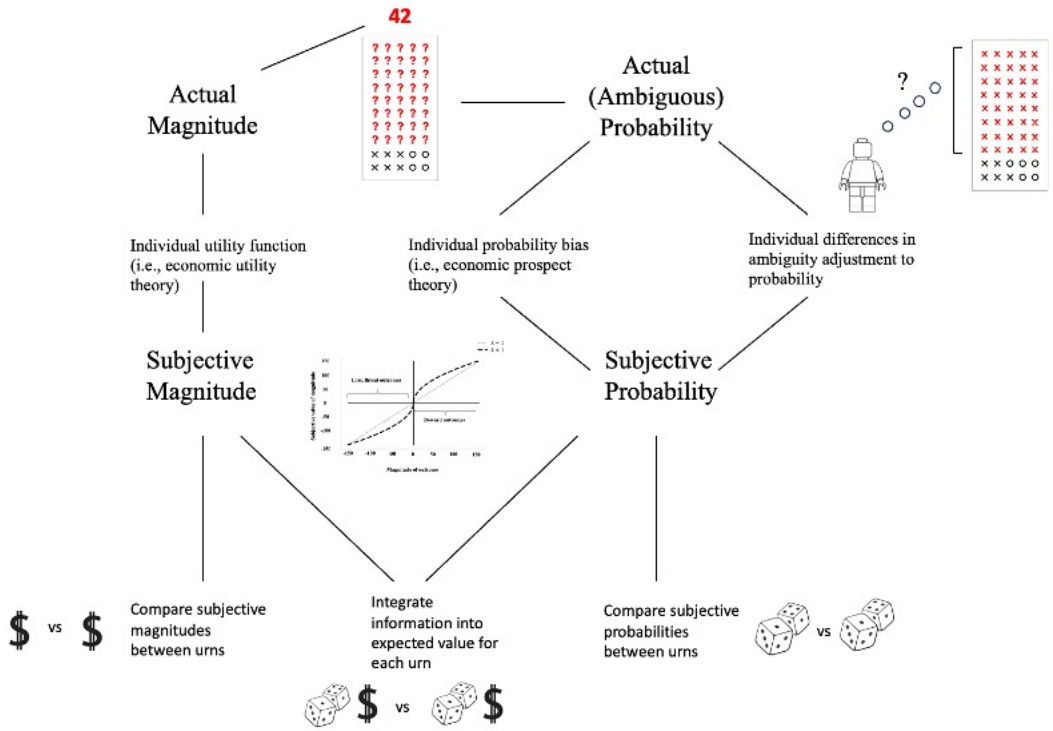


Figure 3.3 Schematic of decision-making under ambiguity. Individuals may vary in subjective valuation of magnitude, subjective valuation of probability/risk, method of updating probability in the presence of ambiguous information, and method for comparison of value between two urns.

Foundational components of risky decision-making under ambiguity

First, participants may differ in their subjective value of outcome magnitudes within and across each valence. For example, in the reward pursuit task, some participants may consider small reward magnitudes to be more valuable than other participants do. The same participant may hold different subjective value relationships for monetary losses and for threatening outcomes. These natural variations in subjective outcome value may be characterized by a mechanism as simple as a relative linear weighting of magnitudes, or by a full subjective utility curve as in expected utility theory of behavioral economics (Kahneman and Tversky, 1979). We consider models which account for all combinations of these possible cognitive approaches.

Second, participants may differ in their subjective sensitivity to the probability of an outcome (sometimes referred to as the "risk" of an outcome) when all information is available for the outcome (i.e., when ambiguity is not present). As with magnitude differences outlined above, a given individual may vary in subjective valuation of risk across paradigms of reward, loss, and threat. These differences in subjective probability distortion may be linear or nonlinear (or both), and our models again take these possibilities into account.

Third, participants may differ in the ways in which they infer probability in the absence of information (i.e., when ambiguity is present). Here a variety of mechanisms are possible. For example, participants may perform an unbiased Bayesian update of the outcome probability based on the level of information shown. Alternatively, they may have optimistic or pessimistic biases

regarding hidden information, or they may use distorted priors or expectations to update probability in the presence of missing information. Our models consider each of these possibilities for the inference of probability under ambiguity.

Finally, participants may perform value comparison in different ways. Previous research in decision-making under ambiguity using the same urn-selection task in a threat avoidance paradigm found that participants compared urn magnitudes and urn probabilities separately during value comparison (Lawrance et al., 2022). However, extensive research in behavioral economics in the (unambiguous) monetary reward domain has indicated that here individuals generally take an expected utility approach in which subjective magnitude and probability information for each choice is integrated prior to comparison (Kahneman and Tversky, 1979; Peterson et al., 2021). Again, our models allow for either or both of these mechanisms in the value comparison process.

For organizational purposes, we structure our model space into categories based on mechanism of value comparison. We consider models which weigh up magnitude differences and probability differences separately (magnitude vs. probability or MP models); models in which information for each urn is integrated prior to value comparison (expected utility or EU models), allowing for nonlinearity on magnitude (EUM) or magnitude and probability (EUMP); and models in which both separate and interactive mechanisms of magnitude and probability in value comparison are performed (fully interactive, or FI models).

Within each model category, we divide models into 2 classes to assess various potential ways of updating estimates of probability under ambiguity prior to value comparison. Here, we consider unbiased Bayesian estimates of ambiguous probability versus optimistically or pessimistically adjusted biases. For ambiguous trials, the probability of the ambiguous urn P_A is set equal to this posterior expectation (in accordance with either model class). For unambiguous trials, $P_A = P_{Urn1}$, where unambiguous urns are randomly pre-assigned to be either Urn 1 or Urn 2.

Within each model class we further explore the influence of changes in bias over task and level of missing information in ambiguous trials. Models were fit separately for reward pursuit, loss avoidance, and threat avoidance valence paradigms to allow for potentially different mechanisms of decision-making across valences.

In the positive valence paradigm (reward pursuit) task, larger outcomes indicate larger potential gains and are therefore favorable. In the negative valence paradigms (loss and threat), larger outcomes indicate larger losses or threat stimulation, and are therefore unfavorable. To provide consistency in the interpretation of parameters across paradigms, we have adjusted the equations where appropriate based on valence, with (for all models):

$$\text{reward gain paradigm:} \quad valence = 1$$

$$\text{loss and threat paradigms:} \quad valence = -1$$

Value comparison in decision-making

Category 1: Models in which magnitude differences and probability differences are evaluated separately (MP models)

Prior research in threat-avoidant ambiguous decision-making has indicated that individuals may compare the relative magnitude of the two urns separately from the relative probability of the two urns. Here we consider the class of models where urn magnitude difference (M_{diff}) and urn probability difference (P_{diff}) are given different parametric weights in the decision-making process (Lawrance et al., 2022).

Base Model: MP_c_A

The simplest MP model considered in class 1 consists of one parameter for differences in magnitude (β_M), one parameter for differences in probability (β_P), and two parameters to characterize individual differences in sensitivity to choice ambiguity (β_C, β_A) representing categorical ambiguity bias and sensitivity to level of missing information, respectively.

The magnitude and probability differences between urns on each trial are calculated as:

$$M_{diff} = (M_U - M_A) \cdot m_{scale} \cdot valence \quad 1(a)$$

$$P_{diff} = (P_U - P_A) \cdot valence \quad 1(b)$$

where M_U is the magnitude of the unambiguous urn, M_A is the magnitude of the ambiguous urn, P_U is the probability of drawing an 'X' token from the unambiguous urn, and P_A is the updated estimate of the probability of drawing an 'X' token from the ambiguous urn (where P_A is updated to equal the expected value of an unbiased Bayesian posterior distribution). The constant term m_{scale} is equal to the inverse of the maximum magnitude possible during the task (150), so $m_{scale} = 1/150$ and scaled magnitudes range from 0 to 1.

For unambiguous trials, M_U and P_U refer to M_{Urn1} and P_{Urn1}, while M_A and P_A refer to M_{Urn2} and P_{Urn2}. Urns on unambiguous trials are randomly pre-assigned to be Urn1 or Urn2, counterbalanced by side.

Additionally, ambiguous trials are characterized by the $CatAmbi$ indicator, which is set to 1 if a trial is ambiguous and 0 otherwise.

Finally, we include a parameter to quantify the impact of level of missing information on decision-making, where A represents the level of missing information for an ambiguous trial. Each urn has a maximum of 50 tokens visible, and n is defined as the total number of tokens visible for the given urn. A is then defined as:

$$A = 1 - \sqrt{n/50} \quad 2$$

This formulation for A is nonlinear for increasing levels of missing information, but approximately linear in resulting increase in posterior uncertainty (per Lawrance et al., 2022).

The probability of selecting the unambiguous urn P(U) in Model MP_c_A is then defined as:

$$P(U) = \frac{1}{(1+\exp\left(-(\beta_C CatAmbi+\beta_A A + \beta_M M_{diff}+ \beta_P P_{diff})\right))} \quad 3$$

Category 2: Expected utility models with subjective nonlinear evaluation of magnitude (EUM models)

The second category of models assumed that individuals calculate an expected utility for each decision option prior to value comparison. This class of models (EUM) incorporated a nonlinear parameter on outcome magnitude as suggested by prior economics research (Kahneman and Tversky, 1979).

Base Model: EUM_c_A

The simplest EUM model considered in category 2 consists of one parameter for nonlinear weighting of magnitude in expected utilities, one parameter for differences in expected utility between choices, one parameter to characterize individual bias against choosing an ambiguous option, and one parameter to quantify the impact of level of missing information on decision-making. The expected utility for each urn was calculated as:

$$EUM_A = (M_A \cdot m_{scale})^{\lambda} P_A / m_{scale} \quad 4$$

$$EUM_U = (M_U . m_{scale})^{\lambda} P_U / m_{scale} \quad 5$$

and the difference in expected utilities was then calculated as:

$$EUM_{diff} = (EUM_U - EUM_A) \cdot valence \quad 6$$

As in the previously defined MP models, here M_U is the magnitude of the unambiguous urn, M_A is the magnitude of the ambiguous urn, P_U is the probability of drawing an 'X' token from the unambiguous urn, and P_A is the updated estimate of the probability of drawing an 'X' token from the ambiguous urn (where in class 1 models, P_A is updated to equal the expected value of the unbiased Bayesian posterior distribution). For unambiguous trials, M_U and P_U refer to M_{Urn1} and P_{Urn1}, while M_A and P_A refer to M_{Urn2} and P_{Urn2}. Urns on unambiguous trials are randomly pre-assigned to be Urn1 or Urn2, counterbalanced by side.

The scaling constant m_{scale} is applied to scale magnitudes between [0,1] prior to application of the nonlinear parameter λ, and the resulting subjective magnitude is then scaled back to the original range [0,150] prior to expected utility calculation to improve parameter estimation.

The probability of choosing the unambiguous urn is then calculated as:

$$P(U) = \frac{1}{(1+\exp\left(-(\beta_C CatAmbi + \beta_A A + \beta_{EUM} EUM_{diff})\right))} \quad 7$$

The parameter λ may be interpreted as a nonlinear under- or over-weighting of magnitude relative to objective magnitude on offer (see Figure 3.4). Alternatively, λ may be interpreted according to expected utility theory as a measure of risk aversion in decision-making. When $\lambda < 1$ the resulting expected utility curve for positive outcomes (i.e., reward task) is concave, and the expected utility curve for negative outcomes (i.e., loss, threat tasks) is convex. Analagously, when $\lambda > 1$ the expected utility curve for positive outcomes in convex, and for negative outcomes is concave. Consistent with foundational literature in behavioral economics, concave expected utility curves correspond to risk-avoidant behavior, while convex expected utility curves correspond to risk-seeking behavior (see Figure 3.4).

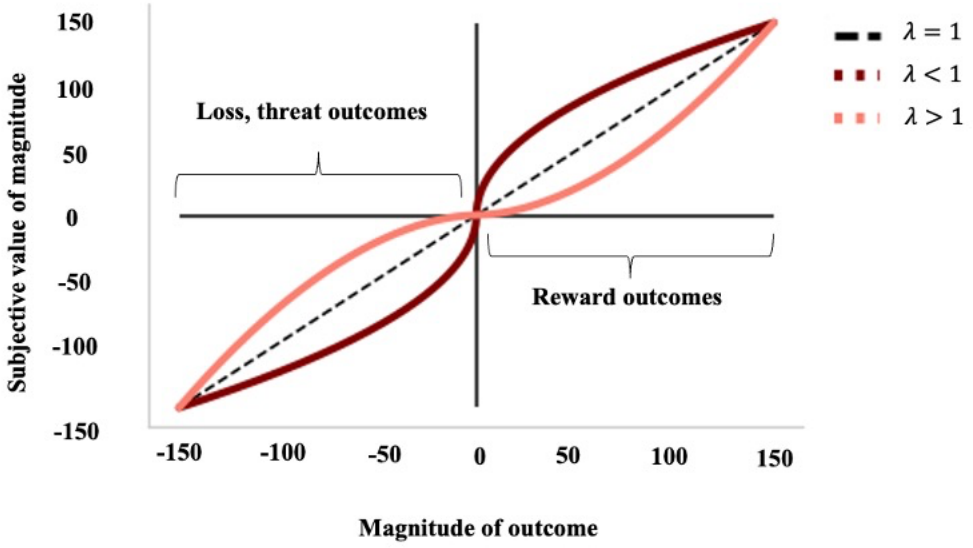


Figure 3.4 Illustration of nonlinear subjective valuation of outcome magnitude. Points outcomes (which range from [0,150] in the task) are first rescaled by $m_{scale} = 1/150$ to the range [0,1]. $(M \cdot m_{scale})^{\lambda} / m_{scale}$ then characterizes the rescaled utility curve, where the parameter $\lambda \in [0,2]$. Magnitudes are then signed for valence (negative outcomes for loss and threat tasks, positive outcomes for reward task). For positive outcomes, $\lambda > 1$ corresponds to a convex valuation curve and describes risk-seeking behavior, while $\lambda < 1$ corresponds to a concave valuation curve and describes risk-avoidant behavior. For negative outcomes, $\lambda > 1$ corresponds to a concave valuation curve and describes risk-avoidant behavior, while $\lambda < 1$ corresponds to a convex valuation curve and describes risk-seeking behavior.

Category 3: Expected utility models with subjective nonlinear evaluation of magnitude and probability (EUMP models)

Recent research on large-sample data sets leveraging neural networks to assess risky decision-making behavior in humans has found that individuals may employ separate nonlinear weighting of magnitude and probability during risky decision-making (Peterson et al., 2021). In category 3 models, we incorporate this possibility parametrically as outlined below.

Base Category 3 Model: EUMP_c_A

For category 3 models, we augmented the EUM model described above to include an additional parameter quantifying subjective nonlinearity of outcome probability as follows:

$$EUMP_A = (M_A \cdot m_{scale})^{\lambda}({P_A}^{\eta})/m_{scale} \quad 8$$

$$EUMP_U = (M_U.m_{scale})^{\lambda}({P_U}^{\eta})/m_{scale} \quad 9$$

and the difference in expected utilities was then calculated as:

$$EUMP_{diff} = (EUMP_U - EUMP_A) \cdot valence \quad 10$$

The probability of choosing the unambiguous urn is then calculated as:

$$P(U) = \frac{1}{(1+\exp\left(-(\beta_C CatAmbi+\beta_A A + \beta_{EUMP}EUMP_{diff})\right))} \quad 11$$

Category 4: Fully interactive models (FI models)

Finally, we considered models in which magnitude differences and probability differences could be considered both separately (additively) and jointly (multiplicatively).

Base Category 4 Model: FI_c_A

The simplest FI model considered in category 4 consists of one parameter for differences in magnitude (β_M), one parameter for differences in probability (β_P), one parameter for differences in expected value (β_{EV}), defined for each urn as magnitude times probability, and two parameters to characterize individual differences in sensitivity to choice ambiguity (β_C, β_A) representing categorical ambiguity bias and sensitivity to level of missing information, respectively.

As in category 1 models, the additive magnitude and probability differences between urns on each trial are calculated as:

$$M_{diff} = (M_U - M_A) \cdot m_{scale} \cdot valence \quad 12(a)$$

$$P_{diff} = (P_U - P_A) \cdot valence \quad 12(b)$$

For FI models, an additional term allowing the calculation of expected value prior to comparison was defined as:

$$EV_{diff} = [(M_U \cdot P_U) - (M_A \cdot P_A)] \cdot valence \quad 13$$

The probability of selecting the unambiguous urn P(U) in Model FI_c_A is then defined as:

$$P(U) = \frac{1}{(1+\exp\left(-(\beta_C CatAmbi + \beta_A A + \beta_M M_{diff} + \beta_P P_{diff} + \beta_{EV} EV_{diff})\right))} \quad 14$$

Models in category 4 can be additionally augmented to include expected utility parameters (i.e., incorporating "EUM" or "EUMP" parameters described above); in each case, the FI parameters for additive M_{diff} and P_{diff}, as well as the simple EV_{diff} defined above are included in addition to the expected utility comparisons specified by the "EUM" or "EUMP" variants, which incorporate nonlinearities as specified.

Model variants: Estimating probabilities under ambiguity

In addition to the overall model categories outlined above which describe the fundamental mechanisms of value comparison that individuals use to make decisions, we tested two potential mechanisms for the inference of probability under ambiguity. These additional variants are denoted as outlined below and can each be used in combination with any of the model categories described above.

Base variant: unbiased Bayesian updating of ambiguous probabilities (no special notation included on models)

When presented with an ambiguous urn (i.e., an urn in which at least some of the information is masked), an unbiased Bayesian decision-maker should combine the data available to them (the actual number of X's and O's shown for the urn and the total number of tokens visible) with an unbiased prior belief about the distribution of the hidden tokens. In the current task, participants are instructed that the distribution of hidden tokens is randomly positioned prior to masking, and cannot be in any way inferred from the information shown.

For the unbiased Bayesian estimate, we assume a uniform prior belief (~Beta(1,1)) over the distribution of hidden tokens in the ambiguous urn. Participants then update this prior belief using the likelihood of the actual probability revealed and the level of missing information. The posterior distribution for the probability of the ambiguous urn is then characterized by ~Beta(a,b), where parameters a and b are defined as follows:

$$k = \text{number of 'X' tokens revealed}$$
$$n = \text{total number of tokens revealed}$$

$$a = 1 + k$$
$$b = 1 + n - k$$

The unbiased Bayesian probability estimate for the ambiguous urn is then given by the expected value of the posterior distribution:

$$P_{A_Bayes} = E[P_{A_posterior}] = \frac{a}{a+b} = \frac{1+k}{2+n} \quad 15$$

In all baseline variant models, participants are assumed to fully update their estimate of ambiguous probabilities in accordance with the ideal Bayesian update. Thus for all models which do not specify a different mechanism for the estimation of ambiguous probability:

$$P_A = P_{A_Bayes} \quad 16$$

Variant "_OP": Models in which participants parametrically update estimates of ambiguous probability towards optimistic or pessimistic beliefs

The next model variant considered cases in which individuals parametrically (partially or fully) update estimates of ambiguous probability in accordance with intrinsic optimistic or pessimistic biases which influence their estimates. These models are denoted by "_OP" in model names to indicate the use of the optimistic/pessimistic bias.

Here, we define the purely optimistic bias as the belief that all hidden tokens are favorable (i.e., all Xs in reward valence, or all Os in threat or loss valence). The purely pessimistic bias is then defined as the belief that all hidden tokens are unfavorable (i.e., all Os in reward valence, or all Xs in threat or loss valence).

For:

$$k = \text{number of 'X' tokens revealed}$$
$$n = \text{total number of tokens revealed}$$

We then define:

$$P_{A_Pessimist} = \begin{cases} k \,/\, 50 & if\ valence = reward \quad 17(a) \\ (k + 50 - n) \,/\, 50 & if\ valence = threat\ or\ loss \quad 17(b) \end{cases}$$

$$P_{A_Optimist} = \begin{cases} (k + 50 - n) \,/\, 50 & if\ valence = reward \quad 18(a) \\ k \,/\, 50 & if\ valence = threat\ or\ loss \quad 18(b) \end{cases}$$

We parameterize the adjustment of ambiguous probability from the unbiased Bayesian value towards optimistic or pessimistic biases using parameter α which ranges from [0,2]. As the estimated value of α moves from 0 towards 2, individuals move from a perfectly pessimistic update towards a perfectly optimistic update (see Figure 3.5).

For ($0 \leq \alpha < 1$):

$$\alpha' = \alpha \quad 19(a)$$

$$P_A = (\alpha')P_{A_Bayes} + (1 - \alpha')P_{A_Pessimist} \quad 19(b)$$

When ($\alpha = 1$):

$$P_A = P_{A_Bayes} \quad 19(c)$$

For ($1 < \alpha \leq 2$):

$$\alpha' = \alpha - 1 \quad 19(d)$$

$$P_A = (1 - \alpha')\, P_{A_Bayes} + (\alpha')P_{A_Optimist} \quad 19(e)$$

Models using the "OP" variant for probability estimation in ambiguity follow the formulas otherwise defined by their class and other notation, using the updated estimate of P_A.

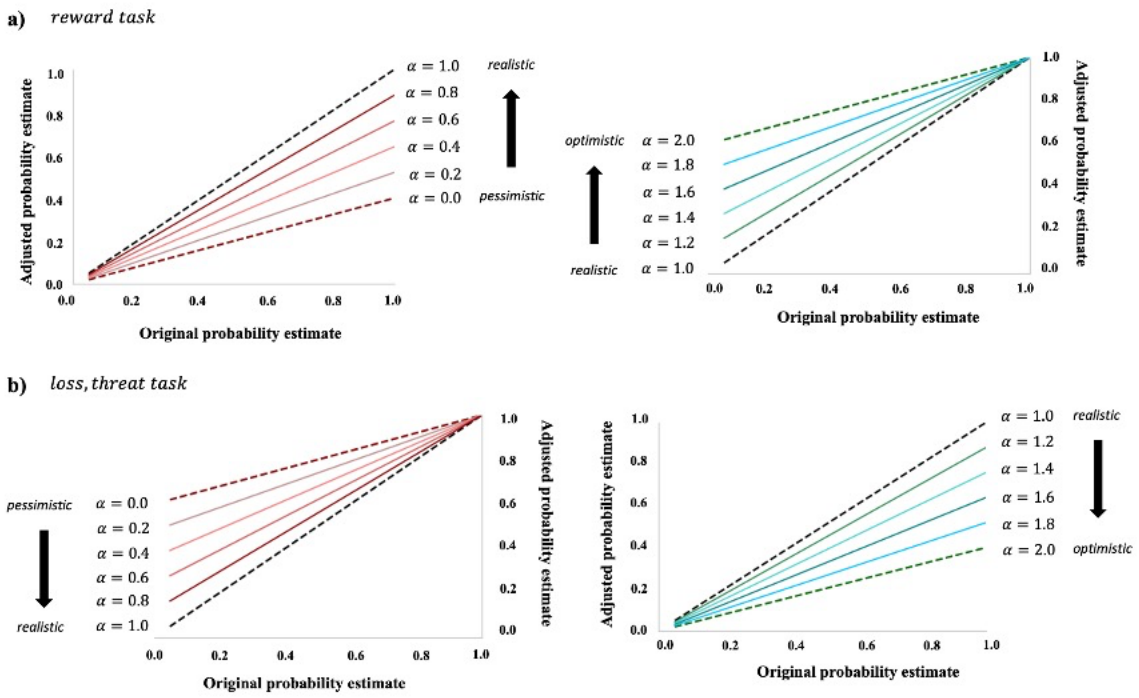


Figure 3.5 Influence of optimistic/pessimistic bias parameter α when 30 tokens are hidden. a) Reward task: higher probability estimates are optimistic (higher chance of reward), and lower probability estimates are pessimistic. As α moves from 0 to 2, individuals shift from a purely pessimistic bias towards a purely optimistic bias. b) Loss, threat task: higher probability estimates are pessimistic (higher chance of loss or threat), and lower probability estimates are optimistic. The interpretation of α is consistent with that in the reward task: as α moves from 0 to 2, individuals shift away from a purely pessimistic bias towards a purely optimistic bias.

Model classes: treatment of relative value	
_MP	Models in which magnitude differences and probability differences are considered separately.
_EUM	Models in which expected utilities are compared (with nonlinearity of magnitude valuation).
_EUMP	Models in which expected utilities are compared (with nonlinearity of both magnitude and probability valuation).
_FI	Fully interactive models in which additive and multiplicative differences in magnitudes, probabilities, and expected utilities may all contribute separately to decision-making.
Model variants: inference of ambiguous probability	
baseline	In the baseline model, participants are assumed to update estimates of probability under ambiguity towards an unbiased Bayesian adjustment of visible information.
_OP	Models in which an optimistic or pessimistic bias allows individuals to over or underweight ambiguous information.
Model notation: basic	
_c	Indicates that model contains an individual bias term categorizing the presence of ambiguity.
_A	Indicates that model parameterized the impact of level of missing information on decision-making.

Table 3.1 Model key for decision-making under ambiguity analysis.

Model fitting

Parameters were estimated at the group and individual level in each valence paradigm using a hierarchical Bayesian approach to sample posterior parameter distributions. Markov Chain Monte Carlo (MCMC) sampling was performed using RStan (Gelman et al., 2014; Stan Development Team., 2022), with two chains of 4000 samples (2000 post burn-in) for each model. Sampling convergence was determined by inspecting the visual convergence of the MCMC chains and by ensuring that the Gelman-Rubin statistic ($\hat{R}$) for parameters was very close to 1.00.

Models were fit hierarchically across participants but within valence paradigm, such that reward pursuit, threat avoidance, and loss avoidance decisions were fit separately. This allowed us to test whether a similar evaluative mechanism best-described individuals across tasks, or whether some aspect of the decision process varied meaningfully by paradigm. See Results for additional details.

Priors were set to be generally uninformative as ~Normal(0,5) or ~Normal(0,10). Where parameters were constrained to be greater than 0, these constraints were applied at the parameter level, which results in stan converting these prior distributions to the renormalized equivalent of the associated Half-Normal distributions.

Model comparison

We compared the fit of each model to the data using the Widely Applicable Information Criterion (WAIC, sometimes referred to as the Watanabe-Akaike Information Criterion; Watanabe S, Opper M, 2010). This measure assesses model fit while penalizing for model complexity in a hierarchical Bayesian framework to mitigate the likelihood of overfitting (Gelman et al., 2014).

Previous literature has pointed to potentially differing mechanisms for decision-making under ambiguity across reward pursuit, loss avoidance, and threat avoidance paradigms. Behavioral economics models in reward and loss tasks, as well as cognitive science, have generally found expected utility models to best fit human subject data, whereas behavioral research in threat avoidance has found strong evidence that participants weigh up magnitude differences versus probability differences separately (Kahneman and Tversky, 1979; Lawrance et al., 2022; Peterson et al., 2021). We compared model fit separately for each valence in order to allow for potential variations in the best-fit algorithms. For each paradigm, all hypothesis testing was performed using the model with the lowest WAIC for that valence.

Model validation

Model validation was performed by simulating data using the generative inversion of the winning model in each paradigm (Model FI_EUMP_OP_c_A). We used mean posterior parameter values for each individual to simulate data. We then compared simulated ambiguity sensitivity at each level of missing information between actual behavioral data and model-simulated data.

We also performed parameter recovery analysis for Model FI_EUMP_OP_c_A in each valence paradigm. Simulated data was fit using hierarchical Bayesian inference over posterior parameter values; we then assessed whether the 95% highest density interval (HDI) of the posterior distribution for each parameter contained the true generating value.

Results

Behavioral analysis

Prior to computational modeling of decision-making behavior, we performed model-independent analysis of the behavioral results in each paradigm for the final in-lab participants (n=71). Across all paradigms, participants selected the unambiguous urn on a greater percentage of trials (across all trials) than the ambiguous urn, despite balanced objective expected value in average (reward pursuit: % unambiguous choices mean =67.79%, SEM = 0.64%; loss avoidance: % unambiguous choices mean = 53.40%, SEM = 0.67%; threat avoidance, % unambiguous choices mean = 53.71%, SEM = 0.65%). However, when ambiguous trial choices were split by level of missing information, participants in both loss avoidance and threat avoidance were ambiguity-seeking at low levels of information, selecting the ambiguous urn more than half the time on these trials (see Figure 3.6). This is consistent with the classic gain-loss framing in first-order uncertainty paradigms of risky decision-making, wherein individuals are risk avoidant in reward pursuit and risk seeking in loss avoidance (Kahneman and Tversky, 1979), which has also been found to hold for decisions made under ambiguity (Kocher et al., 2018). As levels of ambiguity increased in loss and threat avoidance, participants switched behavior to be ambiguity avoidant (see Figure 3.6).

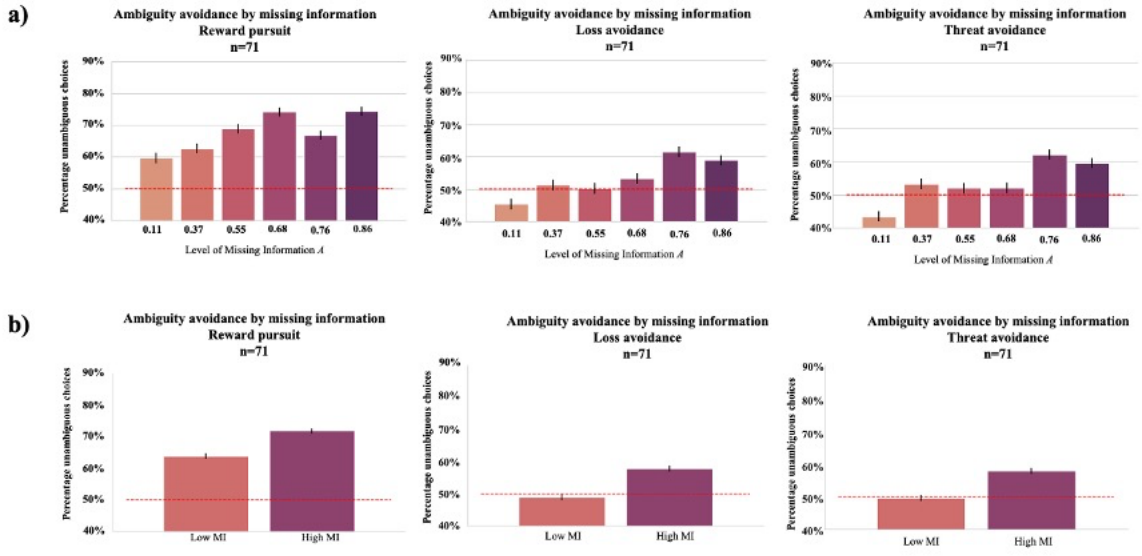


Figure 3.6 Ambiguity avoidance measured by percentage of choices for unambiguous urn. a) Percentage of unambiguous urn choice for each level of missing information *A*. b) Percentage of unambiguous urn choice for low (0.11, 0.37, and 0.55) versus high (0.68, 0.76, 0.86) levels of missing information *A*.

To investigate individual differences in sensitivity to missing level of information, we calculated the simple Information Level Dependent Ambiguity Avoidance (s-ILDAA) metric, defined as the

slope of the regression of unambiguous choice on level of missing information (following Lawrance et al., 2022). Previous research has found that higher levels of trait anxiety (as measured by the STAI-T) were associated with a greater sensitivity to increasing levels of missing information (i.e., a higher s-ILDAA) in a threat avoidance ambiguous decision-making paradigm (Lawrance et al., 2022).

For all paradigms, the group distribution of s-ILDAA slope values was significantly greater than 0, providing a convergent check that across all paradigms individual ambiguity aversion increased as the level of missing information increased (see Figure 3.7; reward pursuit: median s-ILDAA slope = 0.19, Wilcoxon(70) = 222.0, p = 0.000; loss avoidance: median s-ILDAA slope = 0.16, Wilcoxon(70) = 255.0, p = 0.000; threat avoidance: median s-ILDAA slope = 0.19, Wilcoxon(70) = 243.0, p = 0.000) .

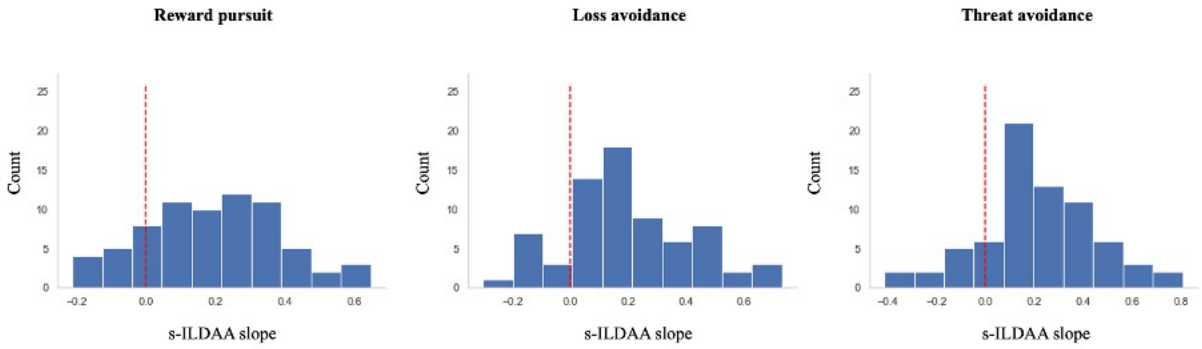


Figure 3.7 s-ILDAA slope. s-ILDAA slope(from regression of unambiguous choice on level of missing information) for each ambiguity avoidance paradigm. In reward pursuit, loss avoidance, and threat avoidance, s-ILDAA slope is significantly greater than 0, indicating that higher levels of missing information are associated with greater ambiguity avoidance.

Model comparison

We fit each model as described in Methods to choice data for each paradigm (reward pursuit, loss avoidance, and threat avoidance) using hierarchical Bayesian inference, and assessed model fit using the widely applicable information criterion (WAIC; Watanabe S, Opper M, 2010). The WAIC quantifies accuracy of fit while penalizing for model complexity within a hierarchical Bayesian framework.

The best-fit model for all three valence paradigms was model FI_EUMP_OP_c_A (reward pursuit: WAIC = 8,770.07, SE = 119.15; loss avoidance: WAIC = 8,816.41 , SE = 118.38; threat avoidance: WAIC = 9,008.64 , SE = 121.17). This model combined the classical findings of prospect theory, including subjective nonlinearity on both magnitude and probability in expected utility calculation (EUMP), with additional parameters to allow separate comparisons to be made for magnitude differences and probability differences as well (FI). Additionally, the model posits that individuals have optimistic or pessimistic biases (OP) which influence their estimates of probability under ambiguity away from the unbiased Bayesian estimate. Finally, the model parameterizes both the

influence of categorical ambiguity (i.e., whether ambiguity is present or not; 'C'), and the level of missing information ('A').

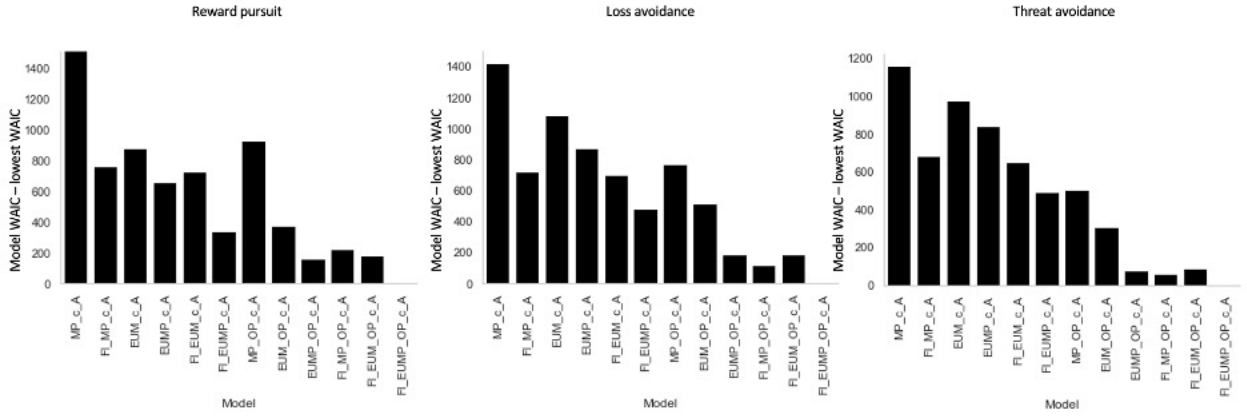


Figure 3.8 Model comparison. Model fit for each paradigm (reward pursuit, loss avoidance, and threat avoidance) was evaluated using the Widely Applicable Information Criterion (WAIC). The same model (Model FI_EUMP_OP_c_A) was the best fit in all valences.

The final parameters of the model (fit separately for each valence) were as follows: categorical ambiguity aversion ($\beta_{CatAmbi}$); sensitivity to level of missing information ($\beta_{LevelMI}$); risk preference (nonlinear subjective valuation of magnitude λ); nonlinear subjective valuation of probability (ω); weight on magnitude difference between urns (β_{MDiff}); weight on probability difference between urns (β_{PDiff}); weight on expected utility difference between urns (β_{EUDiff}); and optimistic or pessimistic inferential bias (α_{Opt_Pess}). See Methods for full model equations. Supplements for distributions of each parameter in each valence.

Model validation

We used the winning computational model for each valence paradigm (Model FI_EUMP_OP_c_A in each valence) to simulate participant data, and compared simulated data to the actual behavioral data for study participants. The model produced a strong replication of participant behavior: see Figure 3.9 for illustration.

We also performed parameter recovery analysis for each parameter of the winning model in each valence paradigm. Hierarchical Bayesian inference was used to estimate the posterior distribution of each parameter, and we then assessed whether the 95% highest density interval (HDI) of the recovered posterior distribution contained the true generating value. Parameters overall recovered fairly well: mean percentage of generative parameters contained in the recovered posterior 95% HDI was 94% for reward pursuit, 98% for loss avoidance, and 98% for threat avoidance paradigms. See Supplements S-5 for full parameter recovery analysis.

We note that although overall parameter recovery was generally good, some specific parameters were unreliably recovered, particularly in the reward pursuit valence (β_{MDiff} recovered 75%). The winning model in the current analysis combines previously successful models for economic choice, including both expected utility estimation and separate magnitude and probability differences.

Parameter recovery should be tested in the simpler individual models in future analysis to ensure that the recovery for the current model is not due to a lack of sufficient power to fit the proposed model structure (see Limitations).

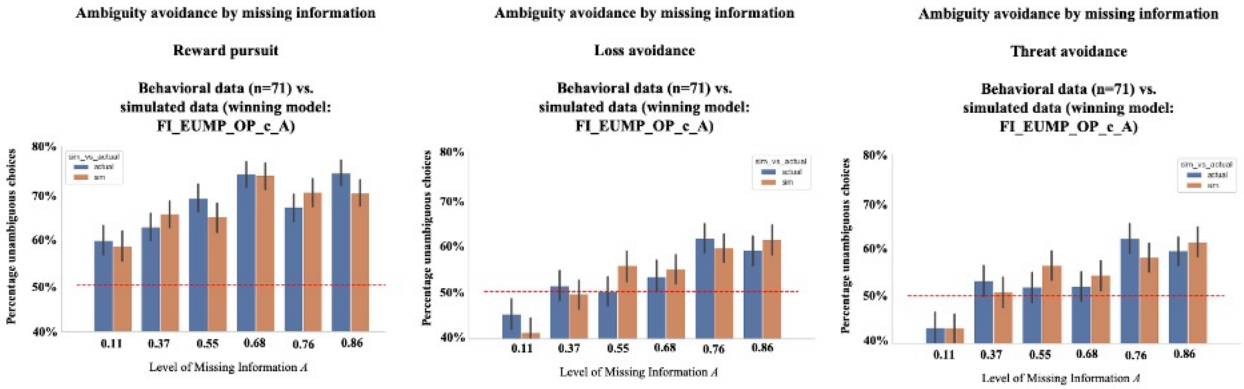


Figure 3.9 Simulated vs. actual data. Behavioral versus simulated data from winning model (Model FI_EUMP_OP_c_A) in each valence paradigm. Simulated data reproduces participant ambiguity avoidance by level of missing information in each valence.

Categorical ambiguity sensitivity versus level of missing information

To assess sensitivity to both categorical ambiguity and level of missing information, we investigated the group-level distribution of posterior parameter means for $\beta_{CatAmbi}$ and $\beta_{LevelMI}$ from the best-fit model in each paradigm (Model FI_EUMP_OP_c_A for all paradigms). In the reward gain paradigm, individuals were broadly avoidant of categorical ambiguity ($\beta_{CatAmbi}$ mean = 0.35; one-sample >0 t-test(70) = 11.266, p = 0.000; see Figure 3.10) and were also significantly sensitive to increases in level of missing information ($\beta_{LevelMI}$ mean = 2.02; one-sample >0 t-test(70) = 10.091, p = 0.000; see Figure 3.10). Individual sensitivity to categorical ambiguity and level of missing information were highly correlated with one another (Spearman rho = 0.502, p = 8.36e-6), consistent with a broad ambiguity avoidance preference in reward gain decisions (see Figure 3.10).

In contrast with reward gain decisions, participants were categorically ambiguity-seeking (when level of missing information was parameterized separately, as in the best-fit model) in both the loss ($\beta_{CatAmbi}$ mean = -0.45; one-sample < 0 t-test(70) = -33.246, p = 0.000) and threat ($\beta_{CatAmbi}$ mean = -0.50; one-sample < 0 t-test(70) = -41.504, p = 0.000) avoidance paradigms. As in the reward paradigm, increases in levels of missing information did cause individuals to become more ambiguity avoidant in both loss ($\beta_{LevelMI}$ mean = 0.87; one-sample >0 t-test(70) = 4.291, p = 0.000) and threat ($\beta_{LevelMI}$ mean = 0.93; one-sample >0 t-test(70) = 4.765, p = 0.000) avoidance. The categorical ambiguity seeking behavior in loss avoidance is consistent with behavioral economics findings that individuals become risk seeking and ambiguity seeking in loss avoidance paradigms (Kahneman and Tversky, 1979; Kocher et al., 2018; Machina et al., 2011) and with previous work by our lab (unpublished data).

As in the reward pursuit paradigm, individual sensitivity to categorical ambiguity and level of missing information were positively correlated with each other in both loss (Spearman rho = 0.274, p = 0.021) and threat (Spearman rho = 0.443, p = 0.0001) avoidance paradigms (see Figure 3.10).

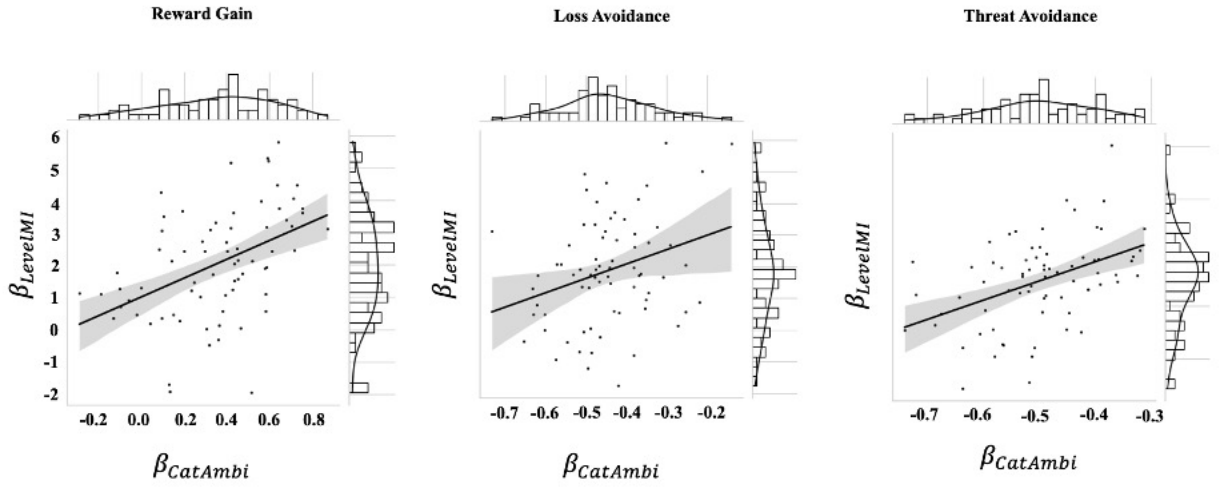


Figure 3.10 Relationship of categorical ambiguity aversion and sensitivity to level of missing information. Individual sensitivity to the presence of ambiguity ($\beta_{CatAmbi}$) was significantly correlated with individual sensitivity to level of missing information ($\beta_{LevelMI}$) in each valence paradigm.

Individual risk preference across valences

The winning computational model for all paradigms was Model FI_EUMP_OP_c_A, which includes a parameter traditionally used in behavioral economics literature to quantify an individual's risk tolerance in risky (i.e., probabilistic) decision-making.

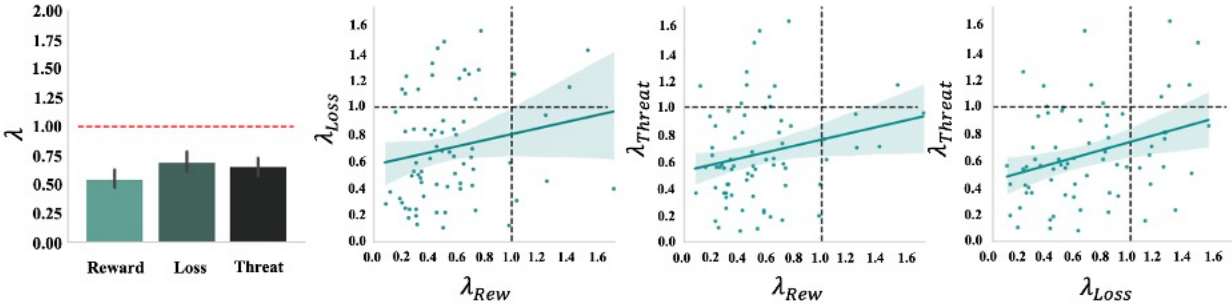


Figure 3.11 Participant risk-preference parameters λ for reward pursuit, loss avoidance, and threat avoidance decision-making under ambiguity. Parameter values less than 1.0 in reward decisions (where outcome valence is positive) indicate a participant is risk avoidant in those decisions, whereas parameter values less than 1.0 in loss or threat decisions (where outcome valence is negative) indicate a participant is risk-seeking in those decisions. Mean participant parameter values were below 1.0 for all three valences, reproducing classical economics findings that individuals tend to be risk-avoidant in reward but risk-seeking in loss. Individual risk preference parameters were significantly correlated across valence paradigms.

In scenarios with positively valenced outcomes (i.e., reward), a risk preference parameter smaller than 1 results in a concave subjective expected utility function and is generally characterized as

risk-avoidant behavior (Kahneman and Tversky, 1979; Tversy and Kahneman, 1974). In scenarios with negatively valenced outcomes (i.e., loss, threat), a risk preference parameter smaller than 1 results in a convex subjective expected utility function over the negative magnitude domain, and indicates risk-seeking behavior. We found that mean risk preference parameters were < 1 for each paradigm. Our results for reward and loss are consistent with past results that individuals tend to be risk-avoidant in reward decisions and risk-seeking in loss decisions (Fox et al., 2015; Kahneman and Tversky, 1979). Risky decision-making preferences in threat were highly correlated with risky decision-making preferences in loss (λ_{Loss}, λ_{Threat}; Spearman rho = 0.301, p = 0.011), indicating that individuals who were risk-seeking in loss avoidance tended to be similarly risk-seeking in threat avoidance. Interestingly, risk preference in threat avoidance was also significantly positively correlated with risk preference in reward (λ_{Rew}, λ_{Threat}; Spearman rho = 0.237, p = 0.047), indicating that individuals who were risk avoidant in reward tended to flip behavior towards risk-seeking in threat. There was no significant correlation between risk preference in reward and loss (λ_{Rew}, λ_{Loss}; Spearman rho = 0.223, p = 0.062), although we note that the relationship is trending towards the same pattern as between reward and threat risk preference.

Optimism/pessimism biases

Individual biases towards optimistic or pessimistic inferences of hidden information were captured in our computational model by parameter α_{opt_pess} (see Methods). Values of $0 < \alpha_{opt_pess} < 1$ imply a pessimistic bias, while values of $1 < \alpha_{opt_pess} < 2$ imply an optimistic bias in all paradigms. There were significant individual differences in optimistic/pessimistic biases in each valence paradigm. In reward, participants tended on average to be slightly optimistic, with a group mean parameter greater than 1 ($\alpha_{opt_pess_reward} = 1.05$). At the group level, however, the bias parameter was not significantly different from 1 (Wilcoxon statistic= 1143.0, p = 0.439).

In loss avoidance decisions, the group mean was less than 1 ($\alpha_{opt_pess_loss} = 0.85$), implying slightly pessimistic inference for hidden probability information in loss decisions. At the group level, the parameters were significantly pessimistic (Wilcoxon statistic = 798.0, p = 0.003). Similarly, participants showed a group-level pessimistic inferential bias for threat avoidance decisions, with mean group parameter less than 1 ($\alpha_{opt_pess_threat} = 0.81$) and individual parameters significantly less than 1 at the group level (Wilcoxon statistic = 718.0, p = 0.001).

Individual optimism pessimism biases were not correlated between reward and loss tasks (Spearman rho = 0.072, p = 0.553) or between reward and threat tasks (Spearman rho = 0.099, p = 0.412). Optimistic/pessimistic biases between the two negatively valenced tasks (loss and threat avoidance) were significantly correlated (Spearman rho = 0.342, p = 0.003), indicating that optimistic or pessimistic inference of missing information may be shared across negatively valenced paradigms.

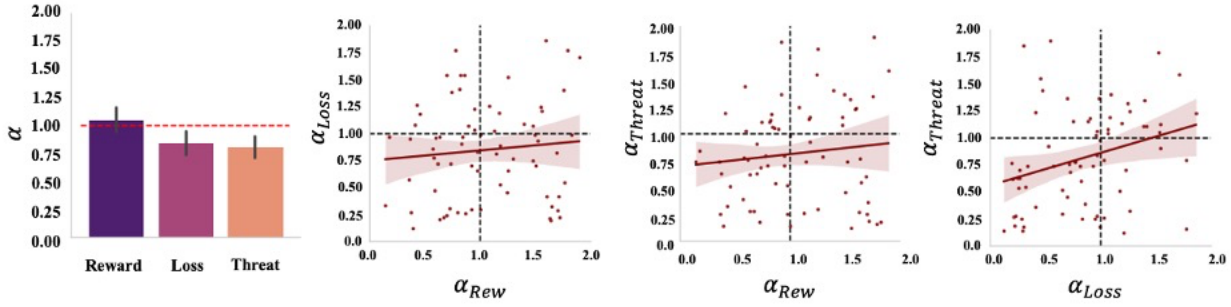


Figure 3.12 Participant optimistic/pessimistic inferential bias parameters α for reward pursuit, loss avoidance, and threat avoidance decision-making under ambiguity. Parameter values less than 1.0 indicate a pessimistic bias in inference over missing information, while parameter values greater than 1.0 indicate an optimistic bias. Mean participant parameter values were below 1.0 for loss avoidance and threat avoidance, indicating an overall pessimistic bias at the group level. Moreover, individual parameter values for loss avoidance and threat avoidance were significantly correlated. Overall parameter values for reward pursuit were not significantly different from 1.0 and were not correlated with either loss avoidance or threat avoidance α values.

Individual differences in optimistic or pessimistic inferential biases were not significantly associated with factor scores from the BI-4 model across the anxiety/depression/mania trait space. See Supplements Figure 3.S.4.

Relationship of mania with ambiguity aversion in decision-making

Our first aim was to test the relationship between hypomania and increased risk or ambiguity tolerance in a reward-based ambiguous decision-making paradigm. We tested individual scores on the latent factor characterizing hypomanic symptoms against parameters describing individual differences in categorical ambiguity sensitivity, sensitivity to level of missing information, and risk preference.

Participant scores on the factor score characterized by mania (F3) were not significantly related to categorical ambiguity sensitivity ($\beta_{CatAmbi}$; Spearman rho = 0.16, p = 0.183), level of missing information ($\beta_{LevelMI}$; Spearman rho = 0.18, p = 0.138), or risk preference parameter (λ; Spearman rho = 0.00, p = 0.980) in reward-based decision-making under ambiguity. As the factor characterization shared some variance with physiological arousal symptoms of anxiety, we additionally tested the relationship of scores on the hypomanic personality scale (HPS) with the same parameters. Scores on the HPS were not significantly related to categorical ambiguity sensitivity ($\beta_{CatAmbi}$; Spearman rho = 0.131, p = 0.274), level of missing information ($\beta_{LevelMI}$; Spearman rho = 0.167, p = 0.163), or risk preference (λ; Spearman rho = 0.041, p = 0.736).

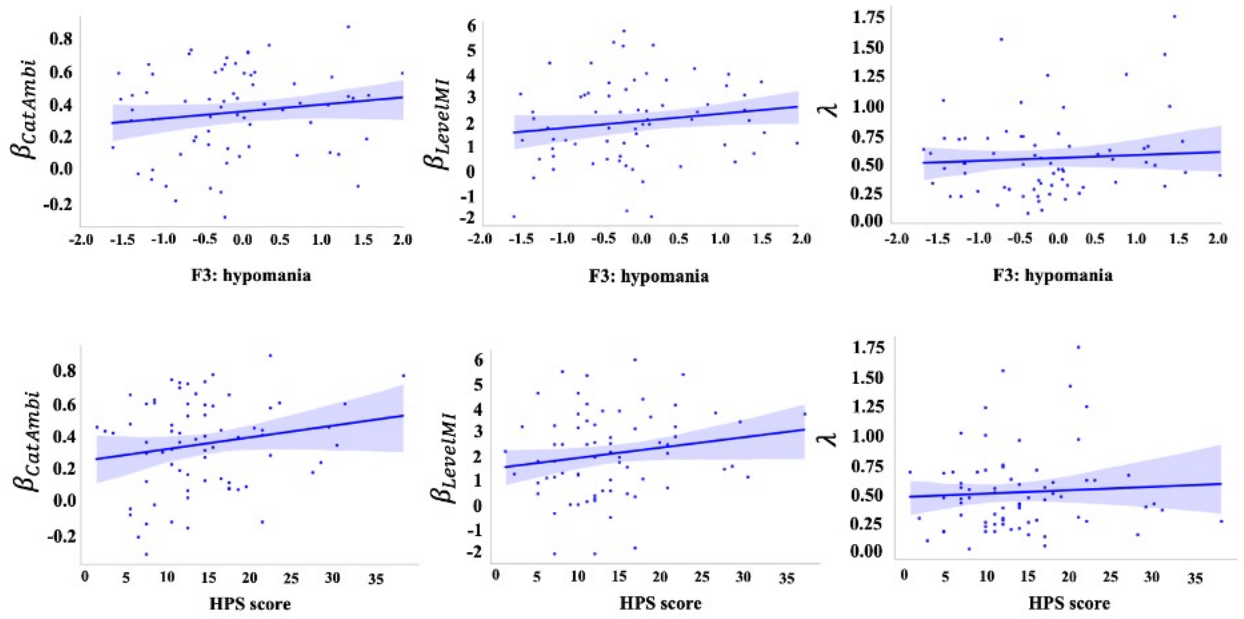


Figure 3.13 Relationship of hypomania symptom levels with ambiguity and risk sensitivity. Hypomania as characterized by latent dimension F3 (first row) and HPS score (second row) had no relationship with categorical ambiguity sensitivity, level of missing information, or risk preference in reward pursuit.

Across the three versions of the decision-making under ambiguity task, factor scores on hypomania (F3) were not significantly related to differences in decision-making under ambiguity.

We noted a trend-level result in threat avoidance decision-making, where higher scores on F3 (hypomania) were associated with a reduced consideration of probability differences between urns during ambiguity avoidant decisions (β_{PDiff}; Spearman rho = -0.21, p = 0.076), in effect, more willing to choose the riskier option. If found to be significant in subsequent studies, this would indicate that individuals with symptoms of hypomania may underweight the consideration of outcome probability in threat avoidance decision-making under second-order uncertainty. Intriguingly, scores on the factor characterized by anhedonic depression (F1) had a trending relationship in the opposite direction with threat avoidance probability differences, in effect suggesting greater risk sensitivity (β_{PDiff}; Spearman rho = 0.21, p = 0.072) in the same decisions. This trending dissociation between depression and hypomania might hint that risk sensitivity in threat avoidance provides a differentiable characterization between the two traits; we therefore further examined the difference in correlation of probability sensitivity β_{PDiff} with hypomania versus anhedonic depression. We performed a bootstrap analysis with 100,000 resamplings of the difference in Spearman correlation between factor F3 scores and β_{PDiff} values relative to the Spearman correlation between factor F1 scores and β_{PDiff} values. The 95% confidence interval of the distribution of differences excluded 0 (95% CI = [-0.72, -0.10]), confirming a significantly lower weighting of probability differences in threat avoidance decisions between dimensions characterized by hypomania versus anhedonic depression (i.e., individuals with higher levels of hypomania place significantly less weight on taking the safer option than individuals with higher levels of depression).

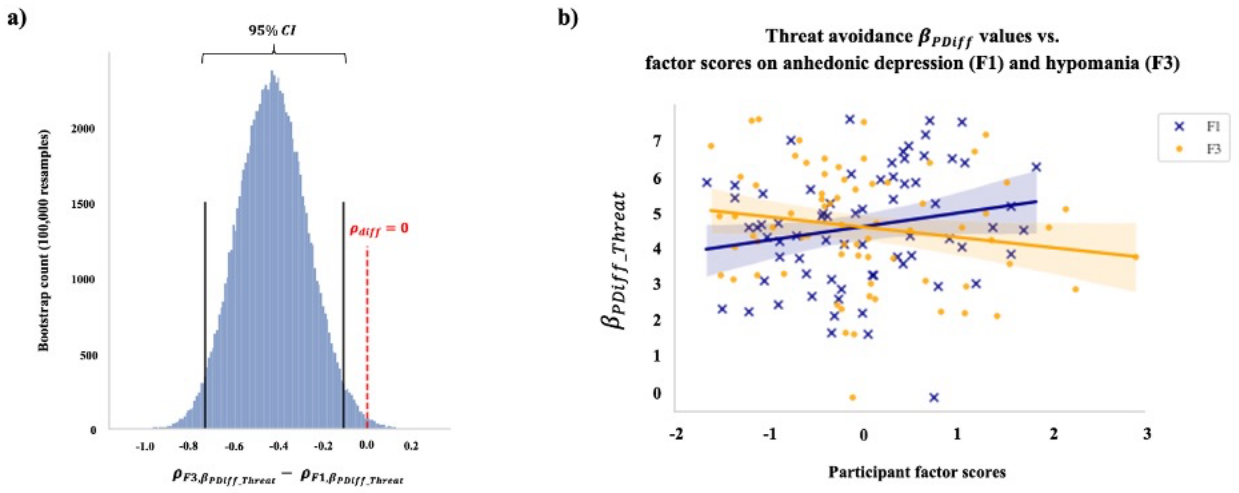


Figure 3.14 Difference between effects of hypomania versus anhedonic depression on probability sensitivity in threat avoidance. a) Bootstrapped (100,000 iterations) difference between correlation of F3 scores with β_{PDiff} and correlation of F1 scores with β_{PDiff} shows significant difference (95% CI excludes 0). High levels of hypomania are less sensitive to riskiness of choice options in threat avoidance than high levels of anhedonic depression. b) Threat avoidance parameter values for probability sensitivity β_{PDiff} for participant factor scores on anhedonic depression (F1) show increased weighting of safer probabilities for those with higher levels of depression, versus increased levels of hypomania (F3) which show reduced risk sensitivity in threat avoidance decisions.

Sensitivity to threat avoidance ambiguity levels: dimensions of physiological anxiety, cognitive anxiety, and shared internalizing variance across anxiety/depression

Our second study aim was to expand upon previous work linking anxiety to a heightened sensitivity to level of ambiguity in threat avoidance decisions. As shown in Lawrance et al., 2022, individuals with higher trait levels of anxiety as measured by the State-Trait Anxiety Inventory (STAI-T) may have a higher sensitivity to level of missing information in threat avoidant decision-making under ambiguity. Studies in learning under volatility, another paradigm of second-order uncertainty, have found a relationship between higher levels of trait anxiety (again using STAI-T) and difficulties in adapting learning rate between volatile and stable environments (Behrens et al., 2007; Browning et al., 2015); however, more recent analysis has shown that this adaptation difficulty is in fact linked to shared internalizing variance across anxiety and depression measures (Gagne et al., 2020). Given that ambiguity and volatility are both forms of second-order uncertainty, we sought to test whether shared internalizing variance across anxiety and depression, or specific dimensional characterizations of anxiety, are related to increased sensitivity to level of missing information in threat avoidant decision-making.

To investigate this, we tested factor scores on the shared variance factor G, the factor characterized by physiological arousal/anxiety (F2), and the factor characterized by cognitive anxiety/worry (F4) with parameters describing categorical ambiguity sensitivity and sensitivity to level of missing information. Scores on factor G were not significantly associated with greater categorical ambiguity sensitivity ($\beta_{CatAmbi}$; Spearman rho = 0.028, p = 0.819) or sensitivity to level of missing information ($\beta_{LevelMI}$; Spearman rho = 0.054, p = 0.652) in threat avoidant decision-making under ambiguity. Similarly, scores on the physiological arousal/anxiety factor F2 were not significantly associated with categorical ambiguity sensitivity ($\beta_{CatAmbi}$; Spearman rho = -0.19, p = 0.117) or

level of missing information ($\beta_{LevelMI}$; Spearman rho = 0.02, p = 0.847), nor were scores on the cognitive anxiety worry factor F4 ($\beta_{CatAmbi}$; Spearman rho = 0.02, p = 0.852; $\beta_{LevelMI}$; Spearman rho = -0.01, p = 0.909).

Across parameters of threat avoidance decision-making, we did not observe any significant associations with scores on G, F2, or F4. Interestingly, we observed a trend-level relationship between scores on the factor characterizing shared variance across anxiety/depression G and increased sensitivity to magnitude differences (β_{MDiff}; Spearman rho = 0.21, p = 0.080) in threat avoidance decision-making under ambiguity. We also found a significant relationship between scores on the STAI-T Anxiety Subscale with this same parameter describing sensitivity to magnitude differences in threat avoidance (β_{MDiff}; Spearman rho = 0.26, p = 0.031). Compared with previous research results relating STAI-T scores to increased sensitivity to level of missing information (Lawrance et al., 2022), we note that the computational model used in the present study is augmented by parameters describing individual differences in risk preference and expected utility weighting. The present results show that individuals with higher levels of trait anxiety as measured by the STAI-T Anxiety Subscale tend to weight differences in threat magnitude more strongly than low anxious individuals when making decisions under ambiguity. Intriguingly, the trending relationship of G with this same parameter may hint that a stronger weighting of magnitude is in fact related to shared variance across anxiety and depression; however, additional study would be needed to assess this.

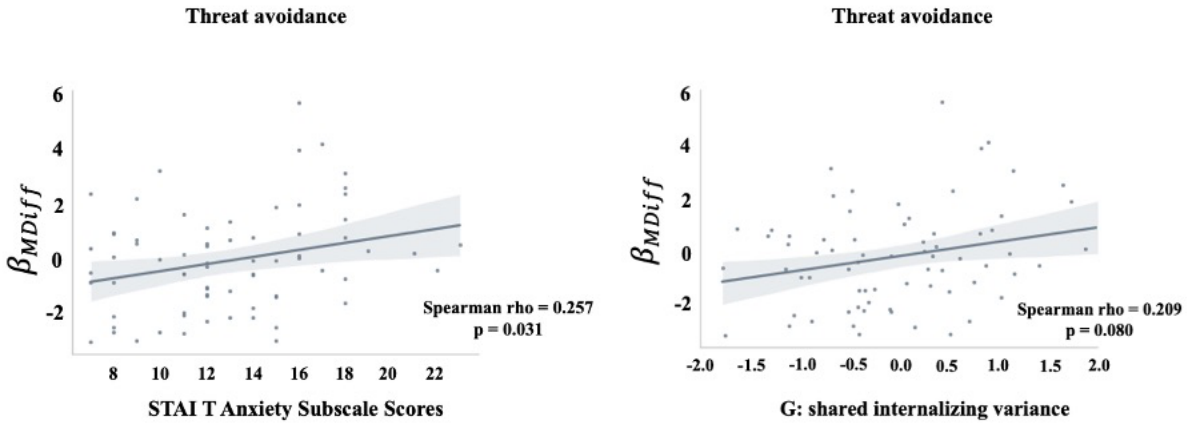


Figure 3.15 Trait anxiety and internalizing factor scores vs. sensitivity to magnitude difference in threat avoidance. Scores on the State Trait Anxiety Inventory anxiety subscale are significantly related to overweighting of magnitude differences during threat avoidance decision-making under ambiguity. When tested under a dimensional characterization of shared versus distinct variance across anxiety and depression, the shared internalizing factor G shows a statistically trending relationship with the same decision-making parameter. Additional study would be required to confirm whether this relationship reaches significance in a larger sample size.

Differential effects of anxiety versus depression on decision-making under ambiguity

In the loss avoidance version of the decision-making under ambiguity task, we found that factor scores on the anhedonic depression factor (F1) were significantly related to increased weighting of probability differences between urns (β_{PDiff}; Spearman rho = 0.28, p = 0.018), indicating that

higher levels of depression may be associated with heightened weighting of probability of negative outcomes in ambiguous loss-based decision-making (i.e., a greater risk sensitivity in loss avoidance). Interestingly, scores on factor F1 were also related at the trend level with increased weighting of probability differences in threat avoidance decision-making (β_{PDiff}; Spearman rho = 0.21, p = 0.072). If found to be significant in additional/later study, this would indicate that higher depression could be associated with increased focus on probability of negative outcomes for negatively valenced ambiguous decisions more generally.

In the loss avoidance paradigm, scores on the factor characterized by physiological arousal/anxiety (F2) were not significantly related to differences in probability difference weighting (β_{PDiff}; Spearman rho = -0.06, p = 0.628), nor were scores on the factor characterized by cognitive anxiety/worry (F4) (β_{PDiff}; Spearman rho = -0.13, p = 0.281). However, scores on the physiological anxiety factor F2 were significantly related to a decreased nonlinear probability weighting in loss-avoidance decisions (ω; Spearman rho = -0.27, p = 0.023), suggesting that anhedonic depression and physiological anxiety relate differentially to treatment of probability during loss-avoidant ambiguous decisions. Physiological anxiety (F2) was also associated with a significantly decreased categorical ambiguity aversion ($\beta_{CatAmbi}$; Spearman rho = -0.26, p = 0.031), which indicates that participants with higher levels of physiological arousal are more ambiguity seeking in loss avoidance than others.

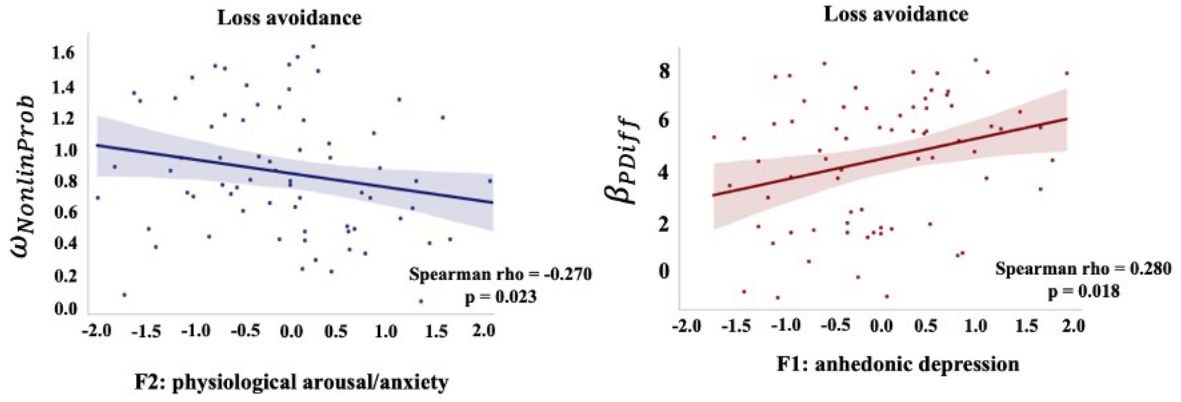


Figure 3.16 Physiological anxiety and anhedonic depression vs. measures of risk sensitivity in loss avoidance. Physiological arousal/anxiety is significantly associated with reduced nonlinear valuation of urn probability in loss avoidance decisions, whereas anhedonic depression is associated with increased linear valuation of probability differences between urns.

There were no other significant relationships between factor scores and other decision-making parameters in loss avoidance or reward pursuit paradigms.

Limitations

The present analysis was subject to several limitations which should be considered carefully for future work. Firstly, we note that the relatively small size of the data set limits our statistical power to draw conclusions. The current results would be better supported with a larger in-lab data set, a

large online data set, or both. Additional data collection for the current study is still ongoing, and we anticipate having a more highly powered data set across both in-lab and online participants for subsequent analysis. The participants recruited for the current study were not recruited based on psychiatric diagnoses. The resulting sample may consequently not reflect the full range of severity across the included symptom measures. Additionally, the measures used are focused on current or ongoing symptom levels; they therefore do not assess levels of functional impairment, and furthermore are not able to cleanly dissociate whether detected profiles are part of vulnerability to the related disorder or the direct effect of current symptoms. A longitudinal investigation would be required to address this question. We also note that there are other internalizing symptom domains not included in the present work, including social anxiety, post-traumatic stress disorder (PTSD), and obsessive-compulsive disorder (OCD) symptoms.

We note a limitation to our current data processing methodology regarding exclusions for so-called "no-brainer" (unambiguous) trials in which both magnitude and probability favor the same urn. Based on a previous study which showed participants selected the correct choice on no-brainer trials over 90% of the time (Lawrance et al., 2022), we excluded anyone from the present analysis who did not select the favorable urn on at least 85% of these trials. However, following review we note that the actual population performance on no-brainer trials for the present task was much lower, likely due to the fact that a different trial structure (with different no-brainer trials) was in place for the current study. Therefore, it is likely we artificially reduced power in the current analysis by excluding individuals according to the current procedures. In final analysis, we will investigate valence-specific performance across the participant group for no-brainer trials, and will exclude participants based on statistical deviance from mean performance.

Additionally, we note that the experimental paradigm for participants was extensive, requiring two separate sessions and three separate outcome valences for the ambiguous decision-making task. Our present analysis shows similarities in ambiguity preference across loss avoidance and threat avoidance paradigms, namely categorical ambiguity seeking behavior coupled with ambiguity avoidance which is increased with level of missing information. Previous studies have indeed found indications that individuals may be ambiguity-seeking in loss scenarios (Kocher et al., 2018; Machina et al., 2011). However, previous investigation of a threat-avoidance paradigm found that participants were broadly ambiguity avoidant at both the categorical level and in relationship to level of missing information. This raises the question of whether baseline ambiguity avoidance/seeking in the threat paradigm is typically similar to that in loss avoidance or not. It is possible that the longer current experimental paradigm, in which participants completed reward, loss, and threat tasks, could cause participants to view threat avoidance in a more economic manner given the uniformity of experiment style and stimuli. This could be investigated in more detail in a large data set of threat avoidance ambiguous decision-making without inclusion of any reward-based decision-making tasks.

Finally, we note a modeling limitation in the current study. Based upon previous literature, we considered decision-making models based on expected utility (Kahneman and Tversky, 1979; Tversy and Kahneman, 1974), as well as models based on independent evaluation of magnitude differences and probability differences between choice options (Lawrance et al., 2022). We also tested models which allowed for both mechanisms to simultaneously apply. These more complex models proved to be the best fit to the behavior for each valence paradigm using a metric of fit

(WAIC) which penalizes for model complexity in an attempt to prevent overfitting. However,, we note that several parameters of the winning models (which were not used in results) did not recover as well as other parameters in the winning model (e.g., β_{MDiff} in reward pursuit, where 75% of generative parameter values were recovered; see Supplements for details of parameter recovery). This may indicate that the current data set does not have enough power to accurately identify the larger set of parameters in the fully interactive models. In future analysis, a careful examination of parameter recoverability could be performed for each model class to assess data set power in parameter identification.

Discussion

Individuals vary widely in their tolerance of uncertainty during decision-making, both with respect to first order uncertainty ("risk", as established by a wealth of behavioral economics literature (Kahneman and Tversky, 1979; Tversy and Kahneman, 1974)) and with respect to second order uncertainty such as ambiguity (Kocher et al., 2018; Trautmann and Van De Kuilen, 2015). Classical findings have established a "gain-loss framing" effect in risky decision-making, whereby individuals are broadly risk-avoidant in gain scenarios but tend to become more risk-seeking in loss scenarios (Kahneman and Tversky, 1979; Tversy and Kahneman, 1974). These findings have also been extended to decision-making under ambiguity (Kocher et al., 2018; Trautmann and Van De Kuilen, 2015). Here, most investigations have used experimental paradigms which evaluated presence versus absence of ambiguity but did not investigate a manipulation of ambiguity level. However, work which separately assessed the effect of categorical ambiguity and level of missing information has found that neural activations in response to the presence of ambiguity are differentiable from neural response to level of missing information (Bach et al., 2011), emphasizing the need to examine both facets of decision-making under ambiguity. Additionally, recent work in the threat avoidance domain has revealed that individual ambiguity aversion tends to be highly sensitive to level of missing information (Lawrance et al., 2022), and that individuals with higher levels of trait anxiety (as measured by STAI-T) show significantly increased sensitivity to level of missing information but not to categorical ambiguity (Lawrance et al., 2022).

Here, we sought to expand the investigation of ambiguity aversion across all three previously investigated valence paradigms (gain, loss, and threat avoidance) in a single study, while manipulating both categorical ambiguity and level of missing information, in an in-lab participant sample. In addition, we aimed to investigate the role of dimensions of psychopathology in decision-making mechanisms using latent dimensions of variance extracted from self-report measures across anxiety, depression, and hypomania.

Our first aim was to test the relationship of symptom-level hypomania with increased tolerance of ambiguity in reward gain environments, either at the categorical level or with respect to level of missing information. Individuals experiencing manic symptomatology are often described as "risk-seeking", and indeed one diagnostic symptom of both mania and hypomania is "excessive involvement in actions that have a high potential for painful consequences" (American Psychiatric Association., 2013). Although studies using gain-based gambling tasks have found evidence of increased risk-taking behavior in individuals with current manic state symptoms (Adida et al., 2011, 2008), it is generally thought that these behaviors arise from increased sensitivity of the behavioral activation system (BAS) and alterations to reward processing (Johnson et al., 2012). In

the present study, we sought to investigate whether heightened levels of hypomanic symptomatology might additionally relate to decreased sensitivity to uncertainty in decision-making. We found no association between higher scores on a latent dimension characterized by hypomania and ambiguity sensitivity (either categorical or level of missing information-dependent) or risk sensitivity across reward pursuit, loss avoidance, and threat avoidance paradigms. We confirmed this finding using scores on the HPS, a self-report questionnaire measure designed to assess levels of hypomania. However, we did find a significant dissociation between risk sensitivity in hypomania versus anhedonic depression in threat avoidance, with individuals with higher levels of hypomanic symptomatology less influenced by risk in these decisions than individuals with higher levels of depression. If replicated, these findings could provide potential evidence for increased risk-taking behavior in threat avoidance decisions among individuals with higher levels of subclinical manic symptomatology. However, the evidence indicates that this finding may be valence specific, with risky decision-making in reward pursuit not arising from changes in sensitivity to first-order or second-order uncertainty. We do note that the trait levels of hypomania considered in the present study might fail to capture the full extent of behavioral changes associated with clinically diagnosed mania, so replication of the null finding in reward pursuit in a participant sample involving individuals experiencing manic symptomatology commensurate with manic episodes would provide broader confirmation.

Our second aim was to build on previous research findings that higher levels of trait anxiety are related to greater sensitivity to level of missing information in threat avoidance paradigms of ambiguity (Lawrance et al., 2022) by using a larger data set and a dimensional characterization of anxiety. Here we tested latent dimensional characterizations of physiological anxiety, cognitive anxiety, and shared anxious/depressive variance against parameters of threat avoidance decision-making from the best-fit computational model. We did not observe any significant relationship between dimensional factor scores and mechanisms of threat avoidance decision-making. We additionally tested the relationship of trait anxiety as measured by STAI-T, which was used in previous research (Lawrance et al., 2022), with the same decision-making parameters in threat avoidance. Here we observed a significant relationship between higher scores on STAI-T and a greater weighting of magnitude during decision-making, in contrast to the previous findings of increased sensitivity to level of missing information. However, we note that in the prior study (Lawrance et al., 2022) participants completed only the threat avoidance decision-making paradigm, whereas in the current study each participant completed all three valence paradigms. It is possible that the inclusion in the task of reward pursuit and loss avoidance decision-making caused participants to also treat the threat avoidance paradigm more like an economic decision-making scenario than they otherwise would have (see Limitations section below).

Additionally, differences in findings may arise due to the larger sample size or the enhanced modeling framework in the current study, which allowed for both linear differences in magnitude and probability and expected utility differences to be considered in the decision-making process. Interestingly, factor scores on the dimension characterized by shared variance across anxiety and depression were related at a statistical trend-level basis to a greater weighting of magnitude during threat avoidance decision-making. This might be suggestive of separable influences of anxiety versus comorbidity between anxiety and depression on aspects of threat-avoidance decision-making and provides an interesting direction for future investigation.

Finally, we examined the relationship of dimensional characterizations of anxiety versus depression (as well as shared variance) with mechanisms of reward pursuit and loss avoidance decision-making under ambiguity. Here we found no associations between factor scores and computational reward pursuit decision mechanisms. In loss avoidance, we observed a differential effect on subjective probability weighting between the factor characterized by depressive symptoms and the factor characterized by physiological anxiety. Individuals with higher factor scores in physiological anxiety had significantly lower nonlinear adjustment to probability than those with lower scores, while higher factor scores in anhedonic depression were associated with an increased weighting of probability differences between urns during loss avoidance. This differential sensitivity to probability in ambiguous loss-avoidance between dimensions of depression and anxiety reveals a dissociable focus on probabilistic outcomes in loss domains, in which individuals with depression-related symptomatology more heavily weight probability during decisions while those with physiological anxiety-related symptomatology place less importance on probability in these same decisions.

Chapter 3: Supplemental materials

Figure S.3.1a: Subscale distributions for final in-lab behavioral participants (n=71)

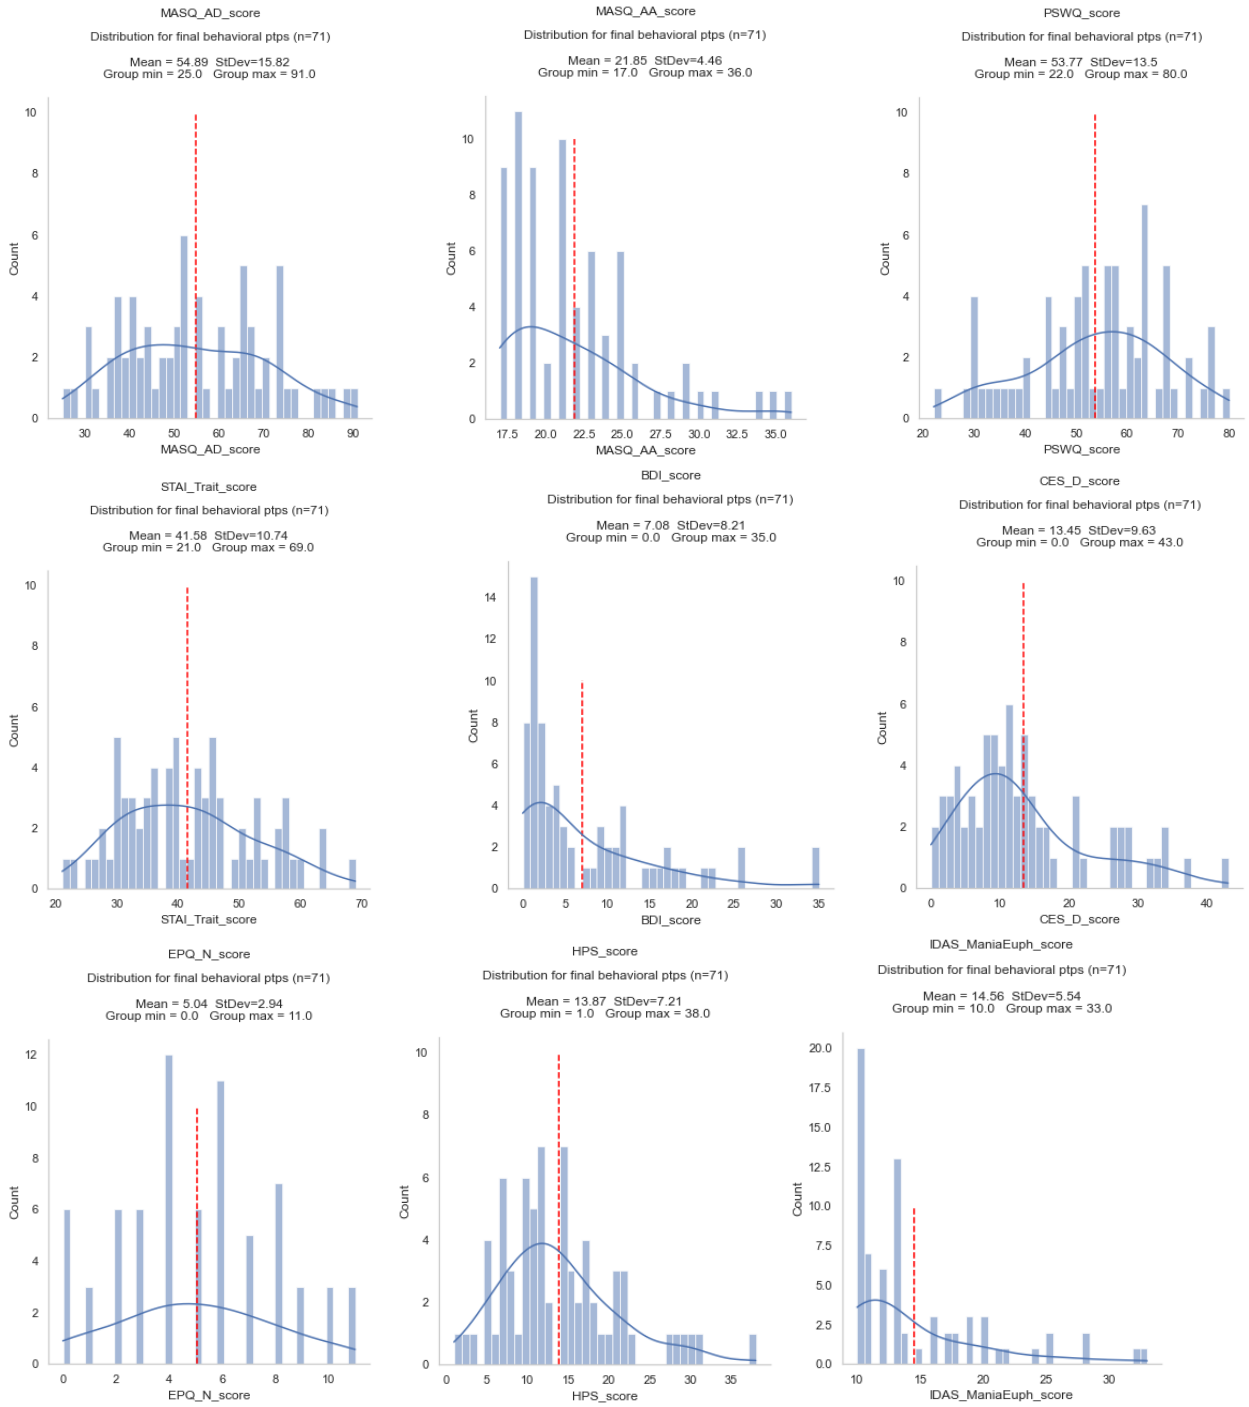


Figure S.3.1b: Within-group subscale score correlations for final in-lab behavioral participants (n=71)

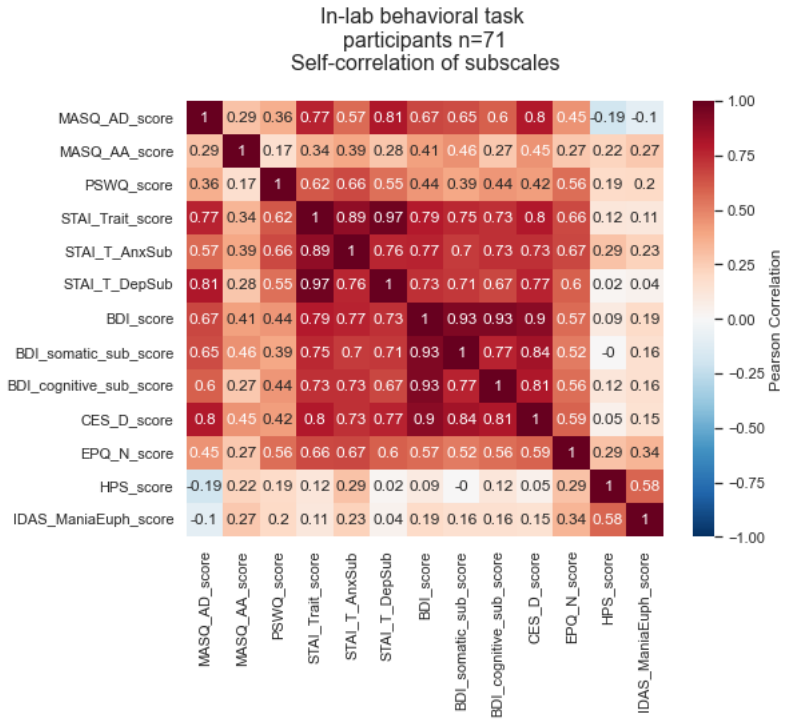


Figure S.3.2a: Subscale distributions for online participants (exploratory bifactor analysis; n=470)

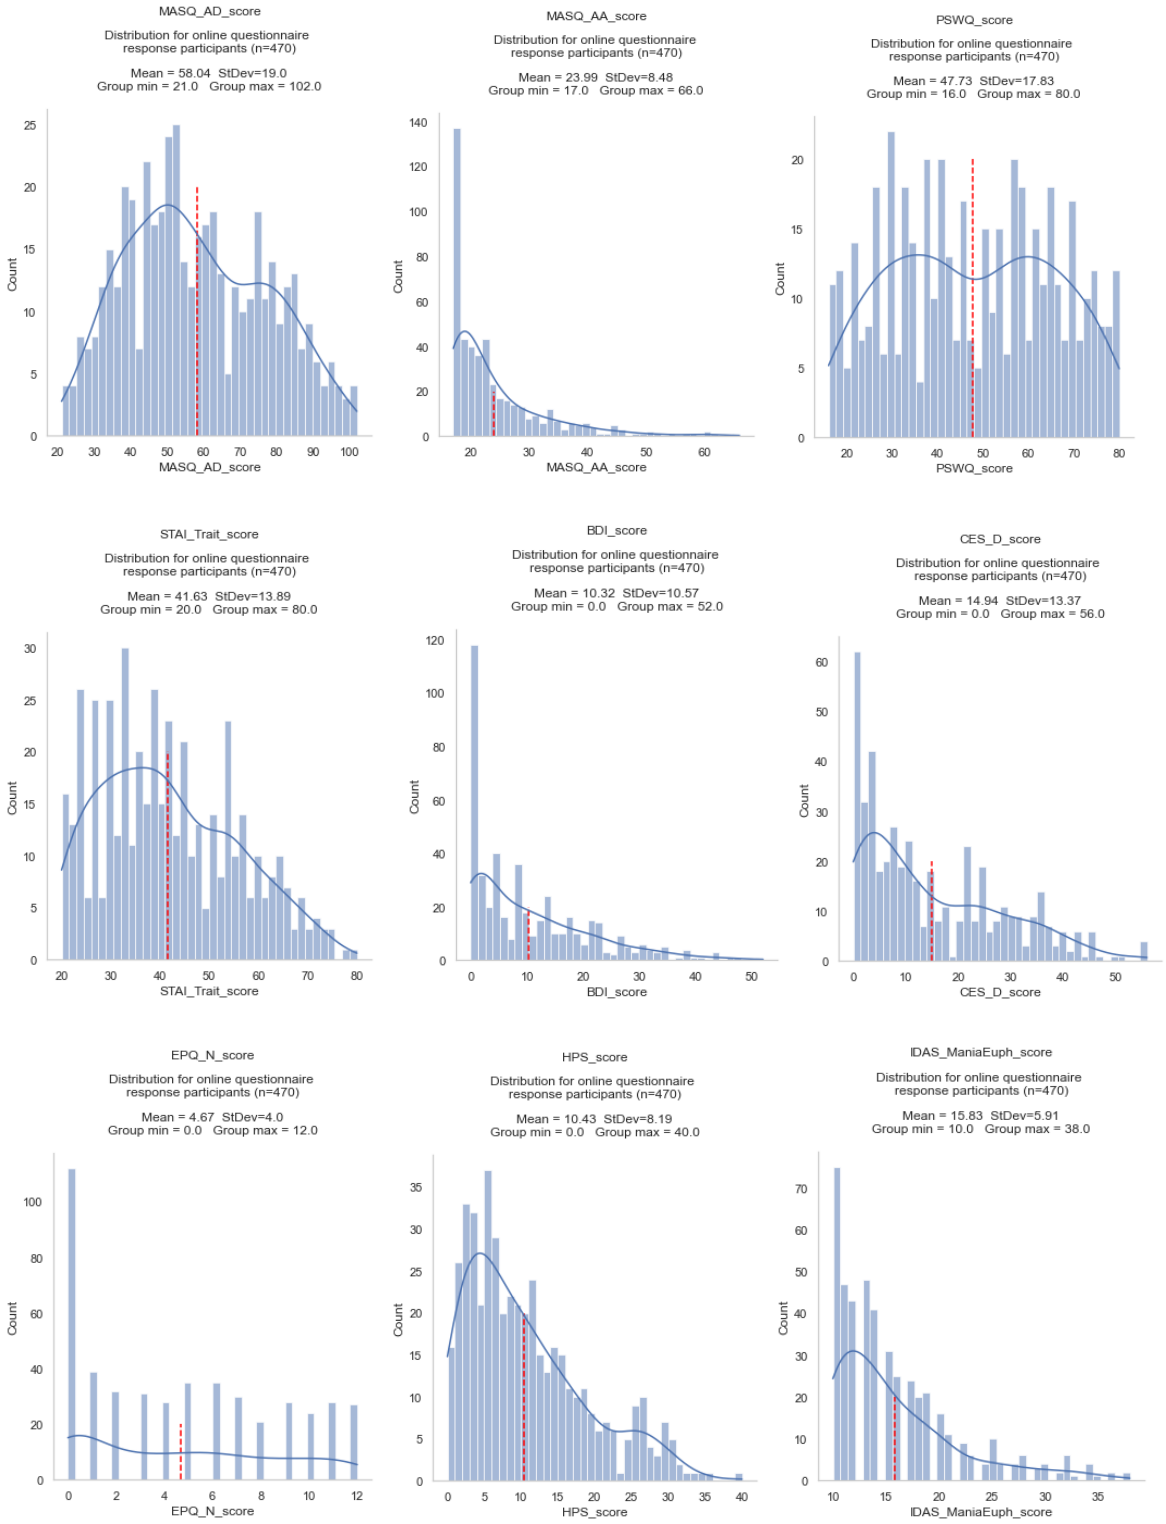


Figure S.3.2b: Within-group subscale score correlations for online participants (exploratory bifactor analysis; n=470)

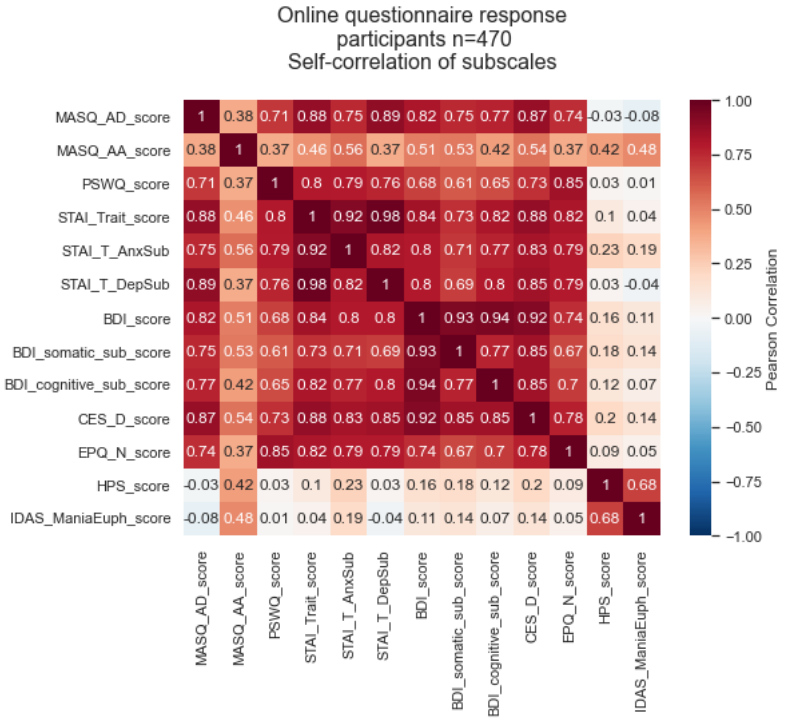


Figure S.3.3a: Subscale distributions for combined in-lab participants (confirmatory bifactor analysis; n=179)

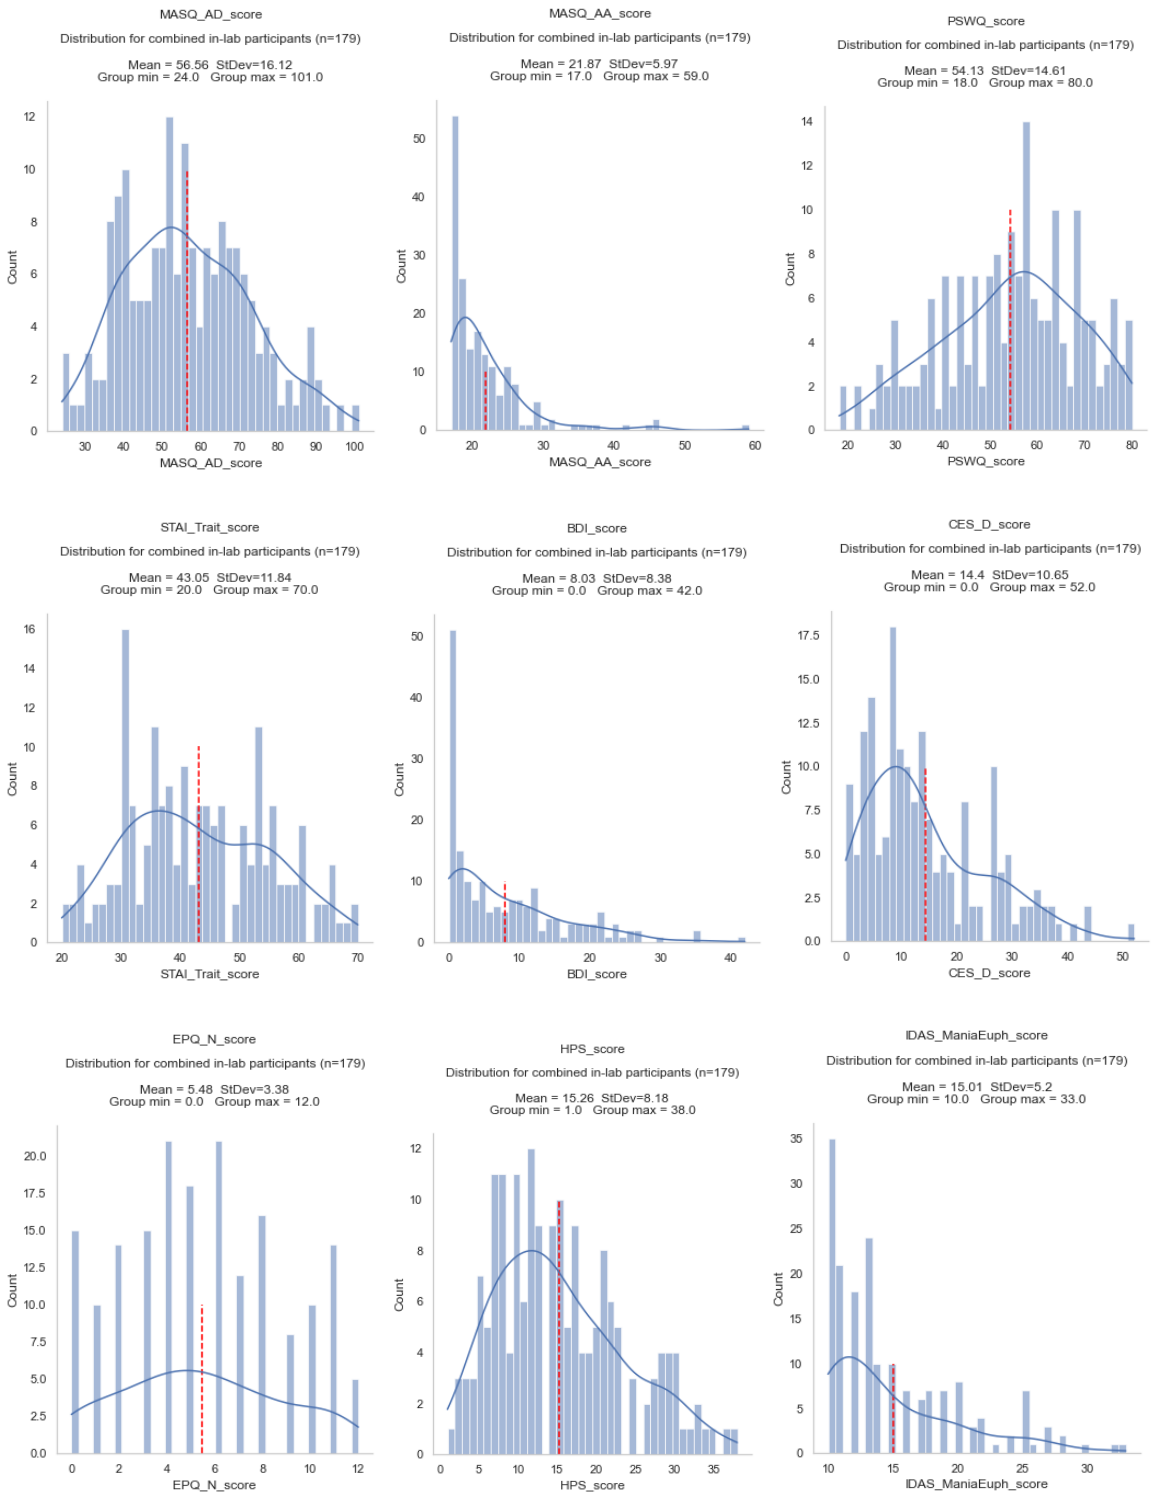


Figure S.3.3b: Within-group subscale score correlations for combined in-lab participants (confirmatory bifactor analysis; n=179)

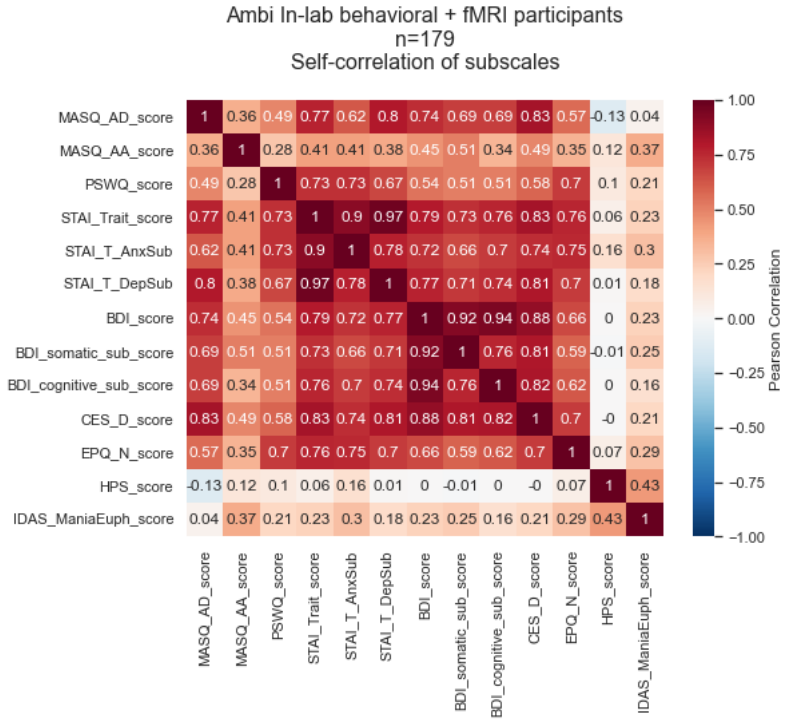
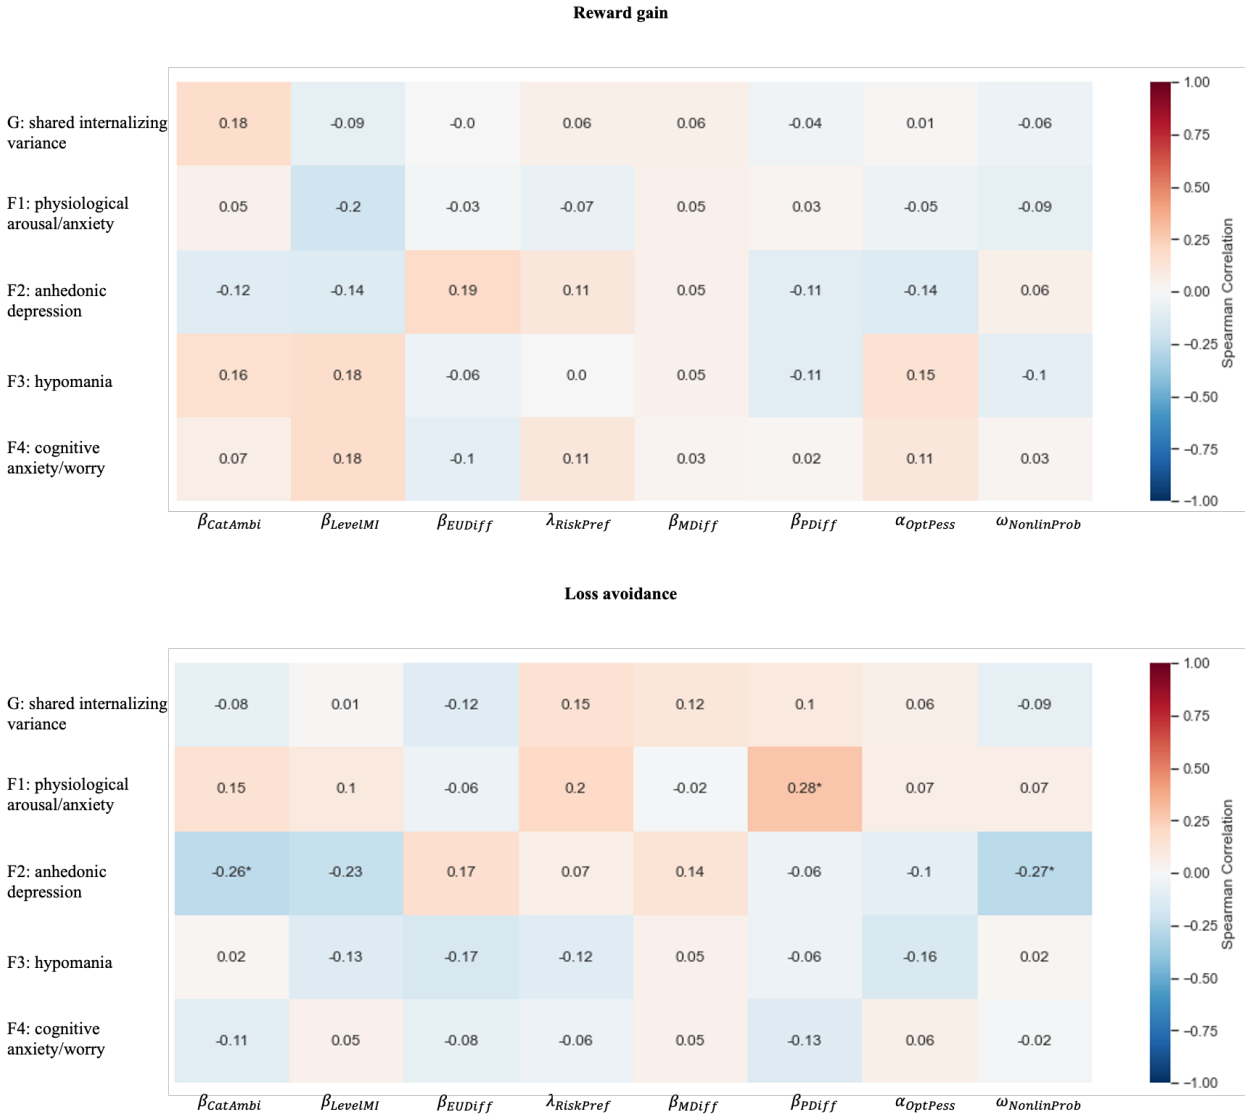


Figure S.3.4: Spearman correlations between posterior mean parameter values and scores on factors of the winning bifactor model.

* denotes p<0.05; ** denotes p<0.01; *** denotes p<0.001.



Threat avoidance

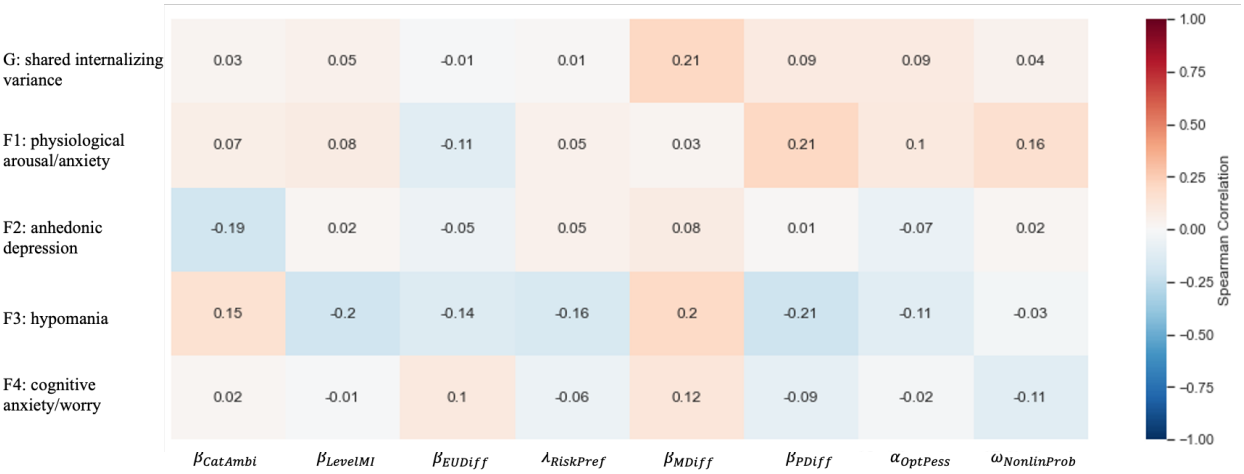
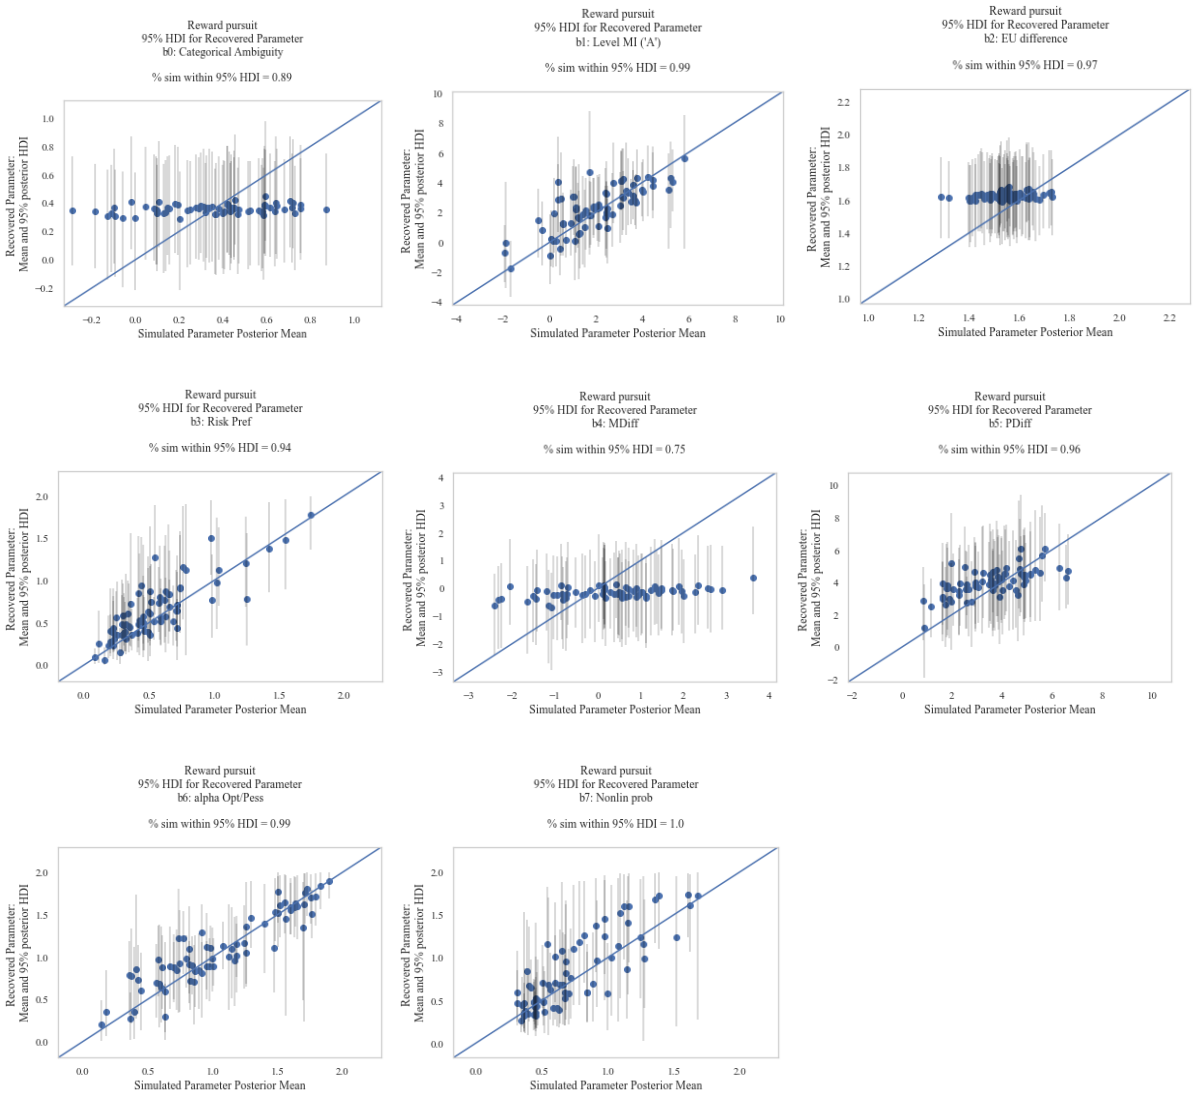
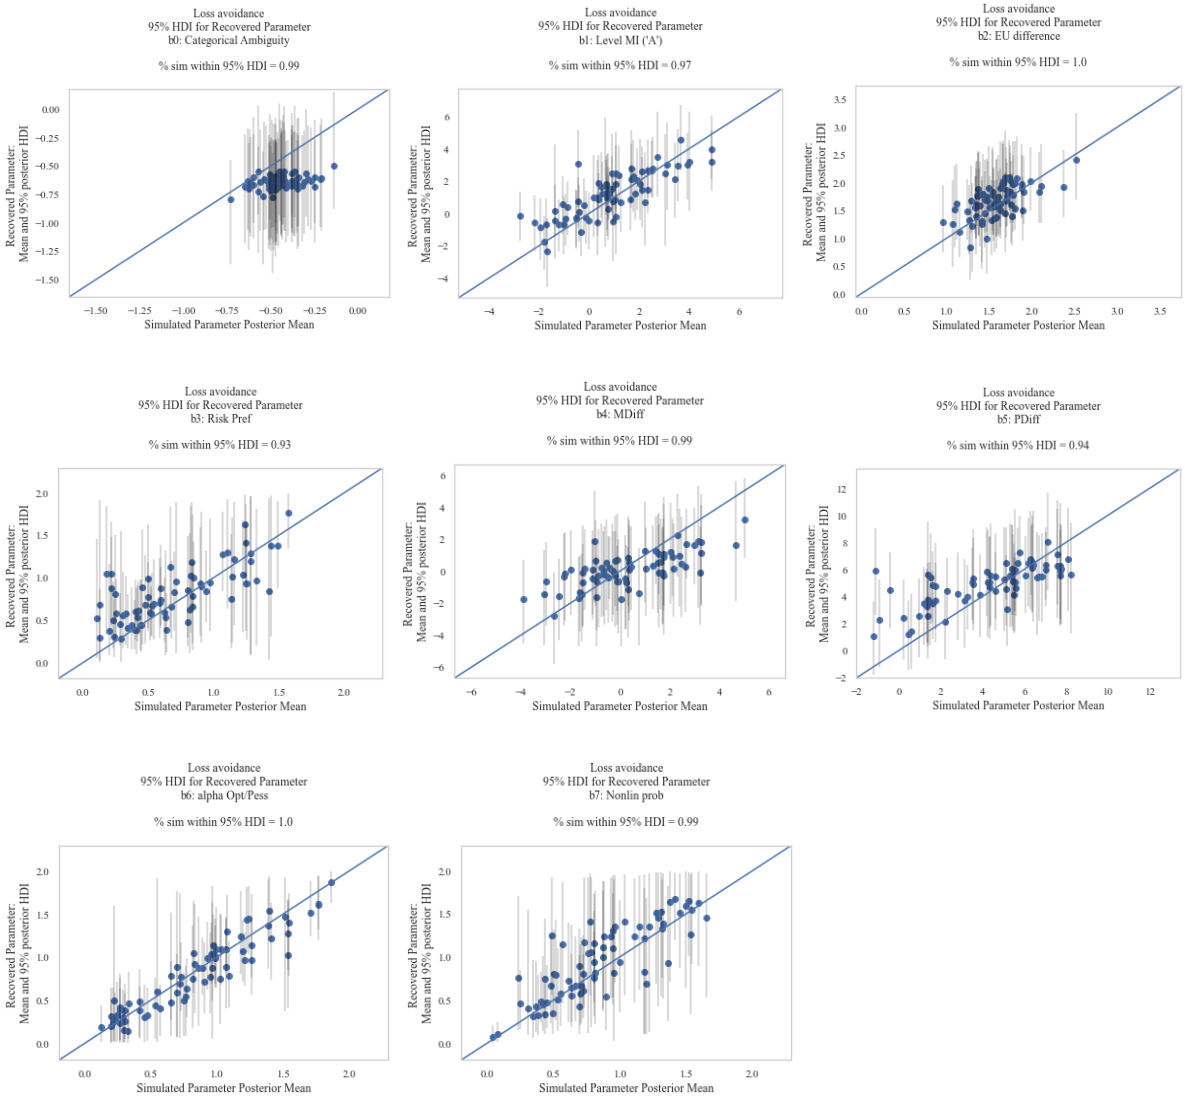


Figure S.3.5: Parameter recovery for winning model (Model FI_EUMP_OP_c_A) in each valence paradigm

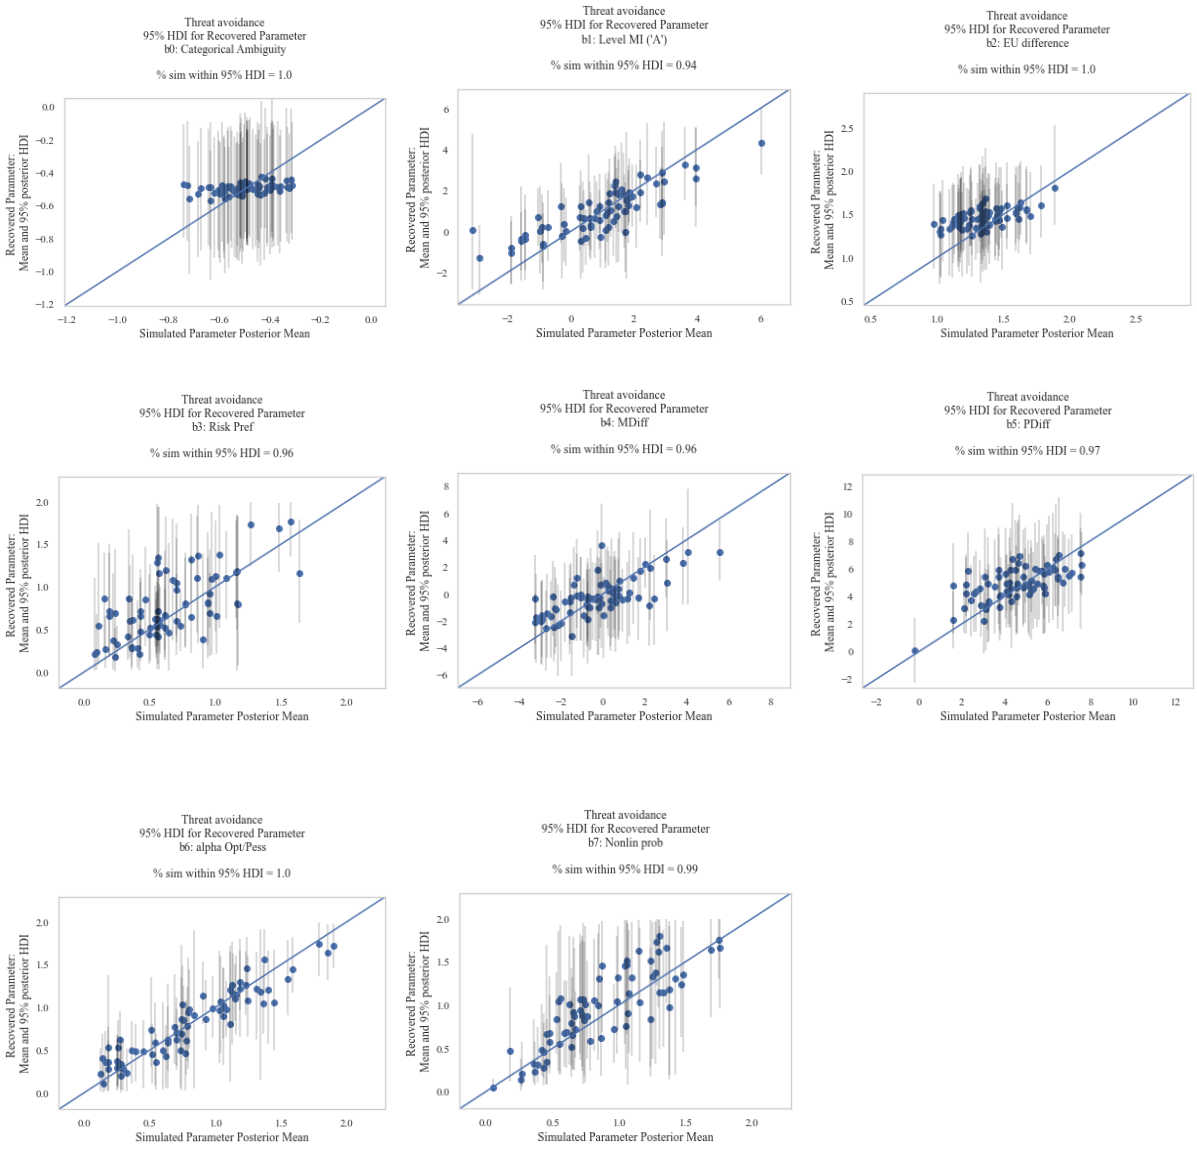
Reward pursuit



Loss avoidance



Threat avoidance



Chapter 4

Reinforcement learning deficits in physiological anxiety are contributed to by both reduced learning rate and increased rate of working memory decay

Abstract

Various forms of anxiety have long been linked to impairments in learning, attentional control, and working memory. However, the influence of working memory on the reinforcement learning process itself has generally not been historically investigated. Recent work in computational cognition has shown that the reinforcement learning and working memory systems are interactively involved during learning, and that modeling the role of both modules during reinforcement learning can provide important insights into the distinction between learning versus working memory system abnormalities. Here we employ an experimental paradigm designed to variationally manipulate the relative loads on working memory and reinforcement learning systems during a learning and retention task. We jointly model the reinforcement learning and working memory systems and their interactions; our findings show that overall impairments in learning accuracy for individuals with higher levels of anxious arousal involve both reduced reinforcement learning rate and faster rate of working memory decay during learning. We do not find any relationship between cognitive anxiety and impaired learning performance, reinforcement learning or working memory mechanisms, indicating that the nature of the anxiety measure used may critically impact the result. These findings may be useful in the consideration of a broad range of reinforcement learning investigations in anxiety, which typically have not considered the relative contributions of working memory to the learning process.

Introduction

Our ability to learn from our experiences of the world is a crucial element of successful decision-making and ultimately survival. Models of reinforcement learning (RL) have been successfully used to investigate cognitive mechanisms of learning in animals and humans (Sutton & Barto, 1998). A key feature of RL models is the use of a parameterized learning rate (often denoted as α) which describes how quickly an individual updates value estimates based upon a prediction error, which is defined as the amount by which experienced outcomes differ from expected outcomes. Findings in both animal electrophysiology and human functional brain imaging have generally established the involvement of the ventral tegmental area (VTA), the striatum, and associated cortico-basal-ganglia circuits in the processes of RL value representation and updating. Some components of these algorithmic representations have shown potential neural instantiation, such as prediction error encoding among the dopaminergic neurons of the VTA (Montague et al., 1996; Schultz, 1997; Bayer & Glimcher, 2005; but also see Dayan & Huys, 2009).

A number of studies have linked various dimensions of psychopathology, including anxiety, with altered reinforcement learning performance or dynamics. Individuals meeting criteria for diagnosis with generalized anxiety disorder (GAD), for example, have been found to show slower learning and worse ultimate performance in an RL task (LaFreniere and Newman, 2019).

A crucial aspect of RL is an individual's ability to modulate learning rate in accordance with the current environment. In particular, various forms of environmental uncertainty, such as an underlying volatility in the association between states or actions/choices and outcome contingencies, require dynamic adjustments to learning by incorporating recent information more quickly into value expectations. Humans are remarkably accurate at appropriately adjusting learning rate in different environments (Behrens et al., 2008, 2007), but several studies have linked higher levels of anxiety to a reduced ability to successfully modulate learning rate between volatile and stable environments, particularly with respect to aversive outcomes (Browning et al., 2015; Gagne et al., 2020). Brain areas including the anterior cingulate cortex (ACC) and amygdala have been associated with the ability to modulate learning rate in volatile versus stable contingency paradigms (Behrens et al., 2007; Bishop and Gagne, 2018a; Li et al., 2011; Roesch et al., 2012).

Intriguingly, although the neural regions involved in RL processes have been reliably linked to value estimation and updating during learning tasks, research has indicated that these regions may not be as strongly involved during the earliest stages of the learning process. Here, functional brain imaging as well as genetic studies have implicated a reduced involvement of cortico-basal-ganglia circuits and heightened involvement of frontal cortical areas in initial learning acquisition and strategic exploration (Collins and Frank, 2012; Frank et al., 2009, 2007; Jocham et al., 2011).

It has been recently suggested that these neural differences in initial versus subsequent learning may be the result of the use of working memory (WM), generally characterized as a short-term, limited-capacity cognitive process involving prefrontal cortical executive function brain areas, in early learning (Collins and Frank, 2012). Working memory serves multiple functions in cognitive processes, including coordinating the ongoing processing of multiple behavioral goals and providing temporary storage and manipulation of limited quantities of information without requiring the full engagement of long term memory processes (D'Esposito and Postle, 2015). While the network of systems underlying working memory is complex, it has been conceptualized as a state-based model which is capable of temporarily accessing distributed information representations via the focusing of attention (Cowan, N., 1995; Lewis-Peacock and Postle, 2008). It has generally been found that elevated activation in PFC is present during (and may coordinate the use of) working memory (Courtney et al., 1997; Funahashi et al., 1989; Fuster and Alexander, 1971).

Studies of WM and attentional control (a closely related process which also involves prefrontal cortical activity) have linked anxiety to deficits in working memory and mitigation of distraction. In a study of N-back task performance in both safe and threatening environments, individuals diagnosed with anxiety disorders had impaired performance and reaction times across both environments (Balderston et al., 2017). In addition, they showed impaired recruitment of dorsolateral prefrontal cortex (dlPFC), an area key in WM activation, and impaired deactivation of neural areas related to the default mode network during the WM tasks. Studies in these areas have also hinted at a dimensional specificity in the role of anxiety in impaired attentional control.

A study using fMRI to investigate performance and cortical activity in a Go/No Go task found that subclinical levels of trait anxiety, but not worry, were linked to impoverished prefrontal control of attention and relatedly reduced frontal-thalamo-striatal connectivity (Forster et al., 2015). In the same study, both worry and trait anxiety showed a higher prevalence of off-task thoughts when increased numbers of task errors were made, indicating that separate investigation of the dimensions of anxiety is crucial for understanding differences in attention and WM.

Although learning and working memory have often been modeled and studied separately, recent evidence has shown that the systems do not operate fully independently, and that in particular reinforcement learning processes are likely to be intertwined with and supported by working memory systems (Collins, 2018; Collins and Frank, 2012). A number of recent studies have used a new reinforcement learning experimental paradigm (Collins and Frank 2012) in which set size of stimuli is varied across the task in order to differentially engage working memory and reinforcement learning system contributions to learning. These studies have shown that models of reinforcement learning must parametrically account for the contributions of working memory in order to accurately reproduce human behavior (Collins, 2018; Collins and Frank, 2012).

Given the established relationship of anxiety with both reinforcement learning and working memory, we sought to apply this enhanced experimental paradigm to determine whether the effects of working memory impairment in anxiety might underlie apparent deficits in reinforcement learning in anxiety. We used computational models of the joint reinforcement learning and working memory process to characterize the mechanisms of interest. Given prior findings indicating that physiological and cognitive dimensions of anxiety may be differentially associated with impairments in processes spanning attentional control, working memory, and reinforcement learning, we sought to test our hypotheses with respect to specific characterizations of each symptom domain. We used scores on self-report measures of anxious arousal (Mood and Anxiety Symptom Questionnaire Anxious Arousal subscale) and worry (Penn State Worry Questionnaire) to evaluate the influence of these two distinct domains of anxiety on learning rate and key parameters of working memory involvement during reinforcement learning.

Methods

Participants

The study was conducted online with participants recruited via the UC Berkeley Research Participation Pool (RPP), which offers partial course credit to undergraduate students for participation in human subjects research. All participants completed online informed consent prior to participation. The study was approved by the UC Berkeley Committee for the Protection of Human Subjects (CPHS). An initial total of n = 229 students [143 = female, mean age = 21.2 +/- 2.4] participated prior to exclusions.

Exclusions

Participants were eligible to participate in the online study if they confirmed that they were not currently taking antidepressant or anxiolytic medications and had not used cannabis within the preceding two weeks. Two attention checks were embedded in the self-report questionnaires (e.g., "Select 2 here to show that you are paying attention") as data validity checks. Additionally, following Meade & Craig, 2012, once participants had completed the task they were informed that their participation credit was now guaranteed, and were asked to answer two questions honestly to ensure that the research would only use credible data: participants were asked to again confirm whether they had used cannabis in the last two weeks, and were asked whether they felt they had paid sufficient attention during the task that their data should be used in our study. 45 participants were excluded based on responses to the end-of-task questions, and 6 additional participants were excluded for missing two or more attention checks during the questionnaires.

Participants were also excluded if they had greater than two standard deviations above the mean number of trial timeouts (>11.9% of trials; 6 participants excluded). Additionally, we calculated mean set size = 2 asymptotic performance across the last 5 same-stimulus presentations across set size = 2 blocks; participants with mean accuracy lower than two standard deviations below the group mean (<64% accuracy; 8 participants) were excluded from further analysis.

The number of participants included in the final analysis was n=164 (104 = female; mean age 21.20 +/- 2.43). 34 participants self-reported Hispanic or Latino ethnicity. Self-identified race distribution of participants was as follows: 94 Asian, 8 Black or African American, 34 White, 2 Native American or Pacific Islander, 12 who identified as "More than one race" and 14 "Unknown".

On a within-participant level, trials that timed out were excluded (mean number of timeout trials = 4.74), and subsequently any trial blocks for which a participant had fewer than 9 presentations of any stimulus remaining in a block after exclusions were also excluded (1 block excluded for each of 5 participants).

Experimental paradigm

The behavioral task used was a variant of the classic RLWM task (Collins & Frank 2012). The main task comprised a learning (or "training") phase, followed by an unrelated distractor task, and then finally a surprise testing phase of the original task stimuli. During the learning phase, participants were presented with a stimulus image on screen for up to 1.5s each trial. There were three possible correct key presses associated with each stimulus; keyboard keys 'j', 'k', or 'l'. Each stimulus within a block was associated with only one correct key press (but multiple stimuli within a block could be associated with the same correct key press). Participants had 1.5s to select an action for each stimulus presented; if no response was selected in the allowed time, a message "You did not make a selection in time!" was displayed for 500ms, and the task advanced to the next trial. If participants responded with a key press, the stimulus was removed from the screen and they were given accurate feedback (a green +1 for correct responses, or a red 0 for incorrect responses) presented for 500ms. A fixation cross was shown for 500ms between each learning trial with a 100ms blank screen immediately before and after the fixation cross.

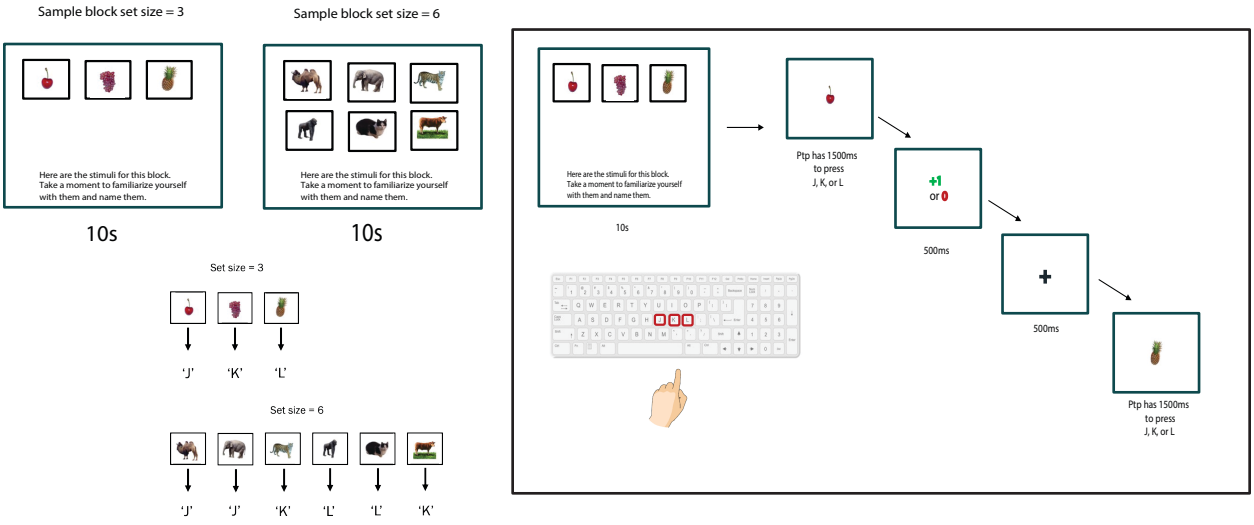


Figure 4.1 RLWM task Participants completed 13 blocks of trials. Each block had a set size of 2, 3, 4, 5, or 6 images to be learned, and each stimulus was shown 13 times within each block for up to 1.5s each presentation (terminating when a key was pressed). Images for a block were related to each other (e.g., "Instruments") and were unique for each block. Each image in a block was associated with one correct response key from 3 possible keys ('j','k','l'). Participants learned the correct key to press for each image via trial and error, with immediate feedback displayed for 500ms after each key press (+1 if correct, 0 if incorrect).

The task consisted of 13 blocks of trials. Within each block, participants learned the correct key press action for each stimulus in the block through multiple trials and feedback. Each trial block included a stimulus set size (nS) varying between 2, 3, 4, 5, or 6 stimuli. Stimuli for each block were drawn from a randomly selected image category (such as 'nature', 'shapes', 'musical instruments', etc.) without replacement, so that each set of unique stimuli within a block were of the same category and no stimuli or categories were repeated across blocks.

Participants were randomly assigned to one of 10 pre-generated learning task trial sequences. In each generated trial sequence, participants started and ended with a block of set size = 2 in order to mitigate the possible conflation of primacy and recency effects with working memory. Block 2 was always of set size = 3; the remaining intermediate blocks were shuffled so that each set size was presented once in blocks 3-7 and once in blocks 8-12. Each stimulus for a given block was shown 13 times within that block, so the total number of trials per block varied with stimulus set size. Stimuli sequences within each block were generated pseudo-randomly controlling for a uniform distribution of delay between each two successive presentations of the same stimulus across every [1:2*nS] trials per block.

Following the learning phase, participants were given an optional break of up to one minute, followed by a distractor task. In the current paradigm participants completed a short n-back task as the distractor task, where they were shown a series of images (shapes which were not used in the learning task) and asked to press the left arrow key of the keyboard if the image on screen was the same as the image that had appeared "n" stimuli ago. Participants completed a 1-back, 2-back, and 3-back task (mean completion time = 18.8 minutes). Following the n-back task, participants were given an optional 30 second break.

Following this break, participants were informed that they would be tested on their learning of the actions associated with the stimuli they had seen during the earlier learning portion of the experiment. Participants were shown each stimulus from their learning phase a total of three times during the testing phase. In each testing trial, the stimulus appeared on screen for 1.5s, and the participant selected 'j', 'k', or 'l' based on their earlier learning for the stimulus. Importantly, participants were not given feedback on accuracy in the testing phase, so additional learning could not occur. If participants did not select an action quickly enough, a message "You did not make a selection in time!" was displayed for 500ms. A fixation cross was shown for 500ms between each testing trial with a 100ms blank screen immediately before and after the fixation cross.

Self-report measures of trait anxiety and depression

Prior to the behavioral task and following informed consent, participants were asked to complete a short series of questionnaires designed to measure aspects of anxiety and depression. The questionnaires used in this study comprised the following: the State-Trait Anxiety Inventory trait subscale (STAI-T; Spielberger, C. D., Gorsuch, R. L., Lushene, R., Vagg, P. R., & Jacobs, G. A., 1983); the Penn State Worry Questionnaire (PSWQ; Meyer et al., 1990); the Beck Depression Inventory (BDI-II; Beck, 1961)); and the Mood and Anxiety Symptom Questionnaire anxious arousal subscale (MASQ AA; Watson & Clark, 1991). Items addressing suicidality were excluded due to the online nature of the study. Final number of items n=126.

Scores on the MASQ AA subscale had a mean of 27.25 (standard deviation=10.86), with a minimum participant score of 17 (minimum possible scale score=17) and a maximum participant score of 68 (maximum possible scale score=85). Scores on the worry measure PSWQ had a mean of 52.12 (standard deviation =13.81), with a minimum participant score of 16 (minimum possible scale score = 16) and a maximum participant score of 79 (maximum possible scale score = 80). These subscale scores differentiate between anxiety as characterized by somatic symptoms versus cognitive symptoms.

Complete score distributions for the participant group, including for questionnaire measures used in exploratory post-hoc analyses, are shown in Supplemental Figure 4.S.1.

Behavioral analysis

Previous investigations of learning under various set sizes have found a robust and highly replicable effect of set size on learning over time. To determine if we replicated these past research findings (Collins, 2018; Collins and Frank, 2012; Master et al., 2019), we first examined the mean learning curves across participants at each set size, measured as mean response percent accuracy by stimulus presentation number. We calculated overall learning accuracy as the mean accuracy across all presentations of all stimuli for each set size, and asymptotic learning accuracy as the mean of the last 5 presentations of each stimulus across blocks with the same set size. Overall testing accuracy was calculated as the mean accuracy across all stimuli presentations (each image was shown 3 times during testing) for each set size.

Where underlying data followed a normal distribution, we used Pearson correlations and t-tests as specified throughout the text. Statistical tests were performed in Python, R, and SPSS. We used

Spearman rank correlation, Wilcoxon, Kruskal-Wallis and Mann-Whitney U testing as specified throughout the text to test various correlations and distributional differences.

Computational modeling

We modeled learning and test performance using a series of two-module reinforcement learning and working memory (RLWM) models of varying levels of complexity. These models have been shown to effectively tease apart the effects of working memory from those of reinforcement learning in RLWM tasks with a set size manipulation as described above (Collins, 2018; Collins et al., 2014; Collins and Frank, 2012; Master et al., 2019).

Baseline model: Reinforcement learning and working memory (RLWM) with cooperative interaction

We first tested a classic version of a two-module reinforcement learning and working memory model which allows working memory weights to influence choice selection and to contribute to RL value updating.

Working memory weights are updated assuming perfect information retention after each trial, but these weights are assumed to decay rapidly at a parameterized rate of φ between each value update to reflect the short-term nature of WM.

The WM weights for possible actions (a) given each stimulus (s) (denoted as $W(s,a)$) are initialized to random chance ($1/nA$, where $nA=3$ is the number of possible actions) for each possible state and action per block. The weight for the current action/stimulus (s_t, a_t) for each trial is then updated based a prediction error from observed trial feedback r_t (1 = correct, 0 = incorrect) with complete retention (learning rate $\alpha_{WM} = 1$) such that:

$$W_{t+1}(s_t, a_t) = W_t(s_t, a_t) + \alpha_{WM} \cdot \delta_t \quad (1)$$

$$\delta_t = r_t - W_t(s_t, a_t) \quad (2)$$

WM weights for all stimuli and actions (s,a) decay back to random chance after each trial according to:

$$W_t(s,a) = W_t(s,a) + \varphi \cdot ((1/nA) - W_t(s,a)) \quad (3)$$

The contribution of working memory to action selection in each block is weighted according to a prior WM confidence weight parameter ρ and the capacity parameter K of WM compared to the number of stimuli in the set size (ns) for the current block.

$$w_{WM_{block}} = \rho \cdot \min(1, K/ns) \quad (4)$$

The reinforcement learning (RL) module tracks action values (denoted $Q_{RL}(s,a)$), also initialized to $1/nA$. Values are updated for the current action/stimulus at each trial using a cooperative approach whereby both RL values and WM weights contribute, with a reward prediction error (*rpe*)

defined as follows and parameter *i* controlling the strength of the interaction between RL and WM modules:

$$rpe_t = r_t - [(w_{WM_{block}} \cdot i \cdot W(s,a)_t) + ((1 - w_{WM_{block}} \cdot i) \cdot Q_{RL}(s,a)_t)] \quad 5$$

Values for RL are then updated using a learning rate parameter α_{RL} governing the rate at which feedback is incorporated into the estimate.

$$Q_{RL}(s,a)_{t+1} = Q_{RL}(s,a)_t + \alpha_{RL} \cdot rpe_t \quad 6$$

To model the probability that an agent selects a given action at a given time, we use a weighted mixture softmax choice policy to evaluate the relative value of each action during the decision process. The probability of selecting action *a* following stimulus *s* is denoted as follows:

$$p(a|s) = (1 - w_{WM}) \cdot \frac{\exp(\beta \cdot Q_{RL}(s,a))}{\sum_i \exp(\beta \cdot Q_{RL}(s,a_i))} + w_{WM} \cdot \frac{\exp(\beta \cdot W(s,a))}{\sum_i \exp(\beta \cdot W(s,a_i))} \quad 7$$

where the parameter β, referred to as the softmax inverse temperature, controls the extent to which relative action values (as opposed to stochastic choice noise) are used in decision choice. During learning, the softmax inverse temperature was set to a fixed value of $\beta = 50$ to improve parameter reliability (Collins, 2018) and choice noise was captured via an undirected noise parameter ε, such that:

$$p(a|s)_{noisy} = {}^{\varepsilon}/_{nA} + (1 - \varepsilon) \cdot p(a|s) \quad 8$$

Modeling of choice selection during the testing phase used the final learning phase Q_{RL} values for the test stimuli, and assumed an epsilon-noisy softmax action selection with no contribution of working memory to choice. The parameter β_{test} was fit in the testing phase softmax action selection policy (reflecting potential weakening of RL-values' influence on test), while ε was fit jointly between the learning and testing phases (reflecting a shared tendency for lapses or inattention).

Model variant #1: RLWM with negative feedback neglect ("bias")

We next tested a model which incorporated the effect of negative feedback neglect on learning and working memory. Previous studies (Collins, 2018) have found that individuals often update values at a lesser rate following negative feedback than positive feedback during learning. Here, we parameterize this by reducing the learning rate in both RL and WM modules following negative feedback according to the rule:

$$\left.\begin{aligned} Q_{RL}(s,a)_{t+1} &= Q_{RL}(s,a)_t + \alpha_{RL} \cdot rpe_t \\ W(s,a)_{t+1} &= W(s,a)_t + 1 \cdot \delta_t \end{aligned}\right\} \quad for\ r = 1 \quad 9(a), 9(b)$$

$$\left.\begin{aligned} Q_{RL}(s,a)_{t+1} &= Q_{RL}(s,a)_t + (1-\eta)\cdot\alpha_{RL}\cdot rpe_t \\ W(s,a)_{t+1} &= W(s,a)_t + (1-\eta)\cdot\delta_t \end{aligned}\right\} \quad for\ r=0 \qquad 10(a),\ 10(b)$$

Model variant #2: Asymmetric negative feedback neglect (total bias for RL and parametric bias for WM, "*asymbias*")

As mentioned above, literature in reinforcement learning has generally found that individuals do not learn from negative feedback as much as from positive feedback. Here we tested the hypothesis that individuals entirely neglect negative feedback in the RL module, and may be biased to neglect it at a parameterized rate in the WM module.

$$\left.\begin{aligned} Q_{RL}(s,a)_{t+1} &= Q_{RL}(s,a)_t + \alpha_{RL}\cdot rpe_t \\ W(s,a)_{t+1} &= W(s,a)_t + 1\cdot\delta_t \end{aligned}\right\} \quad for\ r=1 \qquad 11(a),\ 11(b)$$

$$\left.\begin{aligned} Q_{RL}(s,a)_{t+1} &= Q_{RL}(s,a)_t \\ W(s,a)_{t+1} &= W(s,a)_t + (1-\eta)\cdot\delta_t \end{aligned}\right\} \quad for\ r=0 \qquad 12(a),\ 12(b)$$

Model variant #3: RLWM with two WM prior confidence parameters (*2rhoK*)

We next tested a model which expanded the WM prior confidence parameter ρ into two possible parameters, ρ_{above_K} and ρ_{below_K}. This modification allowed the weight placed on working memory during choice selection to vary based on whether the current block set size exceeded an individual's working memory capacity.

$$w_{WM_{block}} = \rho_{below_K}\cdot\min\left(1,\frac{K}{ns}\right) = \rho_{below_K} \quad for\ ns < K \qquad 13(a)$$

$$w_{WM_{block}} = \rho_{above_K}\cdot\min\left(1,\frac{K}{ns}\right) = \rho_{above_K}\cdot\frac{K}{ns} \quad for\ ns > K \qquad 13(b)$$

Model fitting procedure

Model parameters were fit using maximum likelihood estimation in Matlab with the function fmincon to minimize negative log likelihood. Individual maximum likelihood estimates were performed using 20 independent randomly selected starting points to improve selection of global maxima. All parameters were constrained to [0,1], except the softmax inverse temperature for the testing phase, which was scaled to be constrained on the interval [0,100], and the working memory capacity parameter K, which was constrained to the continuous interval [2,6].

Model comparison

We evaluated the fit of variants of the RLWM model with combinations of each modification described above. The final model space included 5 models as shown in Table 4.1.

Model fit was assessed by comparing the Akaike information criterion (AIC; Akaike, 1974) for each model in aggregate. The AIC measures overall model fit while penalizing for complexity, and we verify that it supports adequate model identification within the RLWM framework (Master et al., 2019). The winning model with the lowest total AIC was Model #5 (RLWM*_asymbias_2rhoK),* with the following 9 parameters: RL learning rate, test phase softmax inverse temperature, epsilon undirected choice noise, WM confidence for low set sizes, WM confidence for high set sizes, WM forgetting, WM neglect of negative feedback, WM capacity, and an interaction parameter for cooperative updating between RL and WM systems.

There were substantial individual differences in best fit model, with Model #4 best-fitting a greater percentage of participants based on individual AIC (see Figure 4.2). Model #4 was nested within the model with the overall lowest AIC (Model #5), indicating that the additional parameters in Model #5 were not necessary to fit some participants but were necessary to best fit others. Of note is that, because some models are nested, the more complex model could yield the exact same maximum likelihood as the simpler model for individuals who did not require the additional parameters by finding null-effect parameter values as optima for those individual parameters, so any improvement in AIC between Models #4 and #5 for a given individual would be due only to the penalty for modeling more parameters than necessary for that individual in the more complex model.

We compared individual differences in model best fit for high versus low anxiety participants using a median split on each of PSWQ and MASQ AA score, which illustrated that individual differences across model fit were qualitatively related to scores on metrics of anxiety (see Figure 4.2). Given that Model #5 showed the overall lowest AIC, best fit a subset of individuals, and contained Model #4 as a nested component, we proceeded with Model #5 as the winning model for purposes of hypothesis testing. We performed post-hoc confirmatory analysis of the significant findings in the analogous parameters in Model #4 to verify that results were robust across the two models.

Model	Notation	Parameters	Description	Total AIC
Model #1: RLWM	RLWM	$\alpha, \beta_{\text{test}}, \varepsilon, \rho, \varphi, \kappa, i$	Basic RL and WM model with parameter i to coordinate interaction between modules	146,311
Model #2: RLWM_*bias*	RLWM η	$\alpha, \beta_{\text{test}}, \varepsilon, \rho, \varphi, \kappa, \eta, i$	Same as RLWM, but updates following negative feedback are partially neglected according to parameter η	143,910
Model #3: RLWM_*bias_2r*	RLWM η $2\rho K$	$\alpha, \beta_{\text{test}}, \varepsilon, \rho_{low_ss}, \rho_{high_ss}, \varphi, \kappa, \eta, i$	Same as RLWM_*bias*, but allows WM choice weight to vary between set sizes below versus above capacity	143,799
Model #4: RLWM_*asymbias*	RLWM η_{asym}	$\alpha, \beta_{\text{test}}, \varepsilon, \rho, \varphi, \kappa, \eta_{asym}, i$	Same as RLWM_*bias*, but negative feedback is neglected asymmetrically for RL (totally neglected) versus WM (partially neglected)	143,718
Model #5: RLWM_*asymbias_2r*	RLWM η_{asym} $2\rho K$	$\alpha, \beta_{\text{test}}, \varepsilon, \rho_{low_ss}, \rho_{high_ss}, \varphi, \kappa, \eta_{asym}, i$	Same as RLWM_*asymbias*, but allows WM choice weight to vary between set sizes below versus above capacity	143,629

Table 4.1 Computational models of reinforcement learning evaluated. For each model: model key, list of parameters, short description, and total group AIC are shown.

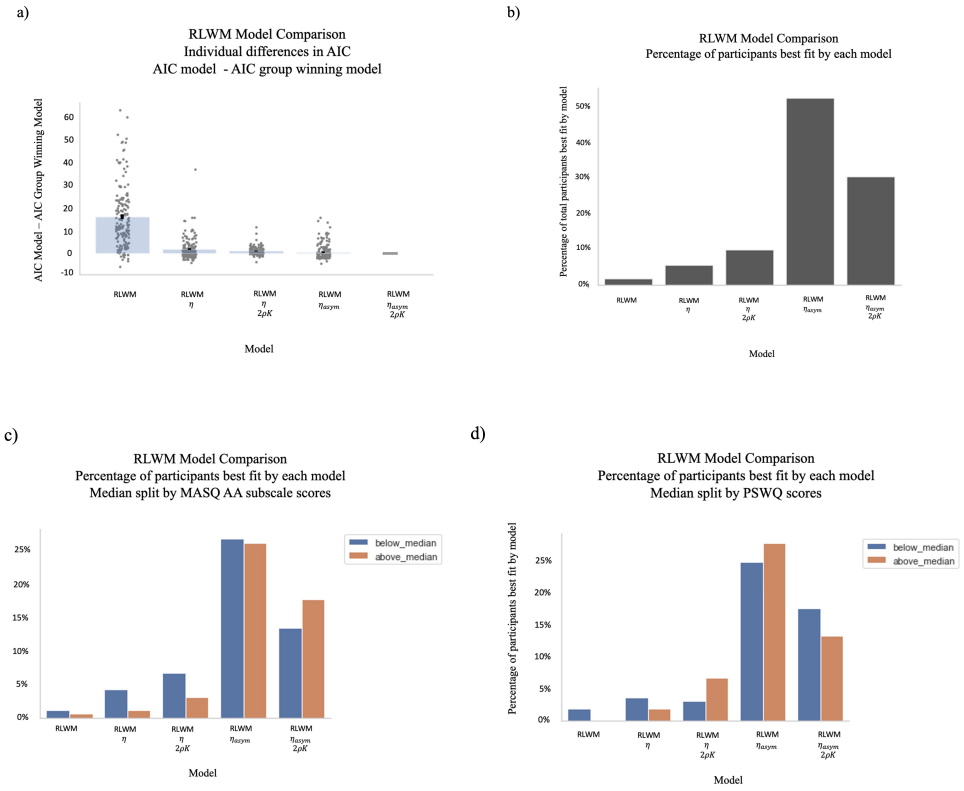


Figure 4.2 Model fit comparison using Akaike Information Criterion (AIC) a) Distribution of individual differences between model fit (AIC) for each model relative to individual AIC for the group winning model (Model #5). b) Percentage of overall participants best fit by each model. c) Percentage of participants best fit by each model by median split on MASQ AA subscale scores. d) Percentage of participants best fit by each model by median split on PSWQ scores.

Model validation

Model validation was performed by simulating data from a generative version of the winning computational model (Model #5, above). Learning curves by set size, difference in performance between low and high set size during learning, and difference in performance between low and high set size during testing were replicated by the model. (See Figure 4.4 for simulated versus actual data).

Parameter recovery was performed for the winning model to test the identifiability of all model parameters. Participant data was simulated using the fit parameter values from the final model, and this simulated data with known underlying parameters was then fit using the optimization process described above. Recovered (fit) parameter values were then compared with the generative (known) parameter values. Parameters generally recovered well, with both significant parameters (learning rate and working memory decay) having correlations between generative and fit parameters greater than or equal to 0.80. See Supplements Figure S.4.3 for parameter recovery analysis for each model parameter.

Model identifiability was confirmed via a model-recovery analysis. Three sets of participant data were simulated from the following models: RLWM_bias, RLWM_asymbias, and RLWM_asymbias_2rhoK. Each simulated data set was then fit by each of the three selected models, and the resulting best-fit model was compared to the generative model for that data. Models recovered well; see Supplements Figure S.4.4 for model recovery analysis.

Results

Accuracy

During learning, all participants performed better than chance (chance $\cong 33\%$), with mean accuracy across all set sizes of 78.2%. As set size increased, accuracy decreased (see Figure 4.3). To quantify the influence of increases in set size on participant performance, we calculated a set size slope for each participant following (Master et al., 2019) according to the following equation:

$$-1.5 \cdot perf_{setsize6} - 0.5 \cdot perf_{setsize5} + 0.5 \cdot perf_{setsize3} + 1.5 \cdot perf_{setsize2} \qquad 14$$

Overall set size slope median was 0.204, and was significantly greater than 0 (Wilcoxon one-sample test statistic = 571.0, p = 0.000) for the group, indicating a significant effect of set size in (decreasing) performance accuracy.

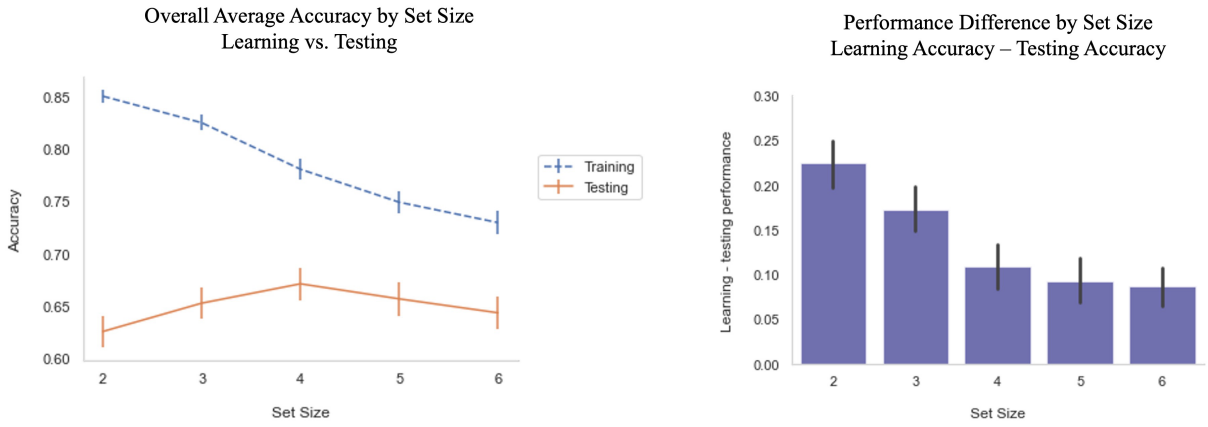


Figure 4.3 Differences between learning and testing performance by set size. Drop between learning and testing performance was significantly higher for lower set sizes than higher set sizes.

During the testing phase, participant performance was substantially lower than during learning (Wilcoxon one-sided test statistic = 13003.0, p = 6.42e-25). Overall performance during testing still exceeded chance (chance ≅ 33%) with a mean testing accuracy of 65.2% across participants. Accuracy during testing was not significantly related to increases in set size (set size slope median = -0.017, Wilcoxon one-sided test statistic=6176.000, p=0.334), replicating previous findings that relative retention of information between learning and testing was greater for higher set sizes and that reinforcement learning and working memory had been differentially engaged across set sizes (Collins, 2018; Rac-Lubashevsky et al., 2023; Rmus et al., 2023; Zou et al., 2022). In line with this, the decreasing difference in performance between learning and testing was significantly related to increases in set size (Kruskal-Wallis = 80.075, p = 1.65e-16; see Figure 4.3).

Learning and testing

At a group level, we examined mean learning curves at each set size, measured as mean response percent accuracy by presentation number averaging across stimuli shown within each set size. In line with previous literature, we found qualitatively different learning curves for each set size, with smaller set sizes associated with a more rapid increase in performance accuracy than larger set sizes during the learning phase of the task. All set sizes had high mean asymptotic performance above 70% accuracy, with set size 2 having the highest asymptotic performance (mean=85%) and set size 6 the lowest (mean=73%).

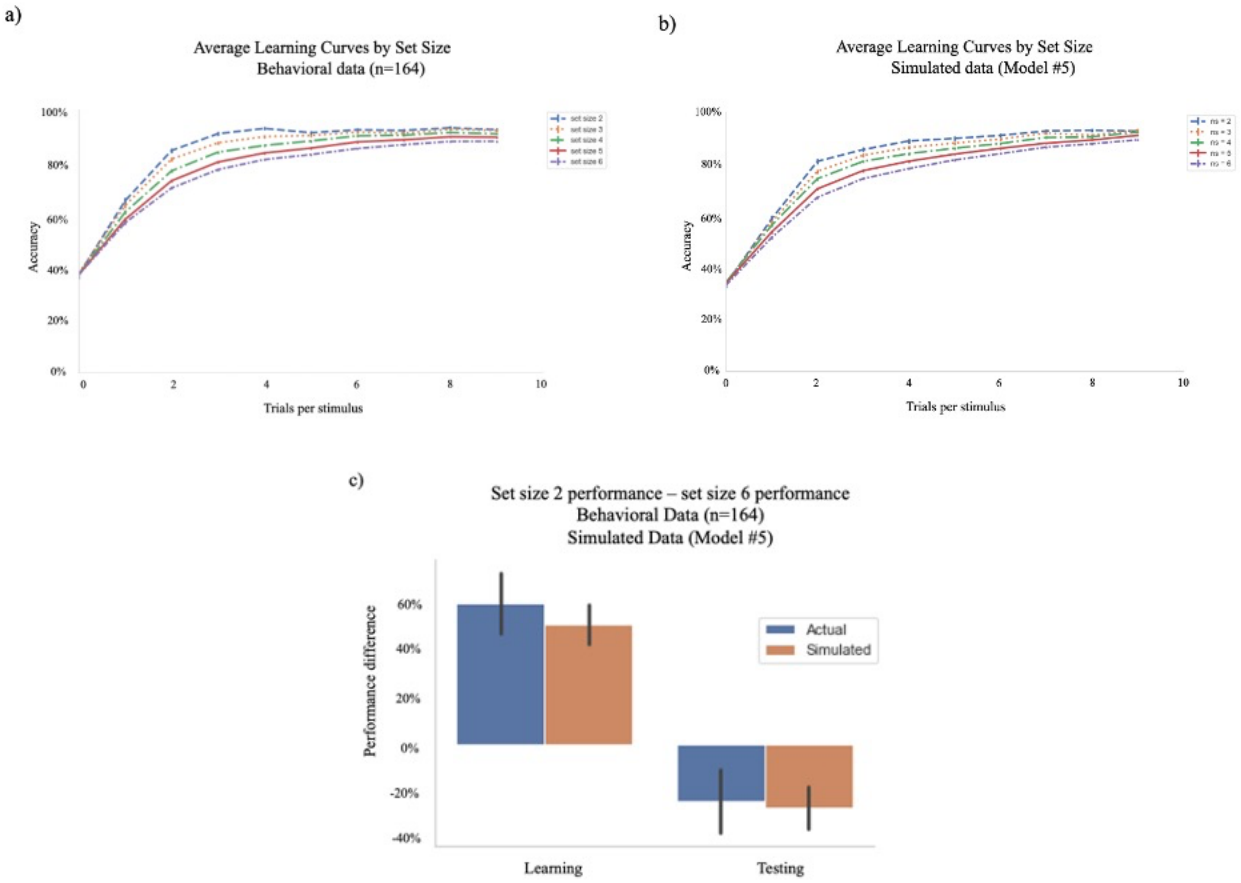


Figure 4.4 Participant behavioral data versus data simulated from winning model (Model #5) Validation of Model #5 in generative capacity to produce realistic behavioral data. a) Average learning curves by set size for actual participant data and b) average learning curves by set size for simulated data from Model #5. c) Difference between set size 2 and set size 6 performance in both learning and testing phases for actual participant data (blue) and data simulated from Model #5 (orange).

We next examined whether our task data showed predicted differences between learning and test phases for high versus low set sizes. The RLWM experimental paradigm makes the prediction that performance for higher set sizes will not decrease as much as performance for lower set sizes between learning and testing phases, given the greater recruitment of reinforcement learning, a longer-term memory process, when set size exceeds working memory capacity during learning.

For each participant, we compared asymptotic learning phase performance (as defined above) to overall mean testing phase performance at each set size. Drop in performance between asymptotic learning and testing was significantly greater than 0 at every set size (SS2: Wilcoxon one-sided statistic = 13358.5, p = 9.65e-29; SS3: Wilcoxon one-sided statistic = 13189.0, p = 2.13e-27; SS4: Wilcoxon one-sided statistic = 12897.5, p = 3.16e-26; SS5: Wilcoxon one-sided statistic = 12286.0, p = 1.07e-24; SS6: Wilcoxon one-sided statistic = 13334.0 p = 2.01e-27). The mean drop in performance between asymptotic learning and testing for set size 2 was 30.4%, and the similarly calculated mean drop in performance for set size 6 was 21.8%. The performance drop for set size 2 was significantly higher than the drop for set size 6 (Wilcoxon one-sided test statistic=9656.0,

p=1.03e-06), indicating that knowledge had been less well retained for the smaller set size between learning and testing phases.

To illustrate the change in relative performance at high and low set sizes between learning and testing, we compared the difference between asymptotic learning performance at set size = 2 versus set size = 6 with the difference between overall testing performance in the same two set sizes (see Figure 4.4). Consistent with previous findings (Collins, 2018; Rac-Lubashevsky et al., 2023; Rmus et al., 2023; Zou et al., 2022), we found that while performance on set size 2 was superior to performance on set size 6 during learning (Wilcoxon one-sided test statistic = 6074.5, p = 4.64e-06), this was not the case during testing (Wilcoxon one-sided test statistic=5943.0, p=0.330).

Scores on the PSWQ were not significantly associated with overall aggregate learning performance (Spearman rho(162) = -0.091, p = 0.248), but had a trending relationship with learning phase set size slope (Spearman rho(162) = 0.149, p = 0.057). A repeated measures ANCOVA with within-subject factor of set size, covariate of z-scored PSWQ scores, and dependent variable of mean learning performance revealed no significant main effect of PSWQ (between-subjects PSWQ F(1,162) = 1.200, p = 0.275) and no significant interaction of PSWQ and set size (within-subject set size x PSWQ effect with Greenhouse-Geisser correction F(3.415, 553.151) = 0.839, p = 0.485). There was a main within-subject effect of set size (within-subject set size effect with Greenhouse-Geisser correction F(3.415, 553.151) = 64.945, p < 0.001). See Supplements Figure S.4.9 for nonparametric correlations and detailed statistics by set size.

A median split by PSWQ (median PSWQ score = 52.0), as shown in Figure 4.5, illustrated a significant group difference in overall learning performance at set size 4 between above-median and below-median PSWQ scores (Mann-Whitney-Wilcoxon-U test, U= 5.512e+7, p=6.653e-3), with higher scores on PSWQ performing significantly worse. This effect was driven by low-performance outliers on set size 4 in the above-median PSWQ group, and there was no significant relationship between PSWQ scores and overall learning performance at set size 4 on a continuous basis (Spearman rho(162) = -0.085, p = 0.281). See Supplements Figure S.4.9 for plot and statistics of PSWQ scores with learning performance at each set size.

There was no significant relationship between higher scores on the PSWQ and overall testing performance (Spearman rho(162) = -0.074, p = 0.347) or testing phase set size slope (Spearman rho(162) = 0.009, p = 0.910). A repeated measures ANCOVA with within-subject factor of set size, covariate of z-scored PSWQ scores, and dependent variable of mean testing performance revealed no significant main effect of PSWQ (between-subjects PSWQ F(1,162) = = 0.672, p = 0.414), and no significant interaction of PSWQ with set size (within-subject set size x PSWQ effect with Greenhouse-Geisser correction F(3.729, 604.140) = 1.321, p = 0.262). There was a main within-subject effect of set size (within-subject set size effect with Greenhouse-Geisser correction F(3.729, 604.140) = 3.982, p = 0.004). See Supplements Figure S.4.10 for nonparametric correlations and detailed statistics by set size.

We additionally tested the relationship between PSWQ scores and drop between asymptotic accuracy over last 3 trials during learning and overall performance during testing. A repeated-measures ANCOVA with within-subject factor of set size, covariate of z-scored PSWQ scores, and dependent variable of drop in performance between asymptotic final 3 trial learning

performance and mean overall testing performance revealed no main effect of PSWQ scores (between-subjects PSWQ $F(1,162) = 0.429$, $p = 0.514$) or interaction between set size and PSWQ (within-subject set size x PSWQ effect $F(4,648) = 1.218$, $p = 0.302$). There was a main effect of set size (within-subjects set size effect $F(4,648) = 9.005$, $p < 0.001$).

In contrast, scores on the MASQ AA subscale were significantly correlated with reduced overall performance during learning (Spearman $rho(162) = -0.299$, $p = 9.79e\text{-}5$). Further, higher MASQ AA subscale scores were also significantly related to higher set size slope during learning (Spearman $rho(162) = 0.220$, $p = 0.005$), confirming greater performance sensitivity during learning to increases in set size. A repeated measures ANCOVA with within-subject factor of set size, covariate of z-scored MASQ AA scores, and dependent variable of mean learning performance revealed a significant within-subject effect of set size (within-subject set size effect with Greenhouse-Geisser correction applied $F(3.430, 555.588) = 66.143$, $p < 0.001$); a significant main effect of MASQ AA (between-subjects MASQ AA effect $F(1,162) = 37.729$, $p <0.001$); and a significant interaction of set size with MASQ AA (within-subject set size x MASQ AA effect with Greenhouse-Geisser correction applied $F(3.430, 555.588) = 3.844$, $p = 0.007$).

We further explored this interaction by conducting parametric correlations between MASQ AA subscale scores and learning performance at each set size. The largest magnitude correlations were for set sizes 5 and 6 (set size 5: Pearson $r(162) = -0.40$, $p = 1.23e\text{-}07$; set size 6: Pearson $r(162) = -0.388$, $p = 2.86e\text{-}07$). However, we also noted a large correlation at set size 2 (Pearson $r(162) = -0.373$, $p = 8.85e\text{-}07$). Correlations for set sizes 3 and 4 were somewhat lower, indicating that the set size effect might be driven largely by this dip at mid-level set sizes (set size 3: Pearson $r(162) = -0.319$, $p = 3.15e\text{-}05$; set size 4: Pearson $r(162) = -0.285$, $p = 0.0002$) and the slightly lower effect at set size 2 compared with set sizes 5 and 6. We note that given the skewed nature of the MASQ AA scores in the present data, parametric correlations may not be highly suitable. Nonparametric correlations did not show substantively different relationships to parametric correlations; see all nonparametric correlations and statistics by set size in Supplements Figure S.4.9.

Higher scores on the MASQ AA subscale were also significantly associated with reduced overall testing performance (Spearman $rho(162) = -0.289$, $p = 0.000$); however, they were not significantly related to testing set size slope (Spearman $rho(162) = 0.099$, $p = 0.205$), suggesting that scores did not relate to increased sensitivity to changes in set size during testing. A repeated measures ANCOVA with within-subject factor of set size, covariate of z-scored MASQ AA scores, and dependent variable of mean testing performance revealed a significant main effect of set size (within-subjects set size effect with Greenhouse-Geisser correction $F(3.744, 606.481) = 4.000$, $p = 0.004$) and a significant main effect of MASQ AA scores (between-subjects MASQ AA effect $F(1,162) = 17.956$, $p < 0.001$), but only a weakly trending effect of interaction between set size and MASQ AA (within-subjects set size x MASQ AA effect with Greenhouse-Geisser correction $F(3.744, 606.481) = 2.046$, $p = 0.091$). See Supplements Figure S.4.10 for plots and nonparametric statistical correlations. .

We additionally tested the relationship between MASQ AA subscale scores and drop between asymptotic accuracy over last 3 trials during learning and overall performance during testing. A repeated measures ANCOVA with within-subject factor of set size, covariate of z-scored MASQ

AA scores, and dependent variable of drop in performance between mean of final 3 learning trials and mean overall testing performance revealed a main effect of set size (within-subjects set size effect $F(4,648) = 8.967$, $p < 0.001$), but no main effect of MASQ AA (between-subjects MASQ AA $F(1,162) = 0.135$, $p = 0.713$) or interaction of MASQ AA with set size (within-subjects set size x MASQ AA effect $F(4,648) = 0.537$, $p = 0.708$).

Higher MASQ AA subscale scores were also significantly associated with reduced relative asymptotic performance between set size 2 and set size 6 during learning (Spearman $rho(162)=0.207$, $p = 0.008$), but not with relative overall performance between set size 2 and set size 6 during testing (Spearman $rho(162) = 0.079$, $p = 0.314$).

e)

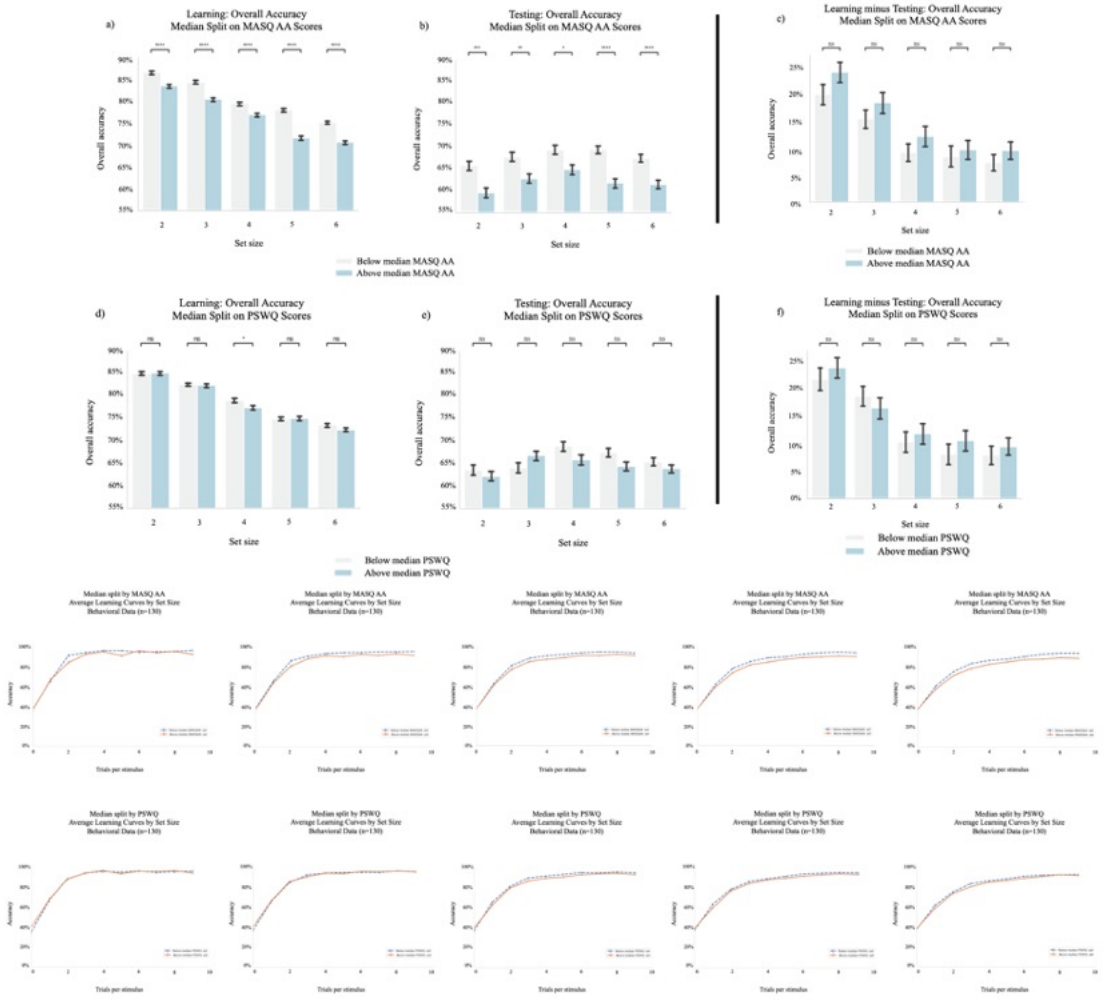


Figure 4.5 Median split of learning and testing performance on MASQ AA and PSWQ a,b) Individuals with above-median scores on MASQ AA subscale have significantly reduced performance accuracy relative to below-median scores at each set size during learning (a) and testing (b) phases. d,e) Individuals with above-median scores on PSWQ do not have significantly different performance relative to below-median scores at any set size during learning (except set size = 4, but see Supplements Figure S.4.9), and no significantly different performance at any set size during testing. e) Learning curves for each set size split by median scores on MASQ AA (top row) and PSWQ (bottom row) reveal relative differences in learning speed and ultimate performance.

Reaction Time

Mean reaction time (RT) across correct learning trials was 599.5ms, and reaction times increased as a function of set size during learning (Kruskal-Wallis statistic(5) = 11308.01, p=0.0). There was a weakly trending effect of PSWQ score on longer overall mean learning reaction time (Spearman rho(162) = 0.139, p = 0.076), but no effect of MASQ AA subscale score on overall mean reaction time (Spearman rho(162) = -0.086, p = 0.276) across set sizes. A nonparametric test of correlation at each set size revealed that higher MASQ AA subscale scores were associated with a significantly shorter reaction time for set size = 5 only (Spearman rho(162) = -0.182, p = 0.02), but found no other significant relationships between MASQ AA or PSWQ scores and learning phase RTs in other set sizes (see Supplements Figure S.4.5 for plots and detailed statistics). Scores on the MASQ AA subscale were significantly related to RT set size slope in learning (Spearman rho(162) = 0.178, p = 0.023), suggesting an increased sensitivity of RT to set size increases for higher MASQ AA subscale scores. PSWQ scores were not related to RT set size slope in learning (Spearman rho(162) = -0.017, p = 0.827).

Mean RT across correct testing phase trials was 557.8ms, and was significantly lower than learning phase RT for correct trials (Mann-Whitney U statistic(1) = 617273595.0, p=1.92e-72). There was no significant effect of set size on testing phase RT (Kruskal-Wallis statistic(5) = 1.36, p=0.851). PSWQ scores were not significantly correlated with longer overall testing phase RT (Spearman rho(162) = 0.118, p = 0.131), nor were MASQ AA subscale scores (Spearman rho(162) = -0.086, p = 0.273). A nonparametric test of correlation at each set size revealed that higher PSWQ scores were associated with significantly longer reaction times for set size = 4 (Spearman rho = 0.224, p = 0.01) and set size = 5 (Spearman rho(162) = 0.174, p = 0.047) only (see Supplements Figure S.4.6 for plots and detailed statistics). There were no other significant relationships between MASQ AA or PSWQ scores and testing phase RTs by set size. RT set size slope in testing was not significantly related to either MASQ AA subscale scores (Spearman rho(162) = -0.038, p = 0.628) or PSWQ scores (Spearman rho(162) = -0.042, p = 0.595).

Computational Modeling

The model which best represented the overall behavioral data as measured by total AIC was Model #5 (RLWM_*asymbias_2rhoK*; see Figure 4.2 and Methods); all results below remain significant when using alternative best-fit Model #4 (see below).

We performed initial hypothesis testing for dimensions of anxiety characterized by physiological arousal (MASQ AA scores) and cognitive worry/anxiety (PSWQ). We tested participant scores on each measure against an initial selection of one RL parameter (learning rate) and four WM parameters (forgetting, WM confidence at low set sizes, WM confidence at high set sizes, and neglect of negative feedback in WM) from the winning model. We performed nonparametric correlation analyses due to the non-normality of the underlying data. We report uncorrected p-values and p-values corrected using the Bonferroni statistical correction for 10 simultaneous comparisons to account for family-wise error across multiple tests.

In initial hypothesis testing, MASQ AA subscale scores were significantly related to lower RL positive learning rate (Spearman rho = -0.249, FWE-corrected p = 0.012) and higher rate of

working memory decay (Spearman rho = 0.255, FWE-corrected p = 0.009). Prior to statistical correction for multiple comparison, greater bias against the use of negative feedback in working memory weights showed a significant relationship with MASQ AA subscale scores, but this did not survive correction for FWE (Spearman rho = 0.188, p = 0.015, FWE-corrected p = 0.151). No significant associations were found between PSWQ scores and any of the initially tested parameter values.

As a robustness check across models, we tested the significant findings between MASQ AA subscale scores and model parameters in the second-closest winning model (Model #4), a nested version of the winning model with fewer parameters. Notably, Model #4 does not include a separate working memory confidence parameter for low versus high set sizes. In Model #4 initial confirmatory testing, higher MASQ AA subscale scores were significantly related to reduced learning rate parameter (Spearman rho = -0.270, p = 0.000); higher rate of working memory decay (Spearman rho = 0.245, p = 0.001); and greater bias against the use of negative feedback in working memory weights (Spearman rho = 0.185, p = 0.017). The relationship between MASQ AA subscale scores and both learning rate and rate of working memory decay survived correction for multiple comparisons over the 3 confirmatory tests (learning rate: 3-comparison FWE-corrected p = 0.001; working memory decay: 3-comparison FWE-corrected p = 0.004), while the relationship with the bias parameter did not (working memory bias parameter: 3-comparison FWE-corrected p = 0.051). We note that if we had performed initial hypothesis testing in Model #4 instead of Model #5, we would have performed 8 total comparisons (4 model parameters from initial hypothesis testing across 2 trait measures). The relationship of MASQ AA scores with both learning rate and with rate of working memory decay also survive this more conservative metric of multiple comparisons (learning rate: 8-comparison FWE-corrected p = 0.003; working memory decay: 8-comparison FWE-corrected p = 0.012).

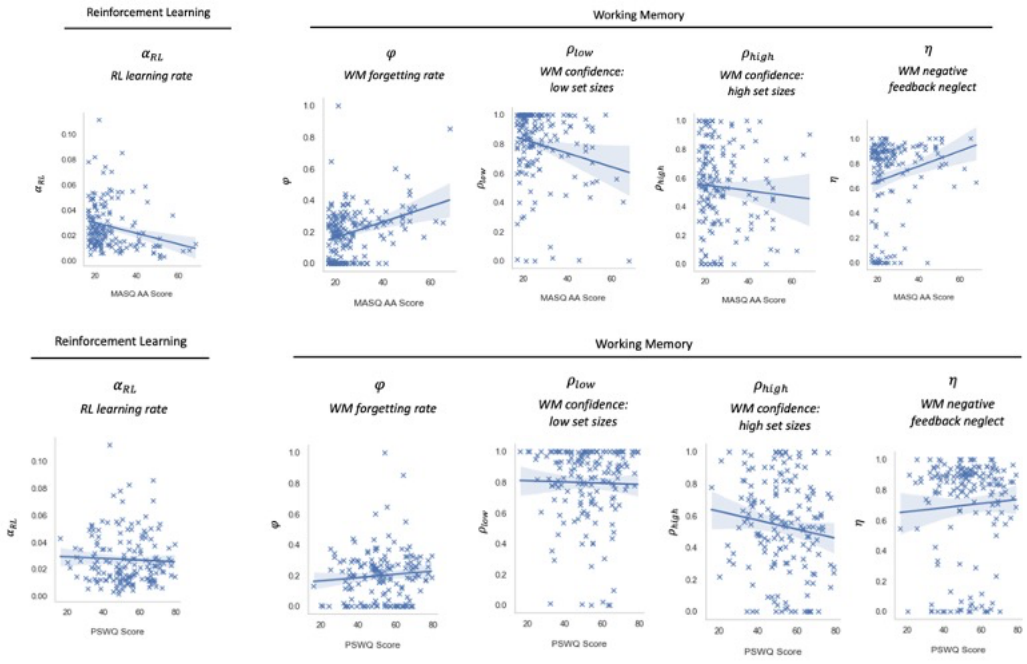


Figure 4.6 Relationship of parameters of learning and working memory with individual scores on MASQ AA and PSWQ From left to right, parameters tested: RL learning rate, WM forgetting rate, WM confidence weight for low and high set sizes, and WM bias against negative feedback updating.

Exploratory modeling analyses

After specific model-based hypothesis testing, we performed an exploratory analysis of the (FWE-uncorrected) relationships between all model parameters and scores on the two measures of anxiety used for hypothesis testing above (MASQ AA and PSWQ), as well as two measures of trait depression (the Center for Epidemiological Studies – Depression (CESD) scale and the Beck Depression Inventory (BDI) II).

There were no significant correlations between MASQ AA, PSWQ, CESD, or BDI scores and any model parameters other than those reported for MASQ AA above.

Discussion

Higher levels of anxiety have been separately associated with impairments in learning and working memory in a number of studies (Balderston et al., 2017; Bishop and Gagne, 2018b; Browning et al., 2015; Forster et al., 2015; LaFreniere and Newman, 2019; Pittig et al., 2018). However, the influence of working memory on reinforcement learning processes in anxiety has not to our knowledge been computationally investigated. In the current analysis, we used an experimental paradigm specifically designed to manipulate the relative load on learning versus working memory systems (Collins and Frank, 2012). We applied a computational model of reinforcement learning which accounts for the supportive role of working memory during learning to test whether relationships between anxiety and learning were attributable to learning processes, working memory processes, or both. We used two measures of anxiety, one specifically characterizing physiological symptoms and one characterizing cognitive symptoms, to test any potential dimensional specificity in results.

Our findings indicate a complex picture of the relationship between anxiety, learning, and working memory. Firstly, we found no behavioral relationship between cognitive anxiety (as measured by PSWQ scores) and reduced performance in learning or testing across the task. This was consistent with the results of computational modeling analysis: no parameters of cognitive mechanisms were found to have significant relationships with increased levels of cognitive anxiety.

In contrast, we found significant behavioral impairment in task performance at every stimulus set size across both learning and testing phases for individuals with higher levels of physiological anxiety (as measured by MASQ AA subscale scores). This performance impairment was mechanistically attributable to multiple parameterized mechanisms in computational modeling analysis. After correcting for multiple comparisons, MASQ AA subscale scores were related to a significantly lower learning rate and a significantly higher rate of working memory decay. This set of results suggests a multi-faceted relationship of anxiety with learning and its supportive processes, whereby both the learning process itself as well as the role of the supporting working memory systems are compromised in high levels of physiological anxiety. In future analyses it would be valuable to characterize whether a traditional RL modeling framework fit to this data would attribute learning deficits to any computational modeling parameter; for instance, it might partition even more explained variance onto the learning rate parameter; alternatively, it might

weaken the specificity of the explanatory power of the model, resulting in no significant parametric results. This test, while not performed here, would be valuable to understand the full import of our findings.

Although anxiety has been found to relate to learning impairments across studies, it has rarely been reliably linked to reduced rate of reinforcement learning in computational modeling studies. Most prior studies of RL in computational psychiatry research have not included a parameterization of the supportive role of working memory during the RL process. We note that including the WM system in the computational modeling process might appropriately allocate variance differently than a pure RL model, thereby revealing significant differences in RL parameters such as learning rate which would not have otherwise been found. Although most RL studies do not specifically manipulate RL and WM loads as the current experimental paradigm does, most RL examinations of trait anxiety have used experimental task with a small number of stimuli (often around 2). In such paradigms, working memory is likely to have a large contributing role to learning which has not typically been evaluated.

Following hypothesis testing of physiological and cognitive anxiety with parametric measures of learning and working memory, we performed a full exploratory analysis across all computational model parameters with both measures of anxiety, as well as with two commonly used measures of depression (CES-D and BDI-II). We found no additional significant relationships between any of these measures and computational model parameters.

Interestingly, at the statistical trend level we observed that scores on both MASQ AA and BDI were related to a reduced confidence in working memory at low set sizes. Although not reaching statistical significance, this could inform future hypotheses regarding potential joint versus distinct influences of depression and anxiety on learning. Future studies might investigate whether a dimension of variance common to depression and anxiety (frequently referred to as "internalizing" psychopathology variance) is related to reduced WM contributions to low set size learning, while the physiological anxiety dimension specifically is related to impairments in learning rate and WM decay.

Additionally, both anxiety and depression have been associated with differential valuation of decision-making outcomes across multiple outcome valences (Bishop and Gagne, 2018a); while depression has been associated with reduced valuation of reward (Esfand et al., 2024; Halahakoon et al., 2020; Whitton et al., 2015), anxiety has been associated with increased valuation of threatening outcomes (Chen and Lovibond, 2020; Geng et al., 2018; Robinson et al., 2013). Given the results of the current analysis, additional studies could investigate whether the same parametric relationships hold in a dual-valence learning task, where both threat avoidance learning and reward learning are tested.

We also note that although we have tested a number of models from the RLWM research space, including one novel parameterization to this analysis (separate WM confidence parameters for low versus high set sizes), there may be additional model formulations which could be appropriate to test. Although this claim could potentially be made with respect to nearly any computational investigation of behavior, in the present case a more careful exploration of the hypothesis space may be warranted given the newer research domain.

Overall, the study findings provide insight to the potentially complicated relationship of anxiety with interactive systems of WM and learning. Further study is required to extend the research hypotheses across valence and internalizing psychopathology more broadly.

Chapter 4: Supplemental materials

Figure S.4.1 Distributions of questionnaire metrics in final participant group (n=164)
Red vertical line denotes group mean; black vertical line denotes group median.

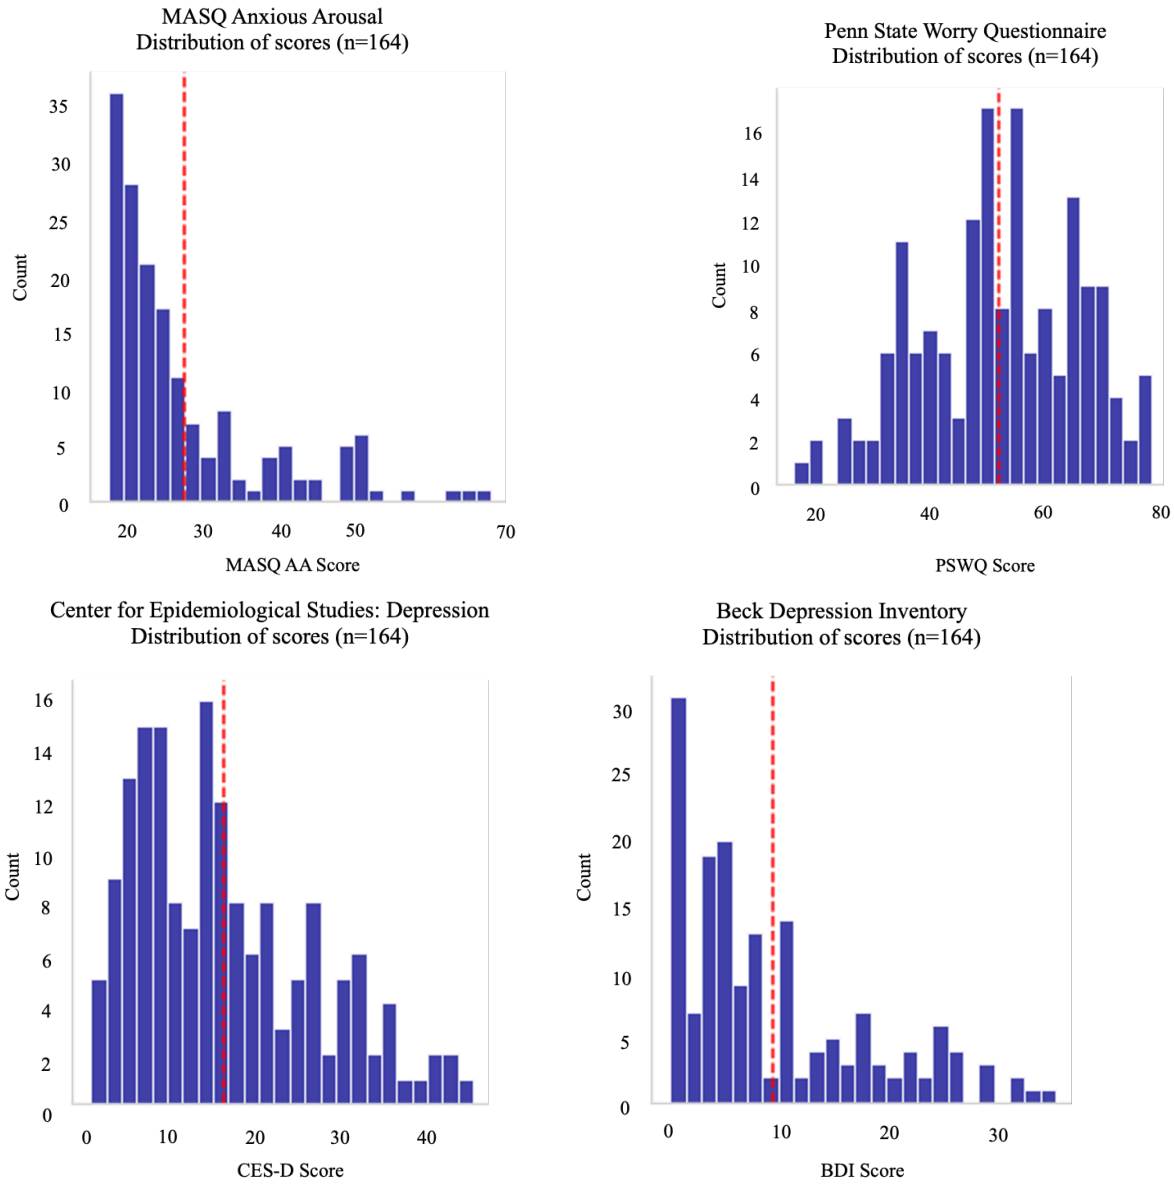


Figure S.4.2: Parameter value distributions: Model #5

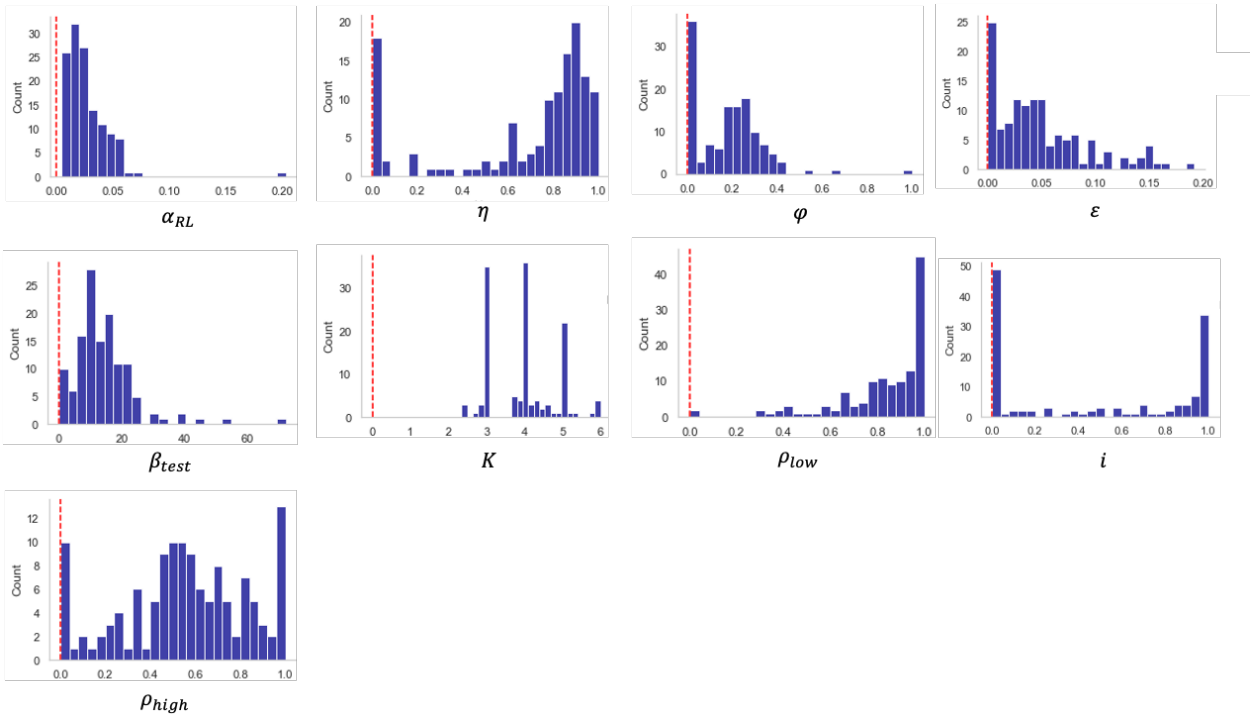


Figure S.4.3: Parameter recovery analyses: Model #5

Error between recovered and actual parameter values for winning model centered at zero for all parameters. Correlation between actual and recovered parameters measured with Spearman rho.

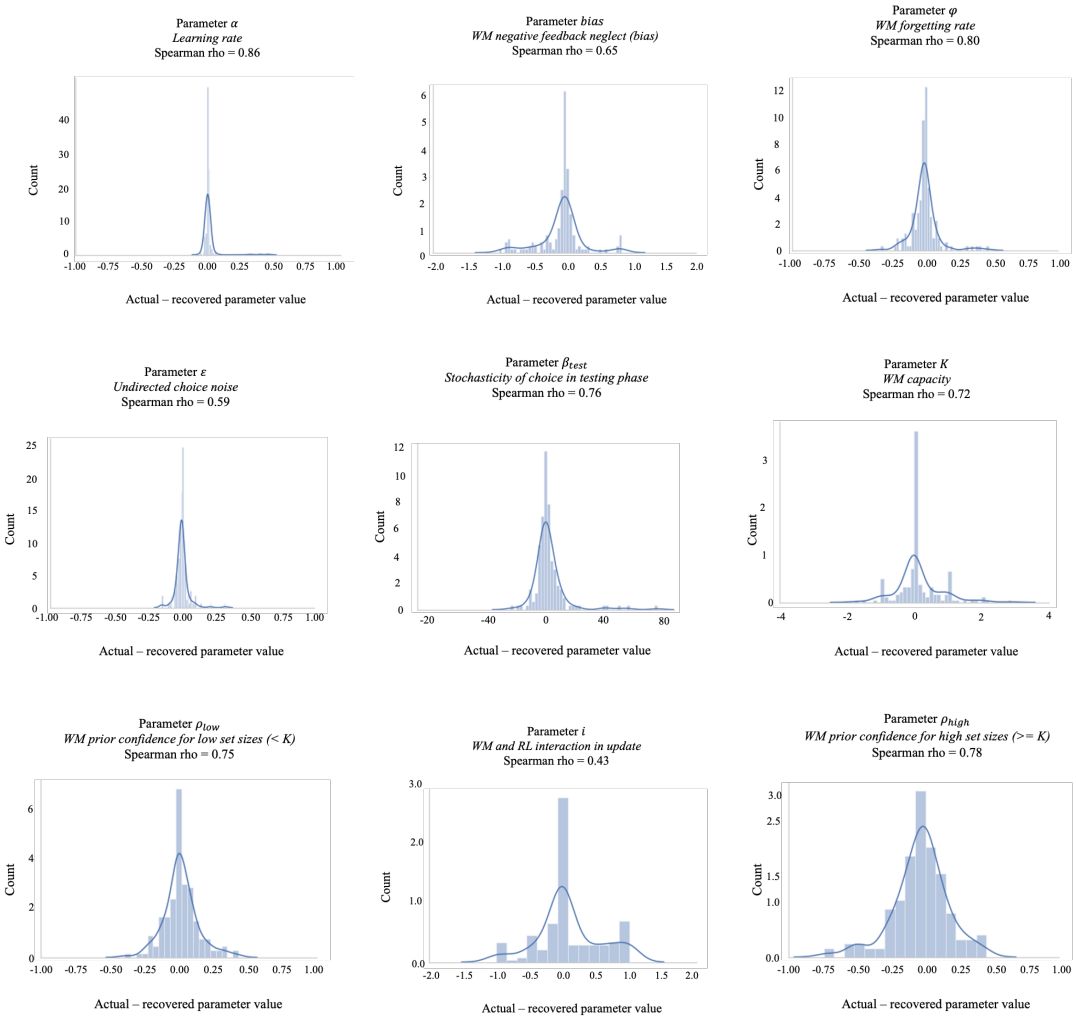


Figure S.4.4: Model recovery analysis

Comparison of frequency with which the generative model provided the best fit to the data across the top 3performing models (based on AIC comparison).

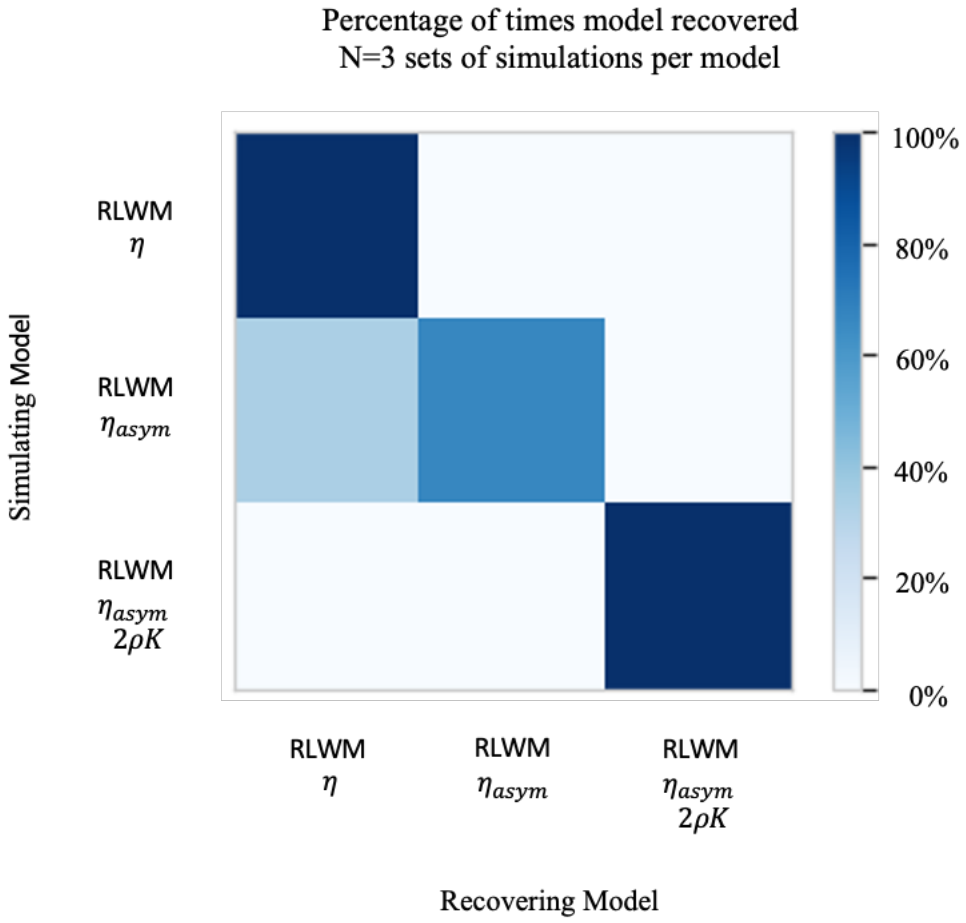


Figure S.4.5: Learning phase reaction times by set size as a function of anxiety scores

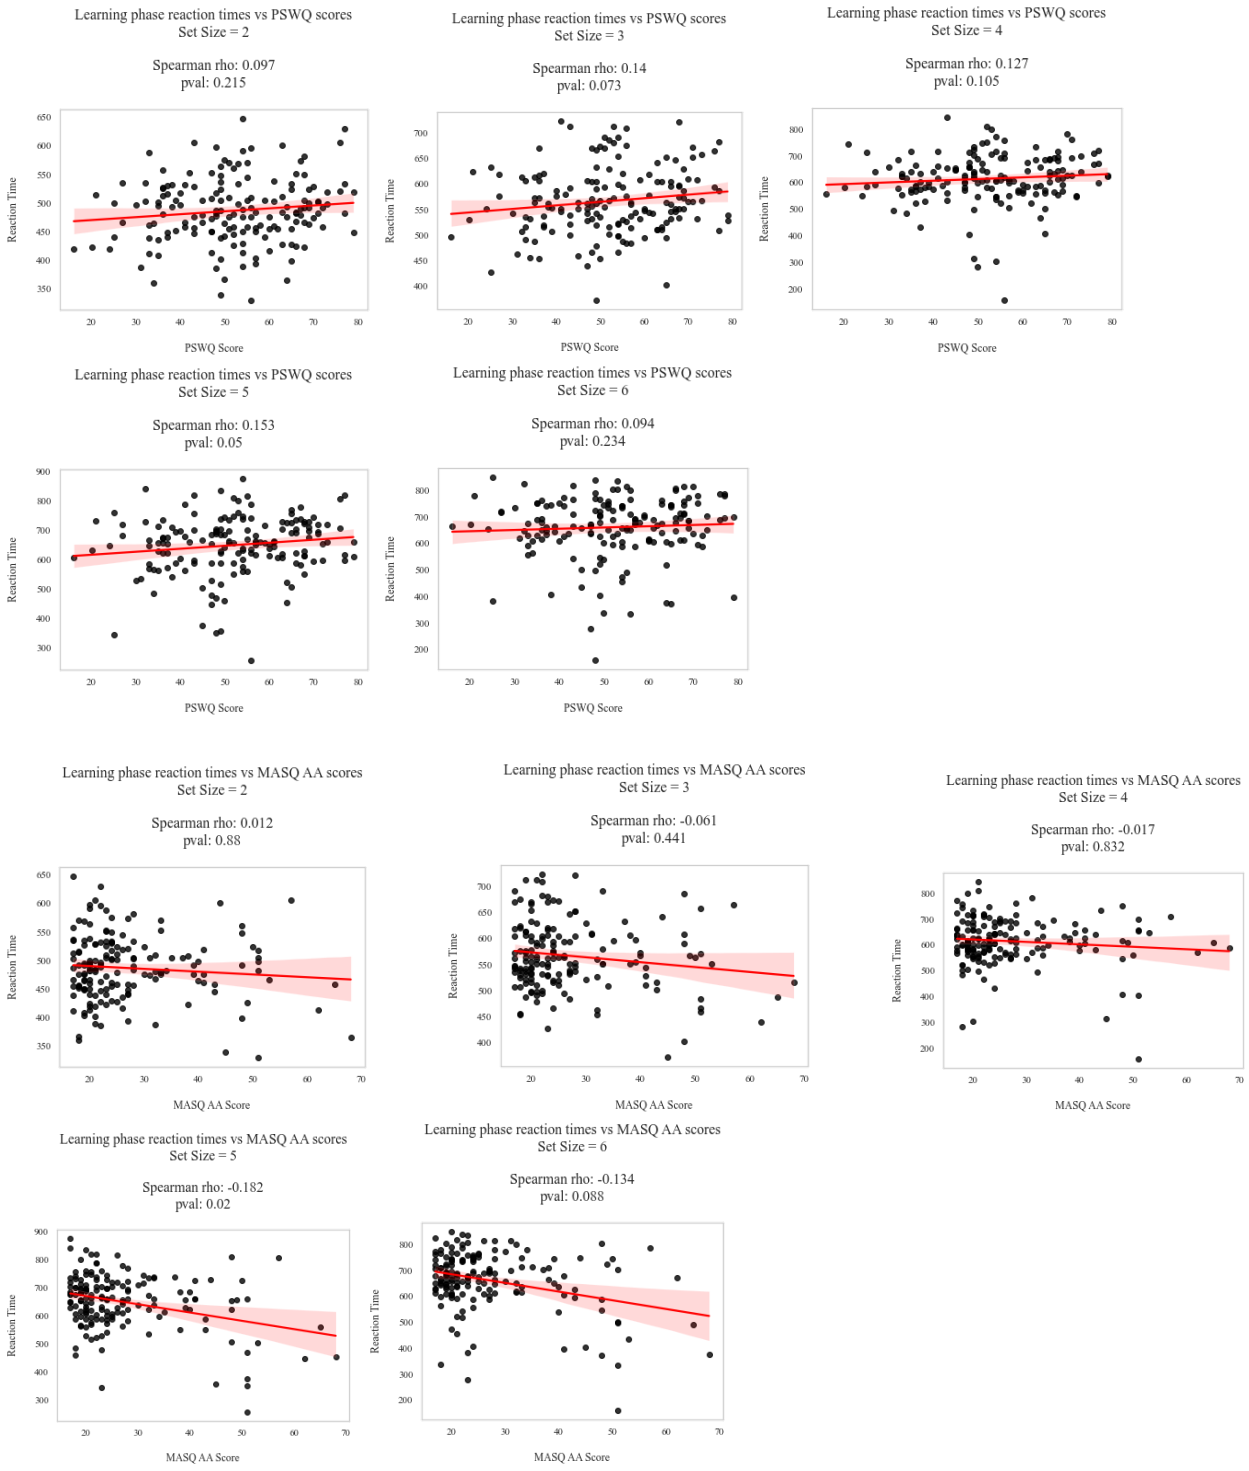


Figure S.4.6: Testing phase reaction times by set size as a function of anxiety scores

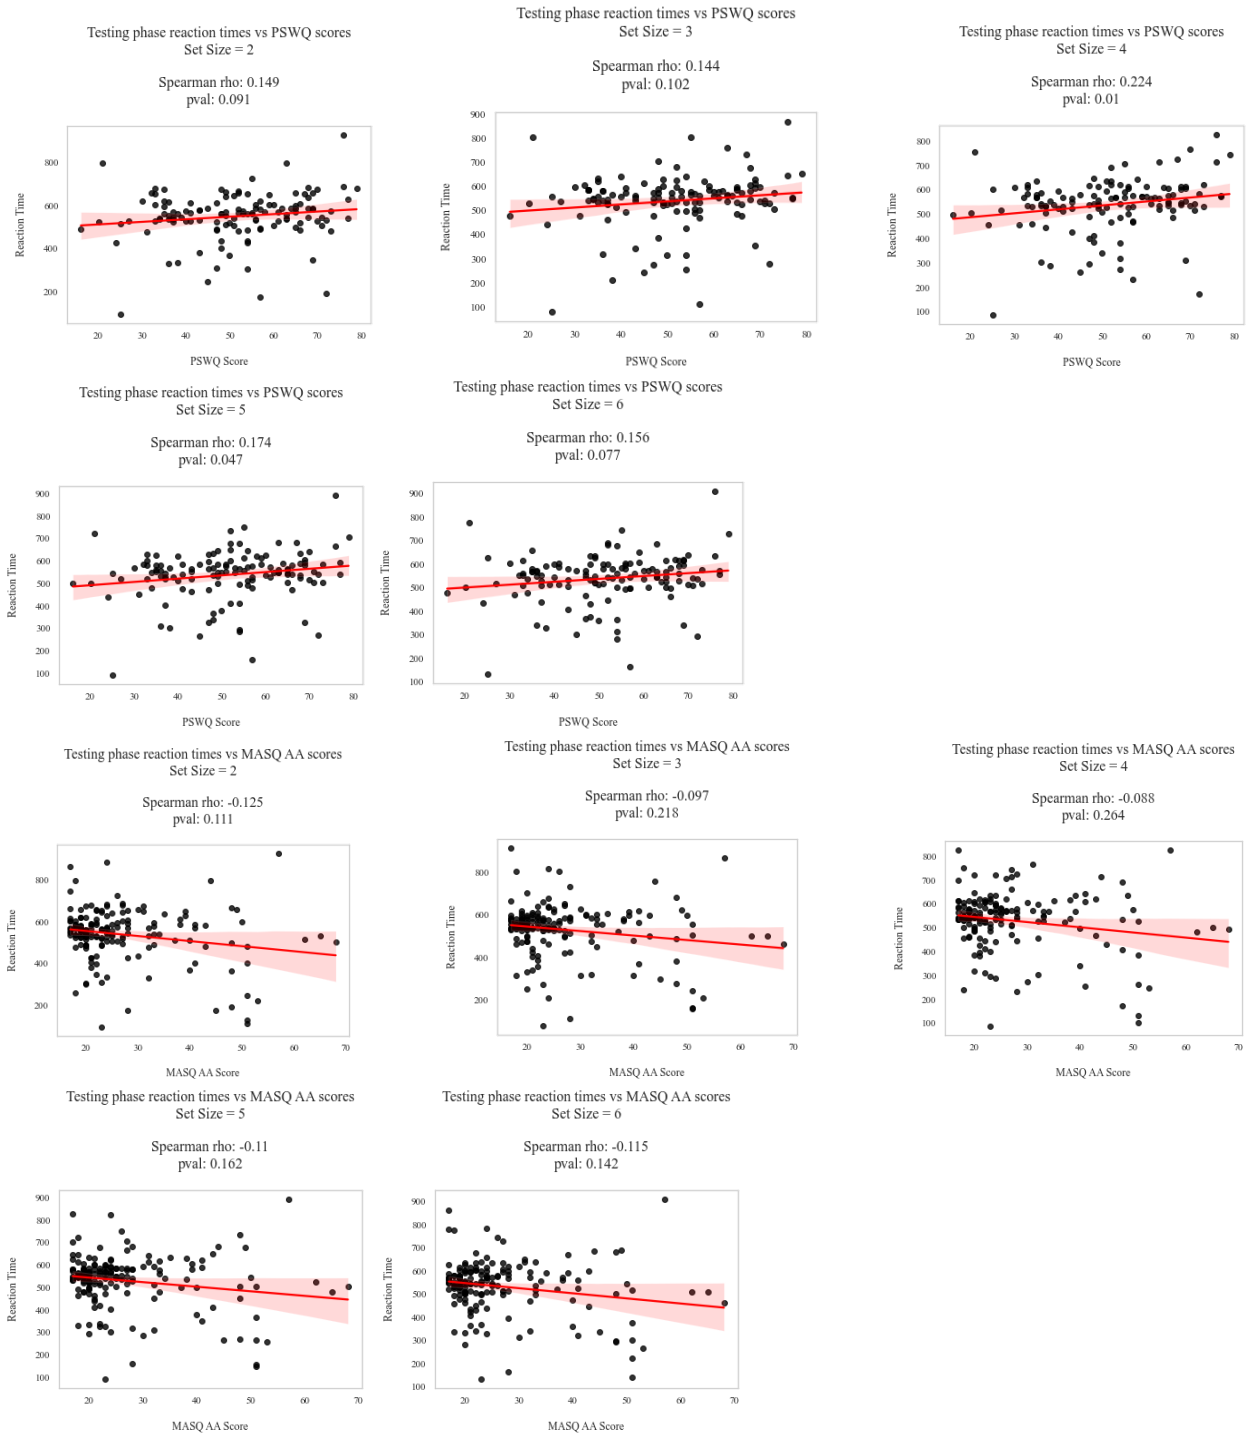


Figure S.4.7: Exploratory parameter testing Model #5

Exploratory testing of Model #5 parameters with MASQ AA and PSWQ scores. p-values uncorrected for multiple comparisons/family-wise error.

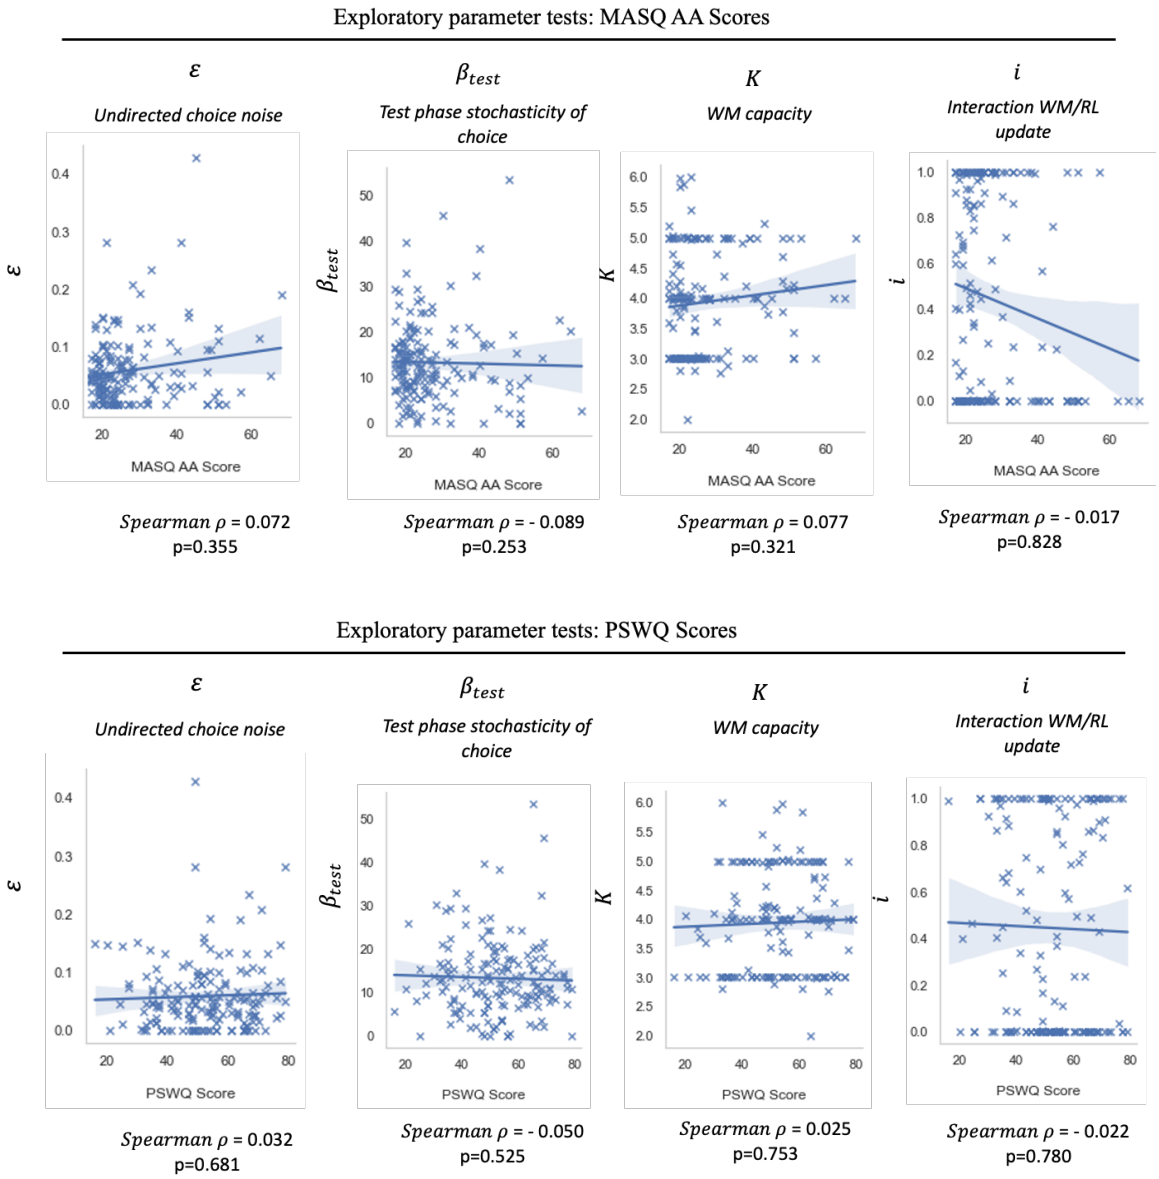


Figure S.4.8: Exclusion statistics

Distribution of number of timeouts per participant prior to timeout-based exclusions (left) and distribution of participant asymptotic accuracy in learning over last 5 stimulus presentations at set size 2 (right). Dashed blue lines denote mean. Exclusion criterion was 2SD from mean (mean + 2sd for timeouts; mean – 2sd for set size 2 performance).

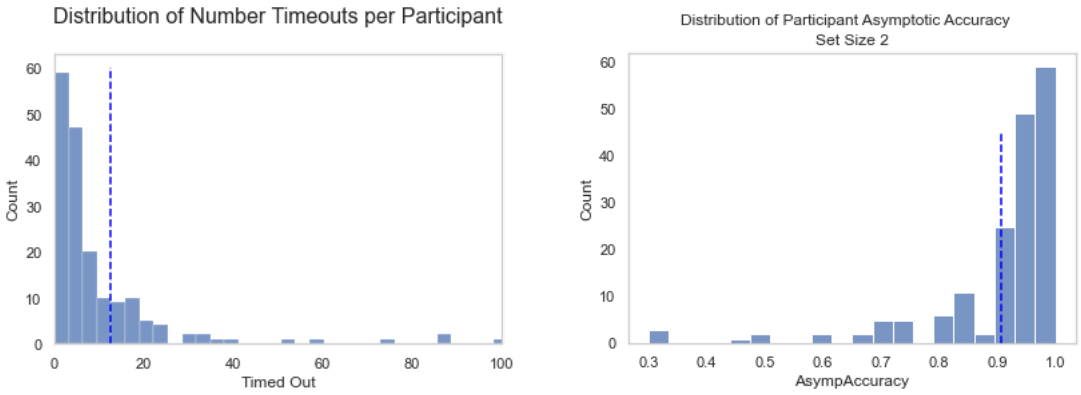
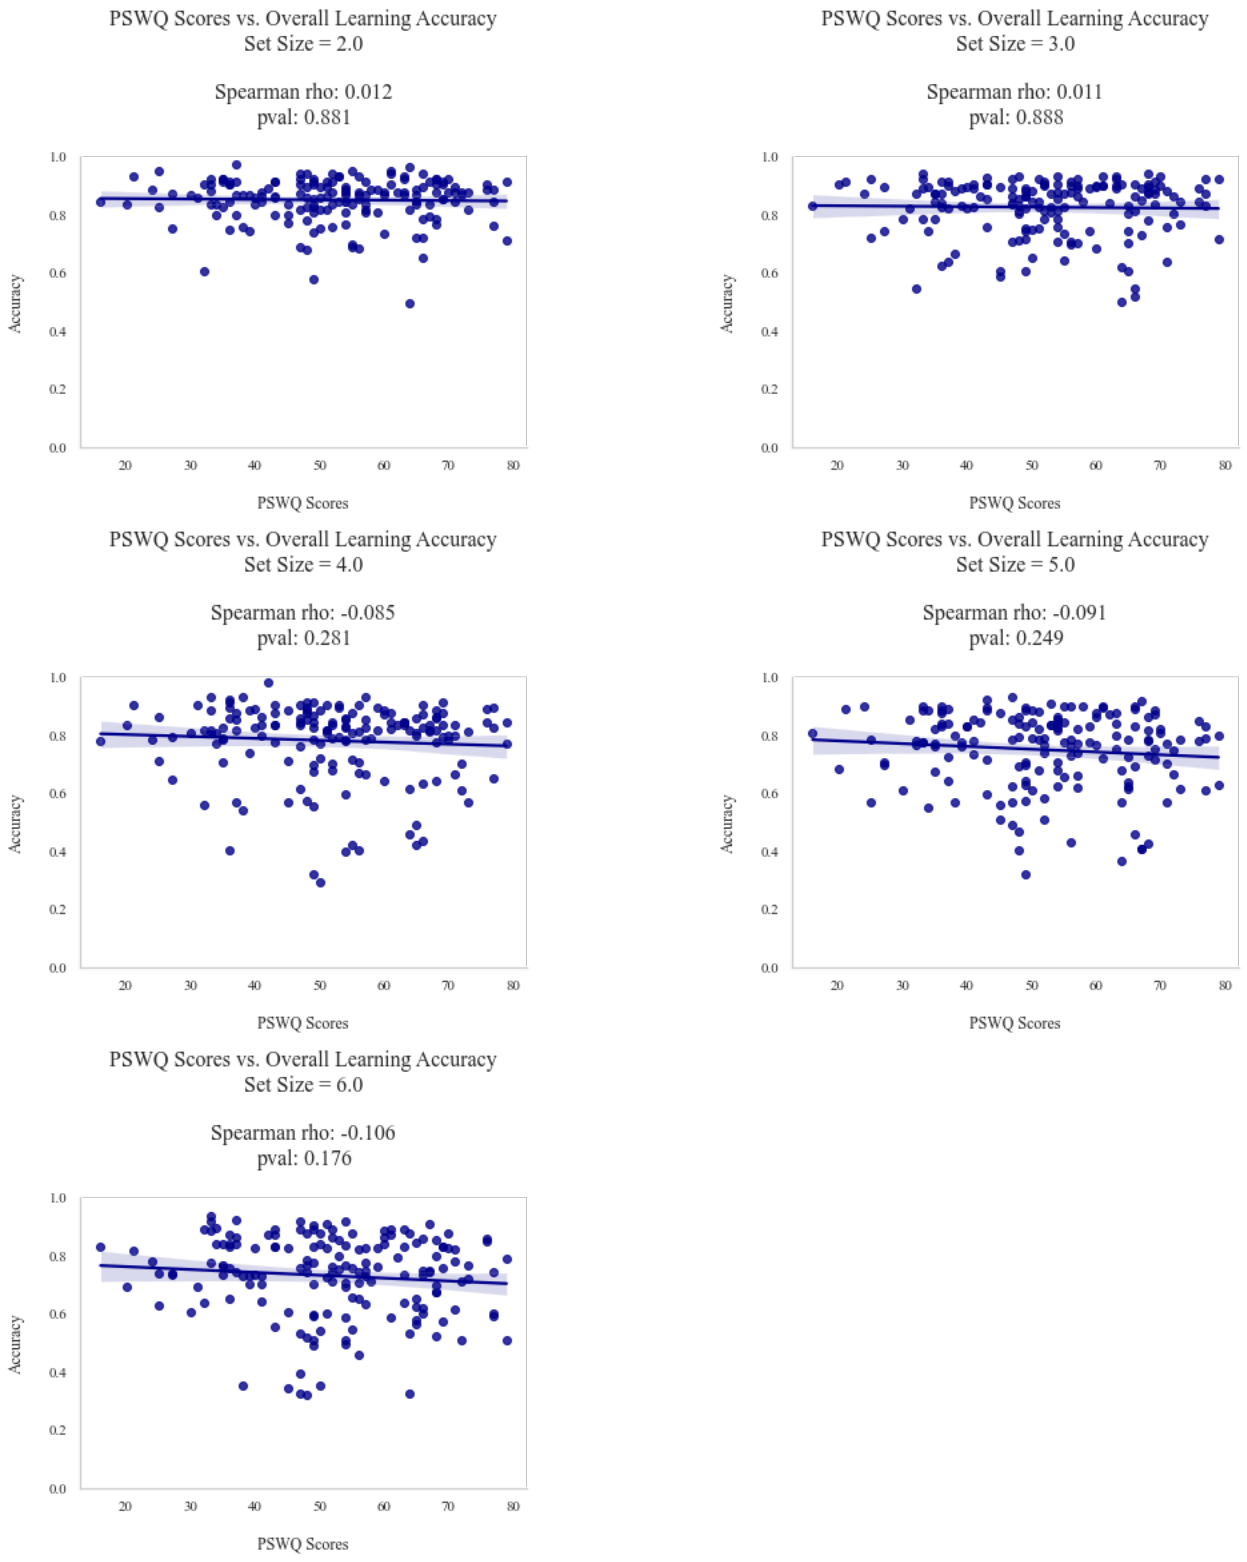
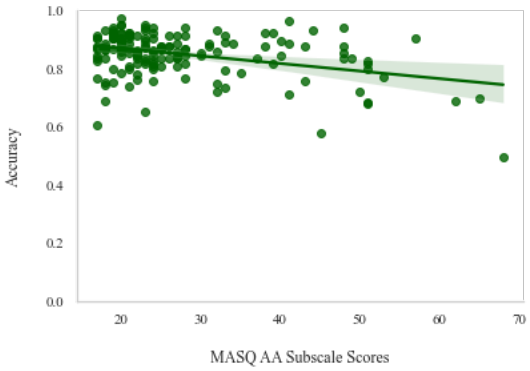


Figure S.4.9: Correlation of PSWQ and MASQ AA subscale scores with overall learning performance



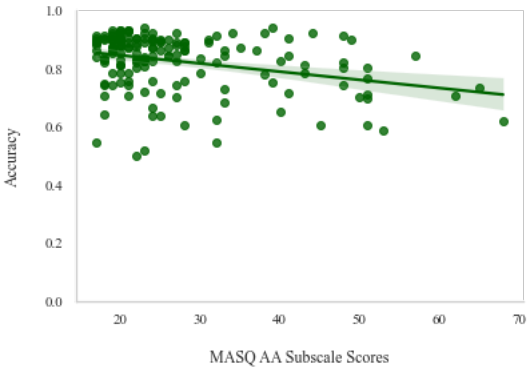
MASQ AA Subscale Scores vs. Overall Learning Accuracy
Set Size = 2.0

Spearman rho: -0.23
pval: 0.003



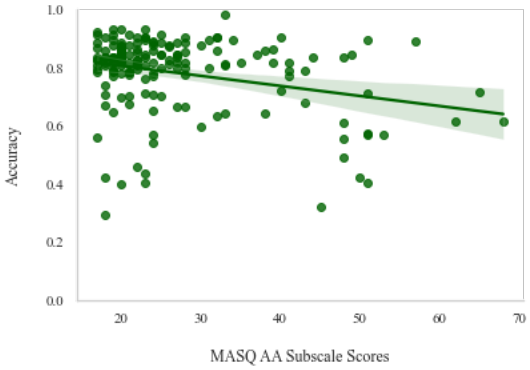
MASQ AA Subscale Scores vs. Overall Learning Accuracy
Set Size = 3.0

Spearman rho: -0.271
pval: 0.0



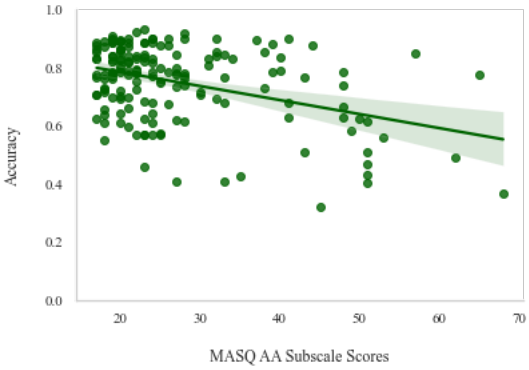
MASQ AA Subscale Scores vs. Overall Learning Accuracy
Set Size = 4.0

Spearman rho: -0.149
pval: 0.058



MASQ AA Subscale Scores vs. Overall Learning Accuracy
Set Size = 5.0

Spearman rho: -0.243
pval: 0.002



MASQ AA Subscale Scores vs. Overall Learning Accuracy
Set Size = 6.0

Spearman rho: -0.26
pval: 0.001

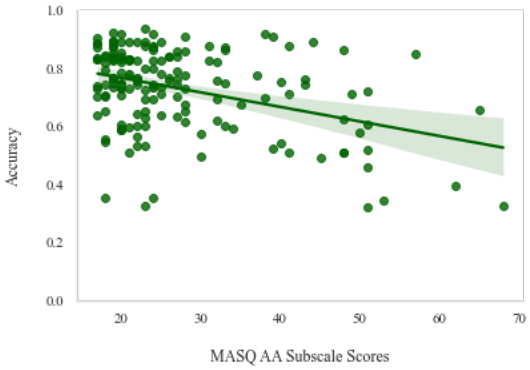
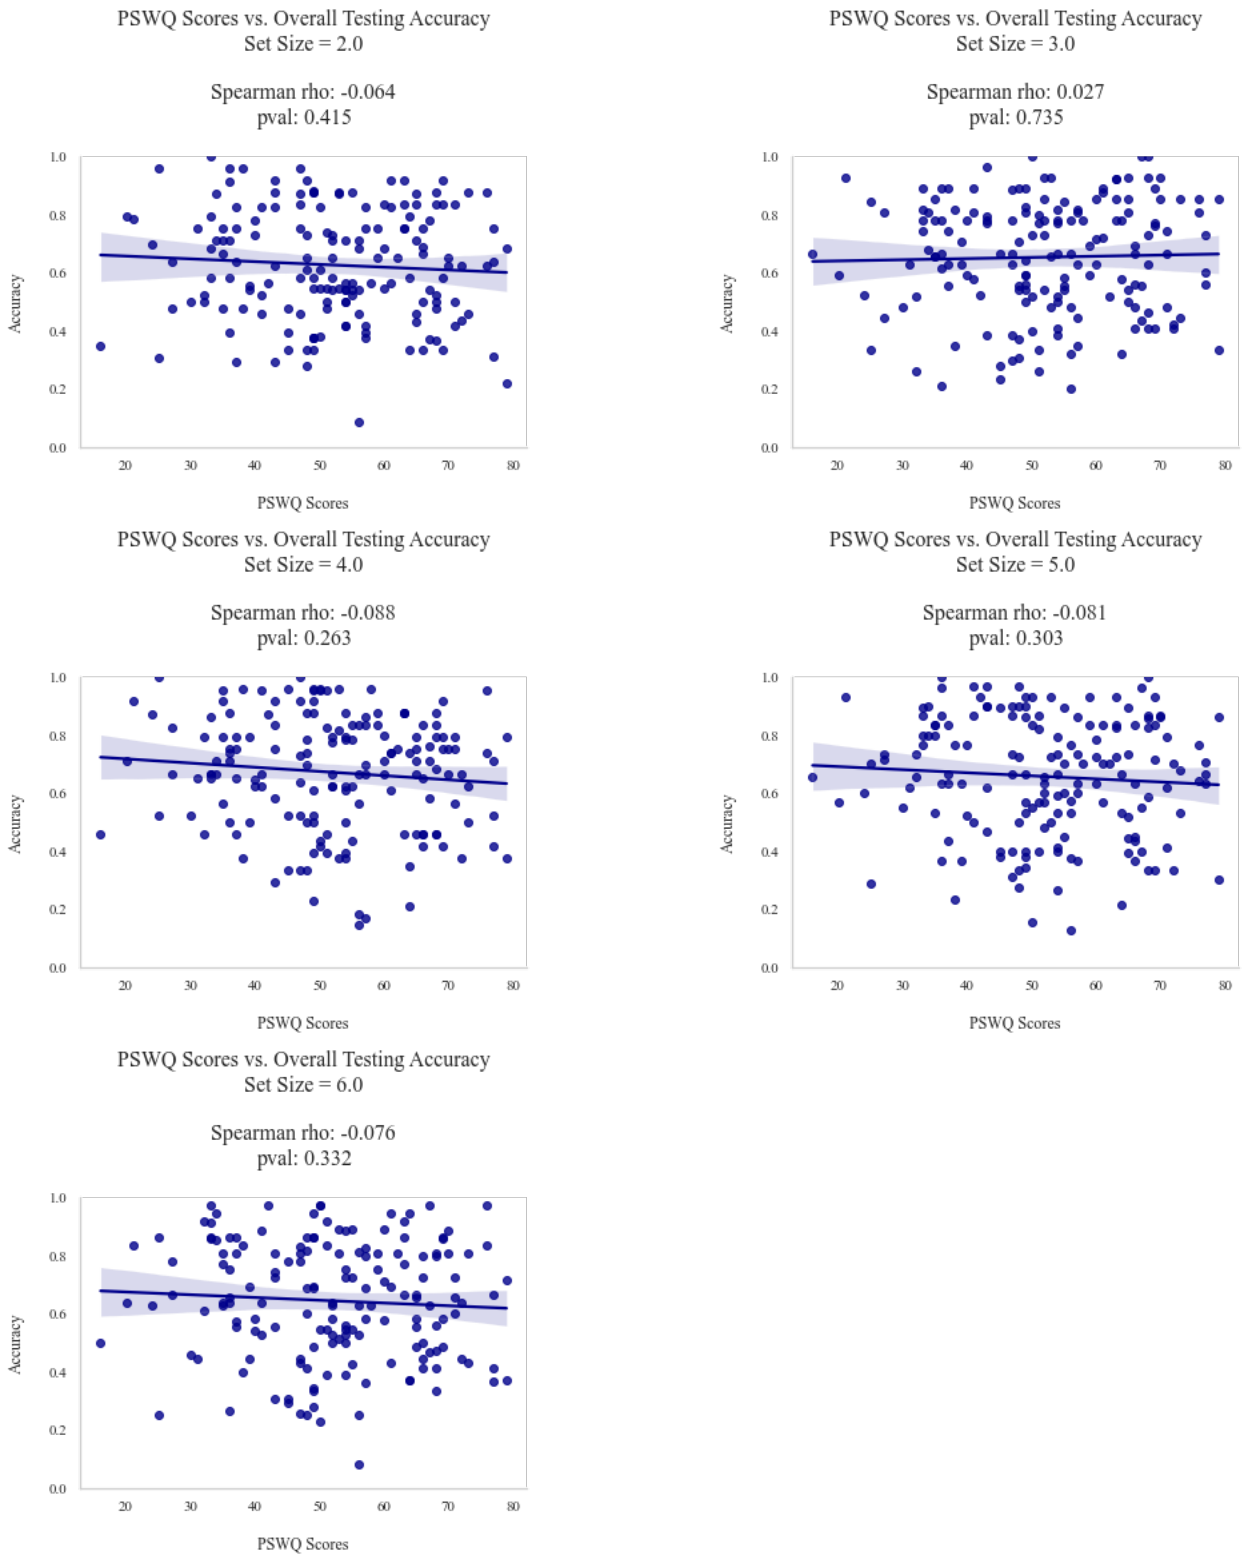
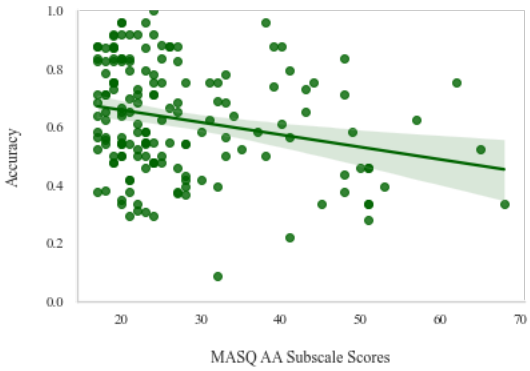


Figure S.4.10: Correlation of PSWQ and MASQ AA subscale scores with overall testing performance



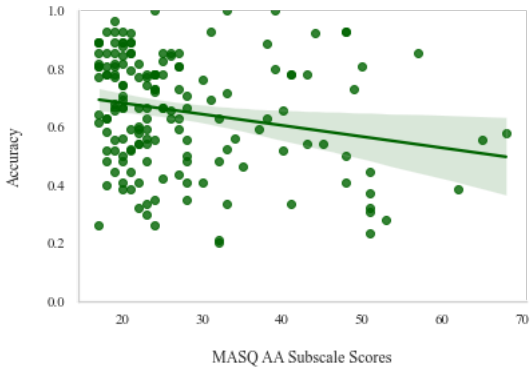
MASQ AA Subscale Scores vs. Overall Testing Accuracy
Set Size = 2.0

Spearman rho: -0.241
pval: 0.002



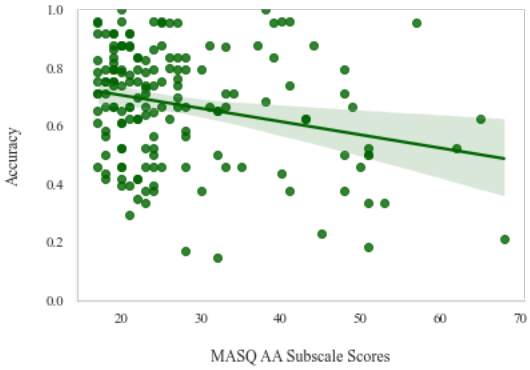
MASQ AA Subscale Scores vs. Overall Testing Accuracy
Set Size = 3.0

Spearman rho: -0.233
pval: 0.003



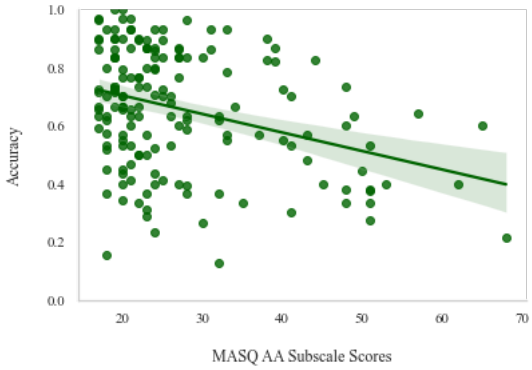
MASQ AA Subscale Scores vs. Overall Testing Accuracy
Set Size = 4.0

Spearman rho: -0.19
pval: 0.015



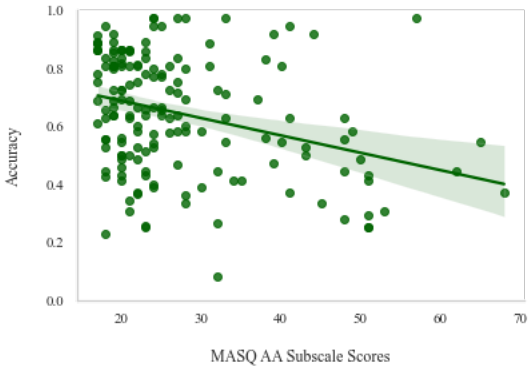
MASQ AA Subscale Scores vs. Overall Testing Accuracy
Set Size = 5.0

Spearman rho: -0.276
pval: 0.0



MASQ AA Subscale Scores vs. Overall Testing Accuracy
Set Size = 6.0

Spearman rho: -0.303
pval: 0.0



Chapter 5

Discussion

In this dissertation, we have explored the computational underpinnings of decision-making and learning with an aim of elucidating how the related cognitive mechanisms are linked to dimensions of internalizing psychopathology. We sought to study this across several different cognitive processes: effortful decision-making in reward pursuit and threat avoidance (Chapter 2); decision-making under ambiguity in reward, loss, and threat paradigms (Chapter 3); and the interactive role of reinforcement learning processes with working memory systems (Chapter 4). Across analyses, we have used varying levels of dimensionality reduction to characterize latent constructs within internalizing psychopathology. In effortful decision-making, we sought to assess common versus symptom-level component dimensions of measures of anxiety and depression, including apathy and anhedonia. In decision-making under ambiguity, our aim was to test common versus distinct dimensions across measures of anxiety, depression, and hypomania, investigating whether mechanistic differences might relate to comorbidity or distinct trait variance. Finally, in our reinforcement learning and working memory investigation, we sought to test two specific characterizations of anxiety (physiological and cognitive) and used well-validated questionnaire measures to capture each of these distinctly.

Our investigation of effortful decision-making, found in Chapter 2 of the dissertation, used a hierarchical Bayesian model of value-based decision-making which incorporated latent dimensions of psychopathology into the hierarchically inferred structure. Our analysis revealed that symptom level characterizations of anhedonia and apathy, while resulting in similar behavioral differences, acted via distinct cognitive mechanisms in the decision-making process for effortful reward pursuit. We also found that higher factor scores characterized by cognitive anxiety were related to greater fatigue in effortful reward pursuit, while those characterized by physiological anxiety were associated with increased differences in valuation of reward magnitude, suggesting the relevance of dimensional investigation of anxiety in reward pursuit paradigms, an area which has not traditionally been well-studied. Higher cognitive anxiety was also related to increased bias toward effortful threat avoidance, and a concurrent reduced sensitivity to changes in magnitude of threat.

In our decision-making under ambiguity work (Chapter 3), we used an experimental paradigm that varied the level of missing information across ambiguous choice trials. Here we replicated the classic gain-loss framing effect of risky decision-making (Kahneman and Tversky, 1979; Tversy and Kahneman, 1974) and the similar valence-based framing effect in ambiguous decision-making (Kocher et al., 2018; Trautmann and Van De Kuilen, 2015). Participants tended to be broadly ambiguity avoidant in reward gain, but ambiguity preference in negative valence domains (loss and threat) was modulated between ambiguity seeking and ambiguity avoidance by level of missing information. The computational model of decision-making which best fit behavior in each task valence (reward, loss, and threat) was found to be the same; the winning model parametrically incorporated both separate and joint consideration of magnitude and probability differences between choice options. In all cases, a model which incorporated an individual optimistic or

pessimistic bias for the inference of uncertain probability information best described participant behavior. Parameters were jointly inferred at the individual and group level using a hierarchical Bayesian approach, and resulting participant mean posterior parameter values in each valence paradigm were tested against scores from a bifactor model of anxiety, depression, and hypomania. Risk-seeking behavior in mania has been associated with increased sensitivity in the behavioral activation system (BAS) and alterations in reward processing (Johnson et al., 2012); here we sought to test whether risk preference or ambiguity sensitivity might also be related to such behavioral differences. While we found no relationship between latent dimension scores characterized by hypomania and reward decisions under ambiguity, we found that ambiguous threat avoidance decisions significantly characterized a dissociation between risk sensitivity in manic tendencies versus anhedonic depression, with higher levels of manic symptom severity associated with reduced risk sensitivity when contrasted with anhedonic depression. We also found significant differences in loss avoidance risk attitudes for individuals with anhedonic depression (more risk sensitive) and individuals with physiological arousal (reduced risk valuation), suggesting that loss avoidance under ambiguity may be a valuable mechanistic candidate for further computational study of mechanisms of anxiety versus depression.

Finally, our exploration of reinforcement learning and working memory (Chapter 4) fit an interactive computational model of joint RL and WM processes of learning. This computational model, fit at an individual level using maximum likelihood estimation, revealed that physiological anxiety is related to significant impairment in both learning and the support of the working memory system during learning. In contrast, a measure of cognitive anxiety showed no relationship with learning performance or the underlying computational mechanisms of learning or working memory. These results provide insight to the dual role of reinforcement learning and working memory systems in learning impairment among individuals with heightened physiological anxiety, and moreover emphasize the importance of a dimensional characterization of anxious symptomatology when investigating individual differences in learning and memory.

Taken together, the research results here suggest differentiable behavioral profiles for individuals suffering from classic symptoms of depression with no comorbid symptoms, versus those suffering comorbid depression and anxiety, versus those suffering only from anxiety. While individuals with symptoms of physiological anxiety are likely to show impaired reinforcement learning and working memory support of learning, those with cognitive anxiety and/or depressive symptoms are more likely to show normative learning performance (Chapter 4). Meanwhile, both depressive and cognitive anxiety symptomatology show relationships with respect to effortful decision-making. Here, those with classic symptoms of depression such as apathy and anhedonia are likely to show multiple mechanistic impairments to motivation towards effortful reward pursuit, with apathy contributing to motivational action bias and increased fatigue, and anhedonia contributing to greater reward requirements to incentivize pursuit (Chapter 2). While cognitive worry may also contribute to increased fatigue in this paradigm, it may be uniquely characterized within effortful decision-making by differences in threat avoidance, where it is related to increased bias towards engagement and a decreased sensitivity to increases in threat magnitude (Chapter 2). Physiological anxiety symptoms may contribute to increased reward requirements to motivate engagement in effortful reward pursuit but may be isolated from the other dimensional relationships with effortful decision-making via its influence on learning and ambiguous decision-making (Chapters 3,4). Our results in ambiguous decision-making suggest that loss avoidance decisions may directly

differentiate anhedonic depression from physiological anxiety, with higher levels of depression showing greater risk sensitivity in loss avoidance and higher levels of physiological anxiety showing reduced risk sensitivity in the same paradigm (Chapter 3). Meanwhile, physiological versus cognitive anxiety may be directly differentiated via impact on learning processes as described in Chapter 4.

In addition, attitudes towards risk in threat avoidance decisions under ambiguity may provide differential characterizations of mania from anhedonic depression, with higher levels of mania being significantly related to lower risk sensitivity than higher levels of anhedonic depression (Chapter 3). If this finding is replicated it would provide insight into a potential valence-specific decreased risk sensitivity in mania relative to depression.

As outlined in the introduction to the dissertation, one ultimate application of computational psychiatry research lies in the development of computational phenotypes of dimensions of psychopathology. The differential characterizations of symptom profiles in learning and decision-making algorithms presented in this dissertation are only the first steps in pursuit of computational phenotypes to characterize these or other profiles of symptomatology, and must be replicated to establish reliability. As the field gains increasingly more nuanced information regarding the relationship of latent dimensions of psychiatric symptoms to specific mechanisms of cognitive impairment, we are in turn enabled to propose even more nuanced models to characterize the true underlying processes. Through this growing body of knowledge, and with a focus on replicability and the clear communication of modeling techniques, we are enabled to approach the development of robust computational phenotypes. The work included in this dissertation aims to contribute in some small way to the larger (inter-)disciplinary effort towards computational phenotyping, with the hope that by each standing on the shoulders of giants we might someday realize the goal of individualized psychiatric diagnosis and treatment, and the improvement or cessation of suffering for many individuals affected by mental illness.

References

Actuarial Standards Board, 2019. Actuarial Standard of Practice No. 56: Modeling.

Adida, M., Clark, L., Pomietto, P., Kaladjian, A., Besnier, N., Azorin, J., Jeanningros, R., Goodwin, G.M., 2008. Lack of insight may predict impaired decision making in manic patients. Bipolar Disord. 10, 829–837. https://doi.org/10.1111/j.1399-5618.2008.00618.x

Adida, M., Jollant, F., Clark, L., Besnier, N., Guillaume, S., Kaladjian, A., Mazzola-Pomietto, P., Jeanningros, R., Goodwin, G.M., Azorin, J.-M., Courtet, P., 2011. Trait-Related Decision-Making Impairment in the Three Phases of Bipolar Disorder. Biol. Psychiatry 70, 357–365. https://doi.org/10.1016/j.biopsych.2011.01.018

Akaike, H., 1974. A new look at the statistical model identification. IEEE Trans. Autom. Control 716–723.

American Psychiatric Association., 2013. Diagnostic and statistical manual of mental disorders (5th ed.).

Anderson TW, Rubin H., 1956. Statistical inference in factor analysis. pp. 111–150.

Andreasen, N.C., Flaum, M., 1991. Schizophrenia: The Characteristic Symptoms. Schizophr. Bull. 17, 27–49. https://doi.org/10.1093/schbul/17.1.27

Ang, Y.-S., Lockwood, P., Apps, M.A.J., Muhammed, K., Husain, M., 2017. Distinct Subtypes of Apathy Revealed by the Apathy Motivation Index. PLOS ONE 12, e0169938. https://doi.org/10.1371/journal.pone.0169938

Bach, D.R., Hulme, O., Penny, W.D., Dolan, R.J., 2011. The Known Unknowns: Neural Representation of Second-Order Uncertainty, and Ambiguity. J. Neurosci. 31, 4811–4820. https://doi.org/10.1523/JNEUROSCI.1452-10.2011

Bailey, M.R., Simpson, E.H., Balsam, P.D., 2016. Neural substrates underlying effort, time, and risk-based decision making in motivated behavior. Neurobiol. Learn. Mem. 133, 233–256. https://doi.org/10.1016/j.nlm.2016.07.015

Balderston, N.L., Vytal, K.E., O'Connell, K., Torrisi, S., Letkiewicz, A., Ernst, M., Grillon, C., 2017. Anxiety Patients Show Reduced Working Memory Related dlPFC Activation During Safety and Threat: Research Article: Anxiety Patients Show Reduced dlPFC Activity. Depress. Anxiety 34, 25–36. https://doi.org/10.1002/da.22518

Ball, T.M., Gunaydin, L.A., 2022. Measuring maladaptive avoidance: from animal models to clinical anxiety. Neuropsychopharmacology 47, 978–986. https://doi.org/10.1038/s41386-021-01263-4

Bays, P.M., Schneegans, S., Ma, W.J., Brady, T.F., 2024. Representation and computation in visual working memory. Nat. Hum. Behav. https://doi.org/10.1038/s41562-024-01871-2

Bechara, A., Damasio, A.R., Damasio, H., Anderson, S.W., 1994. Insensitivity to future consequences following damage to human prefrontal cortex. Cognition 50, 7–15. https://doi.org/10.1016/0010-0277(94)90018-3

Bechara, A., Damasio, H., Damasio, A.R., Lee, G.P., 1999. Different Contributions of the Human Amygdala and Ventromedial Prefrontal Cortex to Decision-Making. J. Neurosci. 19, 5473–5481. https://doi.org/10.1523/JNEUROSCI.19-13-05473.1999

Beck, A.T., 1961. An Inventory for Measuring Depression. Arch. Gen. Psychiatry 4, 561. https://doi.org/10.1001/archpsyc.1961.01710120031004

Behrens, T.E.J., Hunt, L.T., Woolrich, M.W., Rushworth, M.F.S., 2008. Associative learning of social value. Nature 456, 245–249. https://doi.org/10.1038/nature07538

Behrens, T.E.J., Woolrich, M.W., Walton, M.E., Rushworth, M.F.S., 2007. Learning the value of information in an uncertain world. Nat. Neurosci. 10, 1214–1221. https://doi.org/10.1038/nn1954

Berg, M., Feldmann, M., Kirchner, L., Kube, T., 2022. Oversampled and undersolved: Depressive rumination from an active inference perspective. Neurosci. Biobehav. Rev. 142, 104873. https://doi.org/10.1016/j.neubiorev.2022.104873

Bergamini, G., Cathomas, F., Auer, S., Sigrist, H., Seifritz, E., Patterson, M., Gabriel, C., Pryce, C.R., 2016a. Mouse psychosocial stress reduces motivation and cognitive function in operant reward tests: A model for reward pathology with effects of agomelatine. Eur. Neuropsychopharmacol. 26, 1448–1464. https://doi.org/10.1016/j.euroneuro.2016.06.009

Bergamini, G., Sigrist, H., Ferger, B., Singewald, N., Seifritz, E., Pryce, C.R., 2016b. Depletion of nucleus accumbens dopamine leads to impaired reward and aversion processing in mice: Relevance to motivation pathologies. Neuropharmacology 109, 306–319. https://doi.org/10.1016/j.neuropharm.2016.03.048

Berridge, K.C., Robinson, T.E., 1998. What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? Brain Res. Rev. 28, 309–369. https://doi.org/10.1016/S0165-0173(98)00019-8

Berwian, I.M., Wenzel, J.G., Collins, A.G.E., Seifritz, E., Stephan, K.E., Walter, H., Huys, Q.J.M., 2020. Computational Mechanisms of Effort and Reward Decisions in Patients With Depression and Their Association With Relapse After Antidepressant Discontinuation. JAMA Psychiatry 77, 513. https://doi.org/10.1001/jamapsychiatry.2019.4971

Bishop, S.J., Gagne, C., 2018a. Anxiety, Depression, and Decision Making: A Computational Perspective. Annu. Rev. Neurosci. 41, 371–388. https://doi.org/10.1146/annurev-neuro-080317-062007

Bishop, S.J., Gagne, C., 2018b. Anxiety, Depression, and Decision Making: A Computational Perspective. Annu. Rev. Neurosci. 41, 371–388. https://doi.org/10.1146/annurev-neuro-080317-062007

Bonnelle, V., Manohar, S., Behrens, T., Husain, M., 2015a. Individual Differences in Premotor Brain Systems Underlie Behavioral Apathy. Cereb. Cortex bhv247. https://doi.org/10.1093/cercor/bhv247

Bonnelle, V., Veromann, K.-R., Burnett Heyes, S., Lo Sterzo, E., Manohar, S., Husain, M., 2015b. Characterization of reward and effort mechanisms in apathy. J. Physiol.-Paris 109, 16–26. https://doi.org/10.1016/j.jphysparis.2014.04.002

Box, George E. P., Draper, Norman R., 1987. Empirical model-building and response surfaces. John Wiley & Sons.

Brodbeck, J., Abbott, R.A., Goodyer, I.M., Croudace, T.J., 2011. General and specific components of depression and anxiety in an adolescent population. BMC Psychiatry 11, 191. https://doi.org/10.1186/1471-244X-11-191

Browning, M., Behrens, T.E., Jocham, G., O'Reilly, J.X., Bishop, S.J., 2015. Anxious individuals have difficulty learning the causal statistics of aversive environments. Nat. Neurosci. 18, 590–596. https://doi.org/10.1038/nn.3961

Buechel, E.C., Morewedge, C.K., n.d. The (Relative and Absolute) Subjective Value of Money, in: The Psychological Science of Money. Springer, New York, NY.

Canonica, T., Zalachoras, I., 2022. Motivational disturbances in rodent models of neuropsychiatric disorders. Front. Behav. Neurosci. 16, 940672. https://doi.org/10.3389/fnbeh.2022.940672

Charpentier, C.J., Aylward, J., Roiser, J.P., Robinson, O.J., 2017. Enhanced Risk Aversion, But Not Loss Aversion, in Unmedicated Pathological Anxiety. Biol. Psychiatry 81, 1014–1022. https://doi.org/10.1016/j.biopsych.2016.12.010

Chen, J.T.-H., Lovibond, P.F., 2020. Threat appraisal and negative affect under ambiguity in generalised anxiety disorder. J. Anxiety Disord. 76, 102299. https://doi.org/10.1016/j.janxdis.2020.102299

Clark, D.A., Steer, R.A., Beck, A.T., 1994. Common and Specific Dimensions of Self-Reported Anxiety and Depression: Implications for the Cognitive and Tripartite Models. J. Abnorm. Psychol. 645–654.

Clark, L., Chamberlain, S.R., Sahakian, B.J., 2009. Neurocognitive Mechanisms in Depression: Implications for Treatment. Annu. Rev. Neurosci. 32, 57–74. https://doi.org/10.1146/annurev.neuro.31.060407.125618

Clark, L.A., Watson, D., 1991. Tripartite Model of Anxiety and Depression: Psychometric Evidence and Taxonomic Implication. J. Abnorm. Psychol. 100, 316–336.

Cléry-Melin, M.-L., Schmidt, L., Lafargue, G., Baup, N., Fossati, P., Pessiglione, M., 2011. Why Don't You Try Harder? An Investigation of Effort Production in Major Depression. PLoS ONE 6, e23178. https://doi.org/10.1371/journal.pone.0023178

Cohen, M., Jaffray, J.Y., Said, T., 1985. Individual behavior under risk and under uncertainty: An experimental study. Theory Decis. 18, 203–228. https://doi.org/10.1007/BF00134074

Collins, A.G.E., 2018. The Tortoise and the Hare: Interactions between Reinforcement Learning and Working Memory. J. Cogn. Neurosci. 30, 1422–1432. https://doi.org/10.1162/jocn_a_01238

Collins, A.G.E., Brown, J.K., Gold, J.M., Waltz, J.A., Frank, M.J., 2014. Working Memory Contributions to Reinforcement Learning Impairments in Schizophrenia. J. Neurosci. 34, 13747–13756. https://doi.org/10.1523/JNEUROSCI.0989-14.2014

Collins, A.G.E., Frank, M.J., 2012. How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis:

Working memory in reinforcement learning. Eur. J. Neurosci. 35, 1024–1035. https://doi.org/10.1111/j.1460-9568.2011.07980.x

Courtney, S.M., Ungerleider, L.G., Keil, K., Haxby, J.V., 1997. Transient and sustained activity in a distributed neural system for human working memory. Nature 386, 608–611. https://doi.org/10.1038/386608a0

Cowan, N., 1995. Attention and Memory: An Integrated Framework. Oxford University Press, New York.

Crowley, T.J., Dalwani, M.S., Mikulich-Gilbertson, S.K., Du, Y.P., Lejuez, C.W., Raymond, K.M., Banich, M.T., 2010. Risky Decisions and Their Consequences: Neural Processing by Boys with Antisocial Substance Disorder. PLoS ONE 5, e12835. https://doi.org/10.1371/journal.pone.0012835

Cryan, J.F., Valentino, R.J., Lucki, I., 2005. Assessing substrates underlying the behavioral effects of antidepressants using the modified rat forced swimming test. Neurosci. Biobehav. Rev. 29, 547–569. https://doi.org/10.1016/j.neubiorev.2005.03.008

Culbreth, A.J., Moran, E.K., Barch, D.M., 2018. Effort-based decision-making in schizophrenia. Curr. Opin. Behav. Sci., Apathy and Motivation 22, 1–6. https://doi.org/10.1016/j.cobeha.2017.12.003

Davis, M., Walker, D.L., Miles, L., Grillon, C., 2010. Phasic vs Sustained Fear in Rats and Humans: Role of the Extended Amygdala in Fear vs Anxiety. Neuropsychopharmacology 35, 105–135. https://doi.org/10.1038/npp.2009.109

Daw, N.D., Dayan, P., 2014. The algorithmic anatomy of model-based evaluation. Philos. Trans. R. Soc. B Biol. Sci. 369, 20130478. https://doi.org/10.1098/rstb.2013.0478

Daw, N.D., Niv, Y., Dayan, P., 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat. Neurosci. 8, 1704–1711. https://doi.org/10.1038/nn1560

Dayan, P., Huys, Q.J.M., 2009. Serotonin in Affective Control. Annu. Rev. Neurosci. 32, 95–126. https://doi.org/10.1146/annurev.neuro.051508.135607

D'Esposito, M., Postle, B.R., 2015. The Cognitive Neuroscience of Working Memory. Annu. Rev. Psychol. 66, 115–142. https://doi.org/10.1146/annurev-psych-010814-015031

Devlin, H.C., Johnson, S.L., Gruber, J., 2015. Feeling Good and Taking a Chance? Associations of Hypomania Risk with Cognitive and Behavioral Risk Taking. Cogn. Ther. Res. 39, 473–479. https://doi.org/10.1007/s10608-015-9679-3

Diederen, K.M.J., Fletcher, P.C., 2021. Dopamine, Prediction Error and Beyond. The Neuroscientist 27, 30–46. https://doi.org/10.1177/1073858420907591

Eckblad, M., Chapman, L.J., 1986. Development and Validation of a Scale for Hypomanic Personality. J. Abnorm. Psychol. 95, 214.

Edge, M.D., Johnson, S.L., Ng, T., Carver, C.S., 2013. Iowa gambling task performance in euthymic bipolar I disorder: A meta-analysis and empirical study. J. Affect. Disord. 150, 115–122. https://doi.org/10.1016/j.jad.2012.11.027

Ellsberg, D., 1961. Risk, Ambiguity, and the Savage Axioms. Q. J. Econ. 75.

Ernst, M., Kimes, A.S., London, E.D., Matochik, J.A., Eldreth, D., Tata, S., Contoreggi, C., Leff, M., Bolla, K., 2003. Neural Substrates of Decision Making in Adults With Attention Deficit Hyperactivity Disorder. Am. J. Psychiatry 160, 1061–1070. https://doi.org/10.1176/appi.ajp.160.6.1061

Esfand, S.M., Null, K.E., Duda, J.M., De Leeuw, J., Pizzagalli, D.A., 2024. Lifetime history of major depressive disorder is associated with decreased reward learning: Evidence from a novel online version of the probabilistic reward task. J. Affect. Disord. 350, 1007–1015. https://doi.org/10.1016/j.jad.2024.01.133

Etkin, A., Klemenhagen, K.C., Dudman, J.T., Rogan, M.T., Hen, R., Kandel, E.R., Hirsch, J., 2004. Individual Differences in Trait Anxiety Predict the Response of the Basolateral Amygdala to Unconsciously Processed Fearful Faces. Neuron 44, 1043–1055. https://doi.org/10.1016/j.neuron.2004.12.006

Faul, L., Stjepanović, D., Stivers, J.M., Stewart, G.W., Graner, J.L., Morey, R.A., LaBar, K.S., 2020. Proximal threats promote enhanced acquisition and persistence of reactive fear-learning circuits. Proc. Natl. Acad. Sci. 117, 16678–16689. https://doi.org/10.1073/pnas.2004258117

Fishbein, D.H., Eldreth, D.L., Hyde, C., Matochik, J.A., London, E.D., Contoreggi, C., Kurian, V., Kimes, A.S., Breeden, A., Grant, S., 2005. Risky decision making and the anterior cingulate cortex in abstinent drug abusers and nonusers. Cogn. Brain Res. 23, 119–136. https://doi.org/10.1016/j.cogbrainres.2004.12.010

Fleishman, E.A., 1978. Relating Individual Differences to the Dimensions of Human Tasks. Ergonomics 21, 1007–1019. https://doi.org/10.1080/00140137808931807

Forster, S., Nunez Elizalde, A.O., Castle, E., Bishop, S.J., 2015. Unraveling the Anxious Mind: Anxiety, Worry, and Frontal Engagement in Sustained Attention Versus Off-Task Processing. Cereb. Cortex 25, 609–618. https://doi.org/10.1093/cercor/bht248

Foster, D.J., Wilson, M.A., 2006. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature 440, 680–683. https://doi.org/10.1038/nature04587

Fox, C.R., Erner, C., Walters, D.J., 2015. Decision Under Risk: From the Field to the Laboratory and Back, in: Keren, G., Wu, G. (Eds.), The Wiley Blackwell Handbook of Judgment and Decision Making. Wiley, pp. 41–88. https://doi.org/10.1002/9781118468333.ch2

Frank, M.J., Doll, B.B., Oas-Terpstra, J., Moreno, F., 2009. Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. Nat. Neurosci. 12, 1062–1068. https://doi.org/10.1038/nn.2342

Frank, M.J., Moustafa, A.A., Haughey, H.M., Curran, T., Hutchison, K.E., 2007. Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. Proc. Natl. Acad. Sci. 104, 16311–16316. https://doi.org/10.1073/pnas.0706111104

Fudge, J.L., Kunishio, K., Walsh, P., Richard, C., Haber, S.N., 2002. Amygdaloid projections to ventromedial striatal subterritories in the primate. Neuroscience 110, 257–275. https://doi.org/10.1016/S0306-4522(01)00546-2

Funahashi, S., Bruce, C.J., Goldman-Rakic, P.S., 1989. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J. Neurophysiol. 61, 331–349. https://doi.org/10.1152/jn.1989.61.2.331

Fuster, J.M., Alexander, G.E., 1971. Neuron Activity Related to Short-Term Memory. Science 173, 652–654. https://doi.org/10.1126/science.173.3997.652

Gagne, C., Zika, O., Dayan, P., Bishop, S.J., 2020. Impaired adaptation of learning to contingency volatility in internalizing psychopathology. eLife 9, e61387. https://doi.org/10.7554/eLife.61387

Gard, D.E., Gard, M.G., Kring, A.M., John, O.P., 2006. Anticipatory and consummatory components of the experience of pleasure: A scale development study. J. Res. Personal. 40, 1086–1102. https://doi.org/10.1016/j.jrp.2005.11.001

Geaney, J.T., Treadway, M.T., Smillie, L.D., 2015. Trait Anticipatory Pleasure Predicts Effort Expenditure for Reward. PloS One 10, e0131357. https://doi.org/10.1371/journal.pone.0131357

Gelman, A., Hwang, J., Vehtari, A., 2014. Understanding predictive information criteria for Bayesian models. Stat. Comput. 24, 997–1016. https://doi.org/10.1007/s11222-013-9416-2

Geng, H., Wang, Y., Gu, R., Luo, Y., Xu, P., Huang, Y., Li, X., 2018. Altered brain activation and connectivity during anticipation of uncertain threat in trait anxiety. Hum. Brain Mapp. 39, 3898–3914. https://doi.org/10.1002/hbm.24219

Giorgetta, C., Grecucci, A., Zuanon, S., Perini, L., Balestrieri, M., Bonini, N., Sanfey, A.G., Brambilla, P., 2012. Reduced risk-taking behavior as a trait feature of anxiety. Emotion 12, 1373–1383. https://doi.org/10.1037/a0029119

Grillon, C., Lissek, S., Rabin, S., McDowell, D., Dvir, S., Pine, D.S., 2008. Increased Anxiety During Anticipation of Unpredictable But Not Predictable Aversive Stimuli as a Psychophysiologic Marker of Panic Disorder. Am. J. Psychiatry 165, 898–904. https://doi.org/10.1176/appi.ajp.2007.07101581

Grupe, D.W., Nitschke, J.B., 2013. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. Nat. Rev. Neurosci. 14, 488–501. https://doi.org/10.1038/nrn3524

Haile, T.M., Prat, C.S., Stocco, A., 2024. One Size Does Not Fit All: Idiographic Computational Models Reveal Individual Differences in Learning and Meta-Learning Strategies. Top. Cogn. Sci. tops.12730. https://doi.org/10.1111/tops.12730

Halahakoon, D.C., Kieslich, K., O'Driscoll, C., Nair, A., Lewis, G., Roiser, J.P., 2020. Reward-Processing Behavior in Depressed Participants Relative to Healthy Volunteers: A Systematic Review and Meta-analysis. JAMA Psychiatry 77, 1286. https://doi.org/10.1001/jamapsychiatry.2020.2139

Haluk, D.M., Wickman, K., 2010. Evaluation of study design variables and their impact on food-maintained operant responding in mice. Behav. Brain Res. 207, 394–401. https://doi.org/10.1016/j.bbr.2009.10.025

Hopwood, M., 2023. Anxiety Symptoms in Patients with Major Depressive Disorder: Commentary on Prevalence and Clinical Implications. Neurol. Ther. 12, 5–12. https://doi.org/10.1007/s40120-023-00469-6

Hosokawa, T., Kennerley, S.W., Sloan, J., Wallis, J.D., 2013. Single-Neuron Mechanisms Underlying Cost-Benefit Analysis in Frontal Cortex. J. Neurosci. 33, 17385–17397. https://doi.org/10.1523/JNEUROSCI.2221-13.2013

Hsee, C.K., Blount, S., Loewenstein, G.F., Bazerman, M.H., 1999. Preference Reversals Between Joint and Separate Evaluations of Options: A Review and Theoretical Analysis. Psychol. Bull. 125, 576–590.

Huys, Q.J., Pizzagalli, D.A., Bogdan, R., Dayan, P., 2013. Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. Biol. Mood Anxiety Disord. 3, 12. https://doi.org/10.1186/2045-5380-3-12

Huys, Q.J.M., Daw, N.D., Dayan, P., 2015. Depression: A Decision-Theoretic Analysis. Annu. Rev. Neurosci. 38, 1–23. https://doi.org/10.1146/annurev-neuro-071714-033928

Hyojin Kim, Daeyeol Lee, Shin, Y.-M., Jeanyung Chey, 2007. Impaired strategic decision making in schizophrenia. Brain Res. 1180, 90–100. https://doi.org/10.1016/j.brainres.2007.08.049

Jocham, G., Klein, T.A., Ullsperger, M., 2011. Dopamine-Mediated Reinforcement Learning Signals in the Striatum and Ventromedial Prefrontal Cortex Underlie Value-Based Choices. J. Neurosci. 31, 1606–1613. https://doi.org/10.1523/JNEUROSCI.3904-10.2011

Johnson, A., Redish, A.D., 2007. Neural Ensembles in CA3 Transiently Encode Paths Forward of the Animal at a Decision Point. J. Neurosci. 27, 12176–12189. https://doi.org/10.1523/JNEUROSCI.3761-07.2007

Johnson, A., Redish, A.D., 2005. Hippocampal replay contributes to within session learning in a temporal difference reinforcement learning model. Neural Netw. 18, 1163–1171. https://doi.org/10.1016/j.neunet.2005.08.009

Johnson, S.L., Edge, M.D., Holmes, M.K., Carver, C.S., 2012. The Behavioral Activation System and Mania. Annu. Rev. Clin. Psychol. 8, 243–267. https://doi.org/10.1146/annurev-clinpsy-032511-143148

Jöreskog, K.G., 1994. On the estimation of polychoric correlations and their asymptotic covariance matrix. Psychometrika 59, 381–389. https://doi.org/10.1007/BF02296131

Kable, J.W., Glimcher, P.W., 2009. The Neurobiology of Decision: Consensus and Controversy. Neuron 63, 733–745. https://doi.org/10.1016/j.neuron.2009.09.003

Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. Econometrica 47, 263–292.

Kennerley, S.W., Behrens, T.E.J., Wallis, J.D., 2011. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. Nat. Neurosci. 14, 1581–1589. https://doi.org/10.1038/nn.2961

Kessler, R.C., Sampson, N.A., Berglund, P., Gruber, M.J., Al-Hamzawi, A., Andrade, L., Bunting, B., Demyttenaere, K., Florescu, S., De Girolamo, G., Gureje, O., He, Y., Hu, C.,

Huang, Y., Karam, E., Kovess-Masfety, V., Lee, S., Levinson, D., Medina Mora, M.E., Moskalewicz, J., Nakamura, Y., Navarro-Mateu, F., Browne, M.A.O., Piazza, M., Posada-Villa, J., Slade, T., Ten Have, M., Torres, Y., Vilagut, G., Xavier, M., Zarkov, Z., Shahly, V., Wilcox, M.A., 2015. Anxious and non-anxious major depressive disorder in the World Health Organization World Mental Health Surveys. Epidemiol. Psychiatr. Sci. 24, 210–226. https://doi.org/10.1017/S2045796015000189

Klein, R.G., Mannuzza, S., Olazagasti, M.A.R., Roizen, E., Hutchison, J.A., Lashua, E.C., Castellanos, F.X., 2012. Clinical and Functional Outcome of Childhood Attention-Deficit/Hyperactivity Disorder 33 Years Later. Arch. Gen. Psychiatry 69, 1295. https://doi.org/10.1001/archgenpsychiatry.2012.271

Klein-Flügge, M.C., Kennerley, S.W., Friston, K., Bestmann, S., 2016. Neural Signatures of Value Comparison in Human Cingulate Cortex during Decisions Requiring an Effort-Reward Trade-off. J. Neurosci. 36, 10002–10015. https://doi.org/10.1523/JNEUROSCI.0292-16.2016

Kocher, M.G., Lahno, A.M., Trautmann, S.T., 2018. Ambiguity aversion is not universal. Eur. Econ. Rev. 101, 268–283. https://doi.org/10.1016/j.euroecorev.2017.09.016

Kotov, R., Krueger, R.F., Watson, D., Achenbach, T.M., Althoff, R.R., Bagby, R.M., Brown, T.A., Carpenter, W.T., Caspi, A., Clark, L.A., Eaton, N.R., Forbes, M.K., Forbush, K.T., Goldberg, D., Hasin, D., Hyman, S.E., Ivanova, M.Y., Lynam, D.R., Markon, K., Miller, J.D., Moffitt, T.E., Morey, L.C., Mullins-Sweatt, S.N., Ormel, J., Patrick, C.J., Regier, D.A., Rescorla, L., Ruggero, C.J., Samuel, D.B., Sellbom, M., Simms, L.J., Skodol, A.E., Slade, T., South, S.C., Tackett, J.L., Waldman, I.D., Waszczuk, M.A., Widiger, T.A., Wright, A.G.C., Zimmerman, M., 2017. The Hierarchical Taxonomy of Psychopathology (HiTOP): A dimensional alternative to traditional nosologies. J. Abnorm. Psychol. 126, 454–477. https://doi.org/10.1037/abn0000258

Kumar, P., Goer, F., Murray, L., Dillon, D.G., Beltzer, M.L., Cohen, A.L., Brooks, N.H., Pizzagalli, D.A., 2018. Impaired reward prediction error encoding and striatal-midbrain connectivity in depression. Neuropsychopharmacology 43, 1581–1588. https://doi.org/10.1038/s41386-018-0032-x

LaFreniere, L.S., Newman, M.G., 2019. Probabilistic Learning by Positive and Negative Reinforcement in Generalized Anxiety Disorder. Clin. Psychol. Sci. 7, 502–515. https://doi.org/10.1177/2167702618809366

Lavender, A., Watkins, E., 2004. Rumination and future thinking in depression. Br. J. Clin. Psychol. 43, 129–142. https://doi.org/10.1348/014466504323088015

Lawrance, E.L., Gagne, C.R., O'Reilly, J.X., Bjisterbosch, J., Bishop, S.J., 2022. The Computational and Neural Substrates of Ambiguity Avoidance in Anxiety. Comput. Psychiatry 6. https://doi.org/10.5334/cpsy.67

LeDoux, J.E., Moscarello, J., Sears, R., Campese, V., 2017. The birth, death and resurrection of avoidance: a reconceptualization of a troubled paradigm. Mol. Psychiatry 22, 24–36. https://doi.org/10.1038/mp.2016.166

Lee, K.M., Coelho, M.A., Sern, K.R., Class, M.A., Bocz, M.D., Szumlinski, K.K., 2017. Anxiolytic Effects of Buspirone and MTEP in the Porsolt Forced Swim Test. Chronic Stress 1, 247054701771298. https://doi.org/10.1177/2470547017712985

Levy, I., Schiller, D., 2021. Neural Computations of Threat. Trends Cogn. Sci. 25, 151–171. https://doi.org/10.1016/j.tics.2020.11.007

Lewis-Peacock, J.A., Postle, B.R., 2008. Temporary Activation of Long-Term Memory Supports Working Memory. J. Neurosci. 28, 8765–8771. https://doi.org/10.1523/JNEUROSCI.1953-08.2008

Li, J., Schiller, D., Schoenbaum, G., Phelps, E.A., Daw, N.D., 2011. Differential roles of human striatum and amygdala in associative learning. Nat. Neurosci. 14, 1250–1252. https://doi.org/10.1038/nn.2904

Liu, Y., Mattar, M.G., Behrens, T.E.J., Daw, N.D., Dolan, R.J., 2021. Experience replay is associated with efficient nonlocal learning. Science 372, eabf1357. https://doi.org/10.1126/science.abf1357

Lukacs, J.N., Sicilia, A.C., Jones, S., Algorta, G.P., 2021. Interactions and implications of Fuzzy-Trace theory for risk taking behaviors in bipolar disorder. J. Affect. Disord. 293, 305–313. https://doi.org/10.1016/j.jad.2021.06.035

Machina, M., Ritzberger, K., Yannelis, N., Ellsberg, D., 2011. Introduction to the symposium issue. Econ. Theory 48, 219–227.

MacLeod, A.K., Byrne, A., 1996. Anxiety, Depression, and the Anticipation of Future Positive and Negative Experiences. J Abnorm Psychol 105, 286–289.

MacLeod, A.K., Salaminiou, E., 2001. Reduced positive future-thinking in depression: Cognitive and affective factors. Cogn. Emot. 15, 99–107. https://doi.org/10.1080/02699930125776

Maner, J.K., Richey, J.A., Cromer, K., Mallott, M., Lejuez, C.W., Joiner, T.E., Schmidt, N.B., 2007. Dispositional anxiety and risk-avoidant decision-making. Personal. Individ. Differ. 42, 665–675. https://doi.org/10.1016/j.paid.2006.08.016

Mason, L., O'Sullivan, N., Montaldi, D., Bentall, R.P., El-Deredy, W., 2014. Decision-making and trait impulsivity in bipolar disorder are associated with reduced prefrontal regulation of striatal reward valuation. Brain 137, 2346–2355. https://doi.org/10.1093/brain/awu152

Master, S.L., Eckstein, M.K., Gotlieb, N., Dahl, R., Wilbrecht, L., Collins, A.G.E., 2019. Distentangling the systems contributing to changes in learning during adolescence (preprint). Neuroscience. https://doi.org/10.1101/622860

Meade, A.W., Craig, S.B., 2012. Identifying careless responses in survey data. Psychol. Methods 17, 437–455. https://doi.org/10.1037/a0028085

Menger, C., n.d. Notes on Economic Principles, in: Carl Menger Papers. David M. Rubenstein Rare Book and Manuscript Library, Duke University.

Meyer, T.J., Miller, M.L., Metzger, R.L., Borkovec, T.D., 1990. Development and validation of the penn state worry questionnaire. Behav. Res. Ther. 28, 487–495. https://doi.org/10.1016/0005-7967(90)90135-6

Montague, P., Dayan, P., Sejnowski, T., 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. J. Neurosci. 16, 1936–1947. https://doi.org/10.1523/JNEUROSCI.16-05-01936.1996

Mukherjee, D., Kable, J.W., 2014. Value-Based Decision Making in Mental Illness: A Meta-Analysis. Clin. Psychol. Sci. 2, 767–782. https://doi.org/10.1177/2167702614531580

Mukherjee, D., Lee, S., Kazinka, R., D. Satterthwaite, T., Kable, J.W., 2020. Multiple Facets of Value-Based Decision Making in Major Depressive Disorder. Sci. Rep. 10, 3415. https://doi.org/10.1038/s41598-020-60230-z

Nord, C.L., Prabhu, G., Nolte, T., Fonagy, P., Dolan, R., Moutoussis, M., 2017. Vigour in active avoidance. Sci. Rep. 7, 60. https://doi.org/10.1038/s41598-017-00127-6

Nunes, E.J., Randall, P.A., Estrada, A., Epling, B., Hart, E.E., Lee, C.A., Baqi, Y., Müller, C.E., Correa, M., Salamone, J.D., 2014. Effort-related motivational effects of the pro-inflammatory cytokine interleukin 1-beta: studies with the concurrent fixed ratio 5/ chow feeding choice task. Psychopharmacology (Berl.) 231, 727–736. https://doi.org/10.1007/s00213-013-3285-4

Orsini, C.A., Moorman, D.E., Young, J.W., Setlow, B., Floresco, S.B., 2015. Neural mechanisms regulating different forms of risk-related decision-making: Insights from animal models. Neurosci. Biobehav. Rev. 58, 147–167. https://doi.org/10.1016/j.neubiorev.2015.04.009

Padoa-Schioppa, C., Assad, J.A., 2006. Neurons in the orbitofrontal cortex encode economic value. Nature 441, 223–226. https://doi.org/10.1038/nature04676

Parr, T., Holmes, E., Friston, K.J., Pezzulo, G., 2023. Cognitive effort and active inference. Neuropsychologia 184, 108562. https://doi.org/10.1016/j.neuropsychologia.2023.108562

Paulus, M.P., Frank, L., Brown, G.G., Braff, D.L., 2003. Schizophrenia Subjects Show Intact Success-Related Neural Activation but Impaired Uncertainty Processing during Decision-Making. Neuropsychopharmacology 28, 795–806. https://doi.org/10.1038/sj.npp.1300108

Pedersen, A., Göder, R., Tomczyk, S., Ohrmann, P., 2017. Risky decision-making under risk in schizophrenia: A deliberate choice? J. Behav. Ther. Exp. Psychiatry 56, 57–64. https://doi.org/10.1016/j.jbtep.2016.08.004

Peters, J., Büchel, C., 2010. Neural representations of subjective reward value. Behav. Brain Res. 213, 135–141. https://doi.org/10.1016/j.bbr.2010.04.031

Peterson, J.C., Bourgin, D.D., Agrawal, M., Reichman, D., Griffiths, T.L., 2021. Using large-scale experiments and machine learning to discover theories of human decision-making. Science 372, 1209–1214. https://doi.org/10.1126/science.abe2629

Pittig, A., Treanor, M., LeBeau, R.T., Craske, M.G., 2018. The role of associative fear and avoidance learning in anxiety disorders: Gaps and directions for future research. Neurosci. Biobehav. Rev. 88, 117–140. https://doi.org/10.1016/j.neubiorev.2018.03.015

Pizzagalli, D.A., Goetz, E., Ostacher, M., Iosifescu, D.V., Perlis, R.H., 2008. Euthymic Patients with Bipolar Disorder Show Decreased Reward Learning in a Probabilistic Reward Task. Biol. Psychiatry 64, 162–168. https://doi.org/10.1016/j.biopsych.2007.12.001

Pollack, M.H., n.d. Comorbid Anxiety and Depression. J Clin Psychiatry.

Porsolt, R.D., Le Pichon, M., Jalfre, M., 1977. Depression: a new animal model sensitive to antidepressant treatments. Nature 266, 730–732. https://doi.org/10.1038/266730a0

R Core Team, 2018. R: A Language and Environment for Statistical Computing.

Rachman, S., 1976. The passing of the two-stage theory of fear and avoidance: Fresh possibilities. Behav. Res. Ther. 14, 125–131. https://doi.org/10.1016/0005-7967(76)90066-8

Rac-Lubashevsky, R., Cremer, A., Collins, A.G.E., Frank, M.J., Schwabe, L., 2023. Neural Index of Reinforcement Learning Predicts Improved Stimulus–Response Retention under High Working Memory Load. J. Neurosci. 43, 3131–3143. https://doi.org/10.1523/JNEUROSCI.1274-22.2023

Radloff, L.S., 1977. The CES-D Scale: A Self-Report Depression Scale for Research in the General Population. Appl. Psychol. Meas. 1, 385–401. https://doi.org/10.1177/014662167700100306

Raghunathan, R., Pham, M.T., 1999. All Negative Moods Are Not Equal: Motivational Influences of Anxiety and Sadness on Decision Making. Organ. Behav. Hum. Decis. Process. 79, 56–77. https://doi.org/10.1006/obhd.1999.2838

Ramboz, S., Oosting, R., Amara, D.A., Kung, H.F., Blier, P., Mendelsohn, M., Mann, J.J., Brunner, D., Hen, R., 1998. Serotonin receptor 1A knockout: An animal model of anxiety-related disorder. Proc. Natl. Acad. Sci. 95, 14476–14481. https://doi.org/10.1073/pnas.95.24.14476

Ramirez, F., Moscarello, J.M., LeDoux, J.E., Sears, R.M., 2015. Active Avoidance Requires a Serial Basal Amygdala to Nucleus Accumbens Shell Circuit. J. Neurosci. 35, 3470–3477. https://doi.org/10.1523/JNEUROSCI.1331-14.2015

Ramírez-Martín, A., Ramos-Martín, J., Mayoral-Cleries, F., Moreno-Küstner, B., Guzman-Parra, J., 2020. Impulsivity, decision-making and risk-taking behaviour in bipolar disorder: a systematic review and meta-analysis. Psychol. Med. 50, 2141–2153. https://doi.org/10.1017/S0033291720003086

Reise, S.P., 2012. The Rediscovery of Bifactor Measurement Models. Multivar. Behav. Res. 47, 667–696. https://doi.org/10.1080/00273171.2012.715555

Reise, S.P., Moore, T.M., Haviland, M.G., 2010. Bifactor Models and Rotations: Exploring the Extent to Which Multidimensional Data Yield Univocal Scale Scores. J. Pers. Assess. 92, 544–559. https://doi.org/10.1080/00223891.2010.496477

Rescorla, R., Wagner, A., 1972. A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement, in: Classical Conditioning II: Current Research and Theory. Black A, and Prokasy W, eds., pp. 64–99.

Rmus, M., He, M., Baribault, B., Walsh, E.G., Festa, E.K., Collins, A.G., Nassar, M.R., 2023. Age-related differences in prefrontal glutamate are associated with increased working memory decay that gives the appearance of learning deficits. eLife 12, e85243. https://doi.org/10.7554/eLife.85243

Robinson, O.J., Pike, A.C., Cornwell, B., Grillon, C., 2019. The translational neural circuitry of anxiety. J. Neurol. Neurosurg. Psychiatry 90, 1353–1360. https://doi.org/10.1136/jnnp-2019-321400

Robinson, O.J., Vytal, K., Cornwell, B.R., Grillon, C., 2013. The impact of anxiety upon cognition: perspectives from human threat of shock studies. Front. Hum. Neurosci. 7. https://doi.org/10.3389/fnhum.2013.00203

Roesch, M.R., Esber, G.R., Li, J., Daw, N.D., Schoenbaum, G., 2012. Surprise! Neural correlates of Pearce–Hall and Rescorla–Wagner coexist within the brain. Eur. J. Neurosci. 35, 1190–1200. https://doi.org/10.1111/j.1460-9568.2011.07986.x

Sartorius, N., Üstün, T.B., Lecrubier, Y., Wittchen, H.-U., 1996. Depression Comorbid with Anxiety: Results from the WHO Study on Psychological Disorders in Primary Health Care. Br. J. Psychiatry 168, 38–43. https://doi.org/10.1192/S0007125000298395

Sayalı, C., Badre, D., 2021. Neural systems underlying the learning of cognitive effort costs. Cogn. Affect. Behav. Neurosci. 21, 698–716. https://doi.org/10.3758/s13415-021-00893-x

Schacter, D.L., Addis, D.R., Buckner, R.L., 2007. Remembering the past to imagine the future: the prospective brain. Nat. Rev. Neurosci. 8, 657–661. https://doi.org/10.1038/nrn2213

Schmid, J., Leiman, J.M., 1957. The development of hierarchical factor solutions. Psychometrika 22, 53–61. https://doi.org/10.1007/BF02289209

Schwarz, G., 1978. Estimating the Dimension of a Model. Ann. Stat. 6, 461–464. https://doi.org/10.1214/aos/1176344136

Sevy, S., Burdick, K.E., Visweswaraiah, H., Abdelmessih, S., Lukin, M., Yechiam, E., Bechara, A., 2007. Iowa Gambling Task in schizophrenia: A review and new data in patients with schizophrenia and co-occurring cannabis use disorders. Schizophr. Res. 92, 74–84. https://doi.org/10.1016/j.schres.2007.01.005

Simms, L.J., Grös, D.F., Watson, D., O'Hara, M.W., 2008. Parsing the general and specific components of depression and anxiety with bifactor modeling. Depress. Anxiety 25, E34–E46. https://doi.org/10.1002/da.20432

Smith, R., Taylor, S., Stewart, J.L., Guinjoan, S.M., Ironside, M., Kirlic, N., Ekhtiari, H., White, E.J., Zheng, H., Kuplicki, R., Paulus, M.P., 2022. Slower Learning Rates from Negative Outcomes in Substance Use Disorder over a 1-Year Period and Their Potential Predictive Utility. Comput. Psychiatry 6, 117. https://doi.org/10.5334/cpsy.85

Spielberger, C. D., Gorsuch, R. L., Lushene, R., Vagg, P. R., & Jacobs, G. A., 1983. Manual for the State-Trait Anxiety Inventory. Consulting Psychologists Press, Palo Alto, CA.

Stan Development Team., 2022. RStan: the R interface to Stan.

Steer, R.A., Clark, D.A., 1995. Common and Specific Dimensions of Self-Reported Anxiety and Depression: A Replication.

Steer, R.A., Clark, D.A., Beck, A.T., Ranieri, W.F., 1998. Common and specific dimensions of self-reported anxiety and depression: the BDI-II versus the BDI-IA. Behav. Res. Ther.

Subramaniapillai, M., Mansur, R.B., Zuckerman, H., Park, C., Lee, Y., Iacobucci, M., Cao, B., Ho, R., Lin, K., Phan, L., McIntyre, R.S., 2019. Association between cognitive function and performance on effort based decision making in patients with major depressive disorder treated with Vortioxetine. Compr. Psychiatry 94, 152113. https://doi.org/10.1016/j.comppsych.2019.07.006

Sutton, R.S., Barto, A.G., 1998. Reinforcement learning: an introduction, Adaptive computation and machine learning. MIT Press, Cambridge, Mass.

Tolstoy, L., graf, 1828-1910, 1980. Anna Karenina. Oxford University Press, Oxford ; New York.

Trautmann, S.T., Van De Kuilen, G., 2015. Ambiguity Attitudes, in: Keren, G., Wu, G. (Eds.), The Wiley Blackwell Handbook of Judgment and Decision Making. Wiley, pp. 89–116. https://doi.org/10.1002/9781118468333.ch3

Treadway, M.T., Bossaller, N.A., Shelton, R.C., Zald, D.H., 2012. Effort-based decision-making in major depressive disorder: a translational model of motivational anhedonia. J. Abnorm. Psychol. 121, 553–558. https://doi.org/10.1037/a0028813

Treadway, M.T., Buckholtz, J.W., Schwartzman, A.N., Lambert, W.E., Zald, D.H., 2009. Worth the 'EEfRT'? The Effort Expenditure for Rewards Task as an Objective Measure of Motivation and Anhedonia. PLOS ONE 4, e6598. https://doi.org/10.1371/journal.pone.0006598

Treadway, M.T., Salamone, J.D., 2022. Vigor, Effort-Related Aspects of Motivation and Anhedonia, in: Anhedonia: Preclinical, Translational, and Clinical Integration, Current Topics in Behavioral Neurosciences. Springer International Publishing, Cham, pp. 326–353. https://doi.org/10.1007/978-3-031-09683-9

Treadway, M.T., Zald, D.H., 2011. Reconsidering anhedonia in depression: Lessons from translational neuroscience. Neurosci. Biobehav. Rev. 35, 537–555. https://doi.org/10.1016/j.neubiorev.2010.06.006

Tversy, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. Science 185, 1124–1131.

Van Ede, F., Nobre, A.C., 2023. Turning Attention Inside Out: How Working Memory Serves Behavior. Annu. Rev. Psychol. 74, 137–165. https://doi.org/10.1146/annurev-psych-021422-041757

Vichaya, E.G., Laumet, G., Christian, D.L., Grossberg, A.J., Estrada, D.J., Heijnen, C.J., Kavelaars, A., Dantzer, R., 2019. Motivational changes that develop in a mouse model of inflammation-induced depression are independent of indoleamine 2,3 dioxygenase. Neuropsychopharmacology 44, 364–371. https://doi.org/10.1038/s41386-018-0075-z

Von Neumann, J., & Morgenstern, O., 1944. Theory of games and economic behavior. Princeton University Press.

Wang, Yan-yu, Wang, Yi, Huang, J., Sun, X., Wang, X., Zhang, S., Zhu, G., Lui, S.S.Y., Cheung, E.F.C., Sun, H., Chan, R.C.K., 2022. Shared and distinct reward neural mechanisms among patients with schizophrenia, major depressive disorder, and bipolar

disorder: an effort-based functional imaging study. Eur. Arch. Psychiatry Clin. Neurosci. 272, 859–871. https://doi.org/10.1007/s00406-021-01376-3

Wassum, K.M., 2022. Amygdala-cortical collaboration in reward learning and decision making. eLife 11, e80926. https://doi.org/10.7554/eLife.80926

Watanabe S, Opper M, 2010. Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory. J. Mach. Learn. Res. 11, 3571–3594.

Watkins, C.J.C.H., 1989. Learning from delayed rewards. University of Cambridge, England.

Watkins, C.J.C.H., Dayan, P., 1992. Technical Note: Q-Learning. Mach. Learn. 8, 279–292.

Watson, D., Anna, L., Tellegen, A., 1988. Development and Validation of Brief Measures of Positive and Negative Affect: The PANAS Scales. J. Pers. Soc. Psychol. 54, 1063–1070.

Watson, D., Clark, L.A., 1991. The Mood and Anxiety Symptom Questionnaire. Univ. Iowa Dep. Psychol. Iowa City.

Watson, D., O'Hara, M.W., Naragon-Gainey, K., Koffel, E., Chmielewski, M., Kotov, R., Stasik, S.M., Ruggero, C.J., 2012. Development and Validation of New Anxiety and Bipolar Symptom Scales for an Expanded Version of the IDAS (the IDAS-II). Assessment 19, 399–420. https://doi.org/10.1177/1073191112449857

Weiss, J.M., Cierpial, M.A., West, C.H.K., 1998. Selective Breeding of Rats for High and Low Motor Activity in a Swim Test: Toward a New Animal Model of Depression. Pharmacol. Biochem. Behav. 61, 49–66. https://doi.org/10.1016/S0091-3057(98)00075-6

Weiss, J.M., Goodman, P.A., Losito, B.G., Corrigan, S., Charry, J.M., Bailey, W.H., 1981. Behavioral depression produced by an uncontrollable stressor: Relationship to norepinephrine, dopamine, and serotonin levels in various regions of rat brain. Brain Res. Rev. 3, 167–205. https://doi.org/10.1016/0165-0173(81)90005-9

Whitton, A.E., Treadway, M.T., Pizzagalli, D.A., 2015. Reward processing dysfunction in major depression, bipolar disorder and schizophrenia. Curr. Opin. Psychiatry 28, 7–12. https://doi.org/10.1097/YCO.0000000000000122

Wise, T., Liu, Y., Chowdhury, F., Dolan, R.J., 2021. Model-based aversive learning in humans is supported by preferential task state reactivation. Sci. Adv.

Wong, S.C.Y., Ng, M.C.M., Chan, J.K.N., Luk, M.S.K., Lui, S.S.Y., Chen, E.Y.H., Chang, W.C., 2021. Altered Risk-Taking Behavior in Early-Stage Bipolar Disorder With a History of Psychosis. Front. Psychiatry 12, 763545. https://doi.org/10.3389/fpsyt.2021.763545

World Health Organization., 2019. International statistical classification of diseases and related health problems (11th ed.).

Xu, P., Wang, K., Lu, C., Dong, L., Chen, Y., Wang, Q., Shi, Z., Yang, Y., Chen, S., Liu, X., 2017. Effects of the chronic restraint stress induced depression on reward-related learning in rats. Behav. Brain Res. 321, 185–192. https://doi.org/10.1016/j.bbr.2016.12.045

Yang, X., Huang, J., Zhu, C., Wang, Y., Cheung, E.F.C., Chan, R.C.K., Xie, G., 2014. Motivational deficits in effort-based decision making in individuals with subsyndromal

depression, first-episode and remitted depression patients. Psychiatry Res. 220, 874–882. https://doi.org/10.1016/j.psychres.2014.08.056

Zerr, C.L., Berg, J.J., Nelson, S.M., Fishell, A.K., Savalia, N.K., McDermott, K.B., 2018. Learning Efficiency: Identifying Individual Differences in Learning Rate and Retention in Healthy Adults. Psychol. Sci. 29, 1436–1450. https://doi.org/10.1177/0956797618772540

Zou, A.R., Muñoz Lopez, D.E., Johnson, S.L., Collins, A.G.E., 2022. Impulsivity Relates to Multi-Trial Choice Strategy in Probabilistic Reversal Learning. Front. Psychiatry 13, 800290. https://doi.org/10.3389/fpsyt.2022.800290