UNIVERSITY OF CALIFORNIA,
IRVINE

A Cognitive Modeling Analysis of Risk in Sequential Choice Tasks

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Cognitive Sciences

by

Maime Guan

Dissertation Committee:
Michael D. Lee, Chair
Joachim Vandekerckhove
Mark Steyvers

2019

DEDICATION

Dedicated to my parents for their love, encouragement, and endless support.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

Firstly, I would like to express my sincere gratitude to my advisors Professor Joachim Vandekerckhove and Professor Michael D. Lee for the continuous support of my Ph.D. study and related research, for their patience, motivation, and knowledge. Their guidance helped me in all the time of research and writing of this thesis. I could not have imagined having better advisors and mentors for my Ph.D. study.

Besides my advisors, I would like to thank Professor Mark Steyvers, the other member of my thesis committee, for his insightful comments and encouragement. I would also like to thank my fellow labmates for the stimulating discussions, help whenever it was needed, and for all the fun we have had in the past few years together.

Last but not least, I would like to thank my parents for supporting me throughout my time in graduate school, and my life in general. I could not have accomplished any of this without them.

This work was supported by the National Science Foundation Graduate Research Fellowship Program grant #DGE-1321846.

CURRICULUM VITAE

Maime Guan

EDUCATION

Doctor of Philosophy in Psychology **2019**
University of California, Irvine *Irvine, CA*

Master of Science in Statistics **2019**
University of California, Irvine *Irvine, CA*

Bachelor of Arts in Psychology *(magna cum laude)* **2013**
University of California, Irvine *Irvine, CA*

SKILLS

Programming: MATLAB, R, Python, SQL, SPSS, JAGS
Statistical Methods: Bayesian inference, latent variable models, experimental design, hypothesis testing, hierarchical models
Additional languages: Chinese

EXPERIENCE

Graduate Researcher **2013–2019**
Cognition and Individual Differences Lab
University of California, Irvine *Irvine, CA*

Graduate Researcher **2013–2019**
Memory and Decision-Making Lab
University of California, Irvine *Irvine, CA*

Data Science Fellow **2018**
Insight Data Science *San Francisco, CA*

Data Science Intern **2017**
Rocket Software *Austin, TX*

Research Assistant **2011–2012**
Chubb-Wright Lab
University of California, Irvine *Irvine, CA*

FELLOWSHIPS AND AWARDS

National Science Foundation Graduate Research Fellowship	**2014-2019**
Associate Dean Fellowship Award	**2019**
Women of MathPsych Travel and Networking Award	**2015**
Society for Mathematical Psychology Conference Participation Award	**2014**
Society for Mathematical Psychology Conference Participation Award	**2013**
UC Irvine Summer Undergraduate Research Program Grant	**2012**

TEACHING EXPERIENCE

Teaching Assistant
University of California, Irvine

- Research in Experimental Psychology
- Advanced Experimental Psychology
- Principles of Learning Theory
- Psychology Fundamentals

ORGANIZATIONAL ACTIVITIES

- Dean's Fellowship Advisory Committee as a reviewer for Social Sciences graduate research grant proposals. 2014-2018 at UC Irvine.
- Panelist on the Social Sciences Ambassador's Club workshop for undergraduate students in the Social Sciences applying for graduate school. January 25, 2016 at UC Irvine.
- Panelist on the NSF Graduate Research Fellowship Program Workshop for PhD students in the Social Sciences applying for NSF GRFP. 2014 and 2015 at UC Irvine.
- Panelist on the Student Academic Enrichment Program for first-generation, low-income, and under-represented undergraduate students in the Social Sciences applying for graduate school. July 9, 2015 at UC Irvine.
- On the organization committee for the Annual Conference of the Society for Mathematical Psychology. July 17-22, 2015 in Newport Beach, California.
- On the organizational committee, with Michael Lee and Eric-Jan Wagenmakers, for the Bayesian Cognitive Modeling Workshop. August 12-16, 2013 in Amsterdam, the Netherlands.

PUBLICATIONS

Guan, M., & Lee, M. D. (2017). The effect of goals and environments on human performance in optimal stopping problems. *Decision, 5*(4), 339–361.

Guan, M., & Vandekerckhove, J. (2016). A Bayesian approach to mitigation of publication bias. *Psychonomic bulletin & review, 23*(1), 74-86.

Van Elk, M., Matzke, D., Gronau, Q., Guang, M., Vandekerckhove, J., & Wagenmakers, E. J. (2015). Meta-analyses are no substitute for registered replications: A skeptical perspective on religious priming. *Frontiers in psychology, 6,* 1365.

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A Hierarchical Cognitive Threshold Model of Human Decision Making on Different Length Optimal Stopping Problems. *In Proceedings of the 37th Annual Meeting of the Cognitive Science Society*, Pasadena, California.

Guan, M., Lee, M., & Silva, A. (2014). Threshold models of human decision making on optimal stopping problems in different environments. *In Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, Quebec City, Canada.

Vandekerckhove, J., Guan, M., & Styrcula, S. A. (2013). The consistency test may be too weak to be useful: Its systematic application would not improve effect size estimation in meta-analyses. *Journal of Mathematical Psychology, 57*(5), 170-173.

Chubb, C., Dickson, C. A., Dean, T., Fagan, C., Mann, D. S., Wright, C. E., Guan, M., Silva, A. E., Gregersen, P. K., & Kowalsky, E. (2013). Bimodal distribution of performance in discriminating major/minor modes. *The Journal of the Acoustical Society of America, 134*(4), 3067-3078.

PRESENTATIONS

Lee, M.D., Guan, M., Danileiko, I. (2018). *Incorporating cognitive modeling into sports predictions.* Oral presentation at the 51st Annual Conference of the Society for Mathematical Psychology, Madison, Wisconsin.

Guan, M. & Lee, M. D. (2016). *The effects of goals and environments on human performance in optimal stopping problems.* Oral presentation at the 49th Annual Conference of the Society for Mathematical Psychology, New Brunswick, New Jersey.

Guan, M. Lee, M. D., & Vandekerckhove, J. (2015). *A hierarchical Bayesian threshold model of human decision making on different length optimal stopping problems.* Poster presentation at the 37th Annual Conference of the Cognitive Science Society, Pasadena, California.

Guan, M. (2015). *The datafication of HR: benefits from cognitive modeling.* Oral presentation at the Applications of Mathematical Psychology in Industry meeting, Newport Beach, California.

Guan, M. Lee, M. D., & Vandekerckhove, J. (2015). *A cognitive latent variable model of decision making on optimal stopping problems.* Oral presentation at the 48th Annual Conference of the Society for Mathematical Psychology, Newport Beach, California.

Guan, M. & Vandekerckhove, J. (2015). *A Bayesian approach to mitigation of publication bias.* Oral presentation at the 27th Annual Convention of the Association for Psychological Science, New York City, New York.

Guan, M. & Vandekerckhove, J. (2014). *Identification of publication bias through models of publication behavior.* Poster presentation at the 55th Annual Conference of the Psychonomic Society, Long Beach, California.

Guan, M., Lee, M.D., & Silva, A. (2014). *Threshold models of human decision making on optimal stopping problems in different environments.* Oral presentation at the 36th Annual Conference of the Cognitive Science Society, Quebec City, Canada.

Guan, M. & Vandekerckhove, J. (2014). *Identification of publication bias through models of publication behavior.* Oral presentation at the 47th Annual Conference of the Society for Mathematical Psychology, Quebec City, Canada.

Guan, M., Silva, A., & Lee, M.D. (2013). *Threshold models of human decision-making on optimal stopping problems in different environments.* Poster presentation at the 46th Annual Conference of the Society for Mathematical Psychology, Potsdam, Germany.

Guan, M. (2013). *Observed power in a large sample of published articles in psychology.* Oral presentation at the 20th Annual University of California, Irvine Undergraduate Research Symposium, Irvine, California.

Guan, M. & Vandekerckhove, J. (2012). *Observed power in a large sample of published articles in psychology.* Poster presentation at the 45th Annual Conference of the Society for Mathematical Psychology, Columbus, Ohio.

ABSTRACT OF THE DISSERTATION

A Cognitive Modeling Analysis of Risk in Sequential Choice Tasks

By

Maime Guan

Doctor of Philosophy in Cognitive Sciences

University of California, Irvine, 2019

Michael D. Lee, Chair

There exists a variety of instruments that assess risk propensity, or an individual's intrinsic tendency to be risk seeking. This thesis looks at four widely-studied cognitive tasks (the optimal stopping problem, the Balloon Analogue Risk Task, bandit problems, and a preferential choice gambling task) and three commonly used risk questionnaires (Risk Propensity Scale, Risk Taking Index, and Domain-Specific Risk-Taking Scale). Although these decision-making tasks and risk questionnaires have been studied extensively in isolation, there has been less research comparing measures of risk propensity across them. The motivation for examining the relationships between the tasks is that if an individual has a fundamental propensity to take risks, then this trait should be reflected in various questionnaires and cognitive tasks in which behavior is sensitive to risk. Within-subjects data was collected through Amazon Mechanical Turk from 56 participants. As measures of risk from the decision-making tasks, four cognitive models are implemented in which there are psychological variables that can be interpreted as risk propensity. Modeling results, based on Bayesian inferences about parameters and their correlations, show that people's risk behavior is consistent within tasks, but there is less evidence that the way people manage risk in each domain generalizes across tasks and questionnaires.

Chapter 1

Introduction

Humans are faced with risky decision-making situations every day, from financial investments to choosing dating partners. We are constantly evaluating the potential losses and gains in taking a particular action in a certain scenario. An individual's intrinsic tendency to be risk seeking, known as *risk propensity*, has been found to have a dominant influence on their behavior in those situations (Sitkin and Weingart, 1995; Stewart Jr and Roth, 2001; Dunlop and Romer, 2010; Mishra et al., 2010; Lejuez et al., 2004).

In the field of psychology, there exists a variety of instruments that assess risk propensity, many of which are in the form of self-report surveys using Likert scales, but also include cognitive decision-making tasks. Although most tools demonstrate adequate reliability (Harrison et al., 2005), previous research has not compared the measures of risk propensity across some of the most commonly used risk surveys and widely used risky decision-making tasks. If an individual has a fundamental propensity to take risks, then this trait should be reflected in various surveys and cognitive tasks in which behavior is sensitive to risk. The goal of this work is to examine the relationship between the measures of risk propensity from multiple surveys and cognitive tasks.

Among some commonly used cognitive tasks are the optimal stopping problem (Guan et al., 2015; Guan and Lee, 2017; Lee, 2006; Seale and Rapoport, 2000), bandit problems (Lee et al., 2011; Steyvers et al., 2009; Zhang and Lee, 2010b), the Balloon Analogue Risk Task (Lejuez et al., 2002, 2003), and preferential choice gamble tasks (De Martino et al., 2006; Russo and Dosher, 1983; Rieskamp et al., 2006). For all four of these cognitive tasks involving decision making under risk and uncertainty, there are existing cognitive models in the literature with interpretable psychological variables that can represent risk propensity. Among some of the most widely used risk assessment surveys are the Risk Propensity Scale (Meertens and Lion, 2008), Risk Taking Index (Nicholson et al., 2005), and Domain-Specific Risk-Taking scale Blais and Weber (2006). These three surveys have been used in a variety of contexts in the literature and have been found to be reliable in measuring people's risk propensity.

This thesis first considers how the four cognitive tasks measure risk propensity through cognitive models of human behavior. Then, the relationship between the risk measures from the four cognitive tasks and three risk surveys will be explored in a Bayesian framework through correlation models that account for the uncertainty in measurement. All three risk surveys are designed to measure risk propensity and have been shown to be reliable assessments, so positive correlations are expected between these surveys. However, although the four cognitive tasks also involve risky decision making, it is not obvious that risk measures from these tasks will correlate with the survey measures or with each other.

In order to study individual differences and the relationship between risk measures across the four cognitive tasks and three surveys, within-subjects was collected so that each individual completed all tasks and surveys. To my knowledge, there does not yet exist a within-subjects data set that spans these four risky decision-making tasks and risk surveys. Moreover, this is the first time that individual differences between participants as well as differences in the nature of tasks and what they measure will be studied for this set of cognitive tasks and surveys.

The following section of this chapter describes the entire experiment and how the data were collected. Then, the next four chapters provide background and details of each of the four risky decision-making tasks and how cognitive models are used to infer people's risk propensity. After the cognitive tasks, the sixth chapter presents the three risk surveys used in this experiment. The seventh chapter looks at the relationships between risk measures across all cognitive tasks and surveys in a correlation analyses implemented in a Bayesian framework. Lastly, the final chapter concludes with a discussion of the implications of these findings and future directions.

1.1 Overview of Experiments

1.1.1 Participants

A total of 56 participants were recruited through Amazon Mechanical Turk. Each participant was given monetary compensation of $8.00 for their time in completing this experiment. All participants clicked to confirm their agreement with an informed consent form before beginning the experiment. There were 37 male participants and 19 female participants, with ages ranging from 20 to 61 ($M = 36.4$, $SD = 11.6$).

1.1.2 Tasks and Surveys

This experiment consists of four cognitive tasks:

1. Optimal Stopping
2. Balloon Analogue Risk Task (BART)
3. Bandit Problems
4. Gambling Task

and three self-report surveys:

1. Risk Propensity Scale (RPS)
2. Risk Taking Index (RTI)
3. Domain-Specific Risk-Taking scale (DOSPERT)

Each of the cognitive tasks takes about 20-30 minutes to complete. The RPS and RTI take about 5 minutes each, while the DOSPERT takes about 10-15 minutes. Each participant completed all tasks and surveys. Because the entire experiment takes about 2 hours to complete, the experiment was split up into two parts of about 1 hour each. Each part included two cognitive tasks and either the RPS and RTI or the DOSPERT. The RPS and RTI were completed in the same part because these two surveys are much shorter than the DOSPERT. The order of tasks and surveys were randomized across participants, but each participant did the same sets of problems within tasks and answered the same questions in the surveys.

Upon completing Part 1 of the experiment, the participants were given a unique code. This code was necessary for them to come back for Part 2 and receive credit. All participants who completed Part 1 returned and completed Part 2; they were told that they would only receive credit if they complete both parts. Participants were also encouraged to take a break between Part 1 and Part 2, but the only requirement was that they must complete both parts within six days.

The next four chapters go into further detail of the experiment, cognitive model, and modeling results for each of the four cognitive tasks. The chapter following the last cognitive provides details of the three risk surveys and some empirical results of the risk propensity scores. Then, the measures across tasks and surveys will be compared based on Bayesian inferences about their correlations. Finally, implications of these results and further directions of this work will be discussed.

Chapter 2

Optimal Stopping

The optimal stopping problem, also known as secretary problems, is a decision-making task in which people must choose the highest value out of a sequence of numbers, under the constraint that a number can only be chosen when it is presented (Ferguson, 1989; Gilbert and Mosteller, 1966). Optimal stopping problems are interesting for understanding human decision making under risk and uncertainty, because they have two features found in many real-world decision-making settings. The first feature is that there is *no going back*. Oftentimes in real life, it is very difficult or near impossible to return to an earlier option once it is declined. For example, in job searching, it is almost impossible to come back to a job offer once you reject it. The second feature is *only the best will do*. In some real-world situations, there is only one best option ad any other option is completely and equally useless. For example, trying to identify the correct suspect in a police lineup requires that you choose the true guilty suspect. Identifying someone who looks almost identical to the guilty suspect is just as wrong as identifying any other innocent person.

Human decision making on optimal stopping problems has been widely studied in a variety of contexts using a number of different versions of the task. One common version of the optimal stopping problem is the classic rank order version of the problem, in which only the rank of the current option relative to the options already seen is presented (Seale and Rapoport, 1997, 2000; Bearden et al., 2006). Other studies have used the full-information version of the task, in which the actual values of the alternatives are presented (Lee, 2006; Guan et al., 2014, 2015; Shu, 2008). For both of these versions there is a known optimal solution process to which human performance can be compared (Ferguson, 1989; Gilbert and Mosteller, 1966). In the full-information versions that we consider in this paper, the optimal solution is to choose the first number that is both currently maximal or minimal (depending on the goal), and above or below a threshold that depends upon the current position in the sequence.

Furthermore, most existing studies of full-information optimal stopping problems use numbers that are drawn from a uniform distribution (Campbell and Lee, 2006; Kogut, 1990; Lee, 2006; Sonnemans, 1998). An interesting variation, however, involves drawing numbers from non-uniform distributions. Using other distributions creates new *environments* for optimal stopping problems. The psychological motivation for this is that some decision environments are plentiful with desirable alternatives, while others are scarce. Finding a job in a big city, where there are relatively more job opportunities, is different from finding a job in a small town. Intuitively, optimal decision making involves setting higher standards in plentiful environments, and lower standards in scarce environments.

Previous work (Guan and Lee, 2014; Lee, 2006; Goldstein et al., 2007) found evidence that people use a series of thresholds to make decisions, and that there are large individual differences in thresholds. One can imagine that an individual's risk propensity might help determine the thresholds that they use in the optimal stopping problem. Someone who is risk seeking might have a relatively higher threshold, compared to optimal, for accepting a number at a particular position in the sequence. Alternatively, someone who is risk averse might have a relatively lower threshold compared to optimal for accepting a number at a particular position. The goal of this work is to use the optimal stopping problem to estimate the risk propensity of individuals by estimating the thresholds that they use to make decisions in the task.

This work examines decision making on optimal stopping problems with different *lengths* and *environments*. Specifically, we manipulate the length of each problem to be either four or eight number sequences, and the environment of each problem to be either plentiful (relatively more high numbers) or neutral (uniform distribution over all numbers). If risk propensity is a psychological component that determines the thresholds in which people use, then we should expect behavior to be similar between the four tasks. For instance, participants who use thresholds higher than optimal in the neutral length four task should also use thresholds higher than optimal in the plentiful length eight task. A recruiter who is generally picky and willing to hold out until the perfect candidate comes along will have relatively high thresholds for job applicants whether they are choosing from 4 applicants or 8.

In the next section, details of the optimal stopping experiment will be provided. Next, the Bayesian cognitive model for inferring the thresholds that people use will be explained. Lastly, modeling results will be provided for the optimal stopping task.

2.1 Experiment

In the optimal stopping portion of the study, all Amazon Mechanical Turk participants completed four sets of optimal stopping problems: 1) length four problems in the neutral environment, 2) length four problems in the plentiful environment, 3) length eight problems in the neutral environment, and 4) length eight problems in the plentiful environment. In the neutral environments, numbers were generated from a $\text{Uniform}\big(0, 100\big)$ distribution. In the plentiful environments, numbers were generated from a $\text{Beta}\big(4, 2\big)$ distribution. Consequently, the plentiful environments contain more numbers that are relatively larger, while the neutral environment had numbers that were relatively smaller. All participants completed the same 40 problems within each set, while the order of problems within each set was randomized across participants.

Participants were instructed to pick the heaviest cartoon cat out of a sequence of cartoon cats sitting on scales, with their weight in pounds (ranging from 0 to 100) written on the scale they are sitting on. They were told (1) the length of the sequence, (2) that a value could only be chosen when it is presented, (3) that any value that is not the maximum is completely and equally incorrect as the others, and (4) that the last value must be chosen if no values were chosen in all previous positions. Participants indicated whether or not they chose each presented value by pressing either a "select" or "pass" button. The numbers that participants rejected in a sequence were not shown once the next number in the sequence was displayed. In addition, the numbers in the sequence after the one the participant chose were never presented. After each problem, participants were provided with feedback on their response.

Figure 2.1 shows a screenshot of the online experiment. The top of the interface shows how many problems they correctly chose the heaviest cat so far. The bottom of the interface shows which cat in the sequence of cats they are currently on. Feedback is given on the screen after each problem, indicating whether the participant correctly chose the heaviest cat in the sequence with "Correct!', or incorrectly chose any other cat in the sequence with "Wrong".

The following section details the Bayesian cognitive model for inferring the thresholds that people use in optimal stopping problems. The cognitive model has risk propensity parameters that quantify how people's thresholds deviate from optimality.

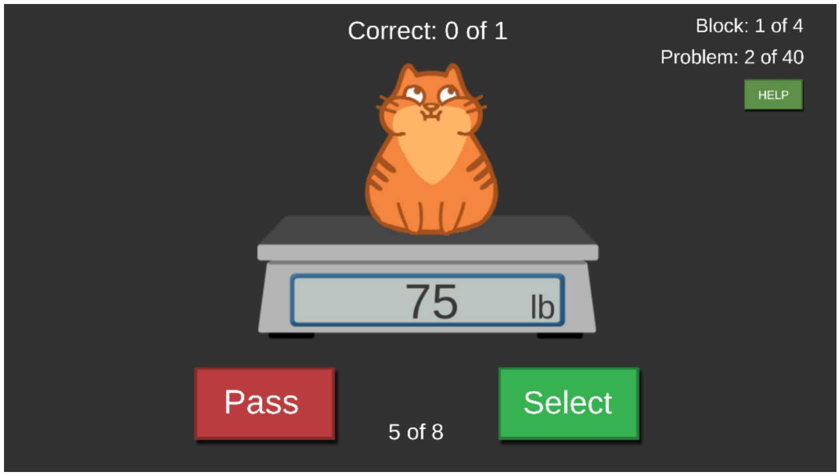

Figure 2.1: A screenshot of the online experiment in the Optimal Stopping task.

2.2 Bias-From-Optimal Model

The goal of the cognitive model is to infer the shape of the thresholds that people use in optimal stopping problems, with parameters that capture each individual's risk propensity. Therefore, the Bias-From-Optimal (BFO) model models each participant's thresholds in terms of how they deviate from the optimal threshold.

We denote the optimal thresholds as $\tilde{\tau}_1, \ldots, \tilde{\tau}_m$ for a problem of length m (Gilbert and Mosteller, 1966, Table 2). Naturally, the last threshold in the sequence must be 0 since the last value must be chosen if nothing else was chosen in the first $m-1$ positions. The ith participant's thresholds depend on a parameter $\beta_i^m \sim \text{Gaussian}\big(0,1\big)$ that determines how far above or below their threshold is from optimal, and a parameter $\gamma_i^m \sim \text{Gaussian}\big(0,1\big)$ that determines how much their bias increases or decreases as the sequence progresses. Formally, the ith participant's thresholds for a problem of length m is

$$\tau_{ik}^m = \Phi\big(\Phi^{-1}(\tilde{\tau}_k^m) + \beta_i^m + \frac{k}{m}\gamma_i^m\big) \tag{2.1}$$

for the first $m-1$ positions, and $\tau_{im}^{m}=0$ for the last. The link functions Φ and Φ^{-1} are the Gaussian CDF and inverse CDF, respectively. According to the threshold choice model, the probability that the ith participant will choose the value they are presented in the kth position on their jth problem is

$$\theta_{ijk}^{m}=\begin{cases}\alpha_i^m & \text{if } v_{ijk}^m > \tau_{ik}^m \ \&\ v_{ijk}^m = \max\left\{v_{ij1}^m,\ldots,v_{ijk}^m\right\} \\ \frac{1-\alpha_i^m}{m} & \text{otherwise}\end{cases} \tag{2.2}$$

for the first m positions and $\theta_{ijm}^{m}=1-\sum_{k=1}^{m-1}\theta_{ijk}^{m}$ for the last position, where $\alpha_i^m \sim \text{Uniform}\big(0,1\big)$ is the individual-level "accuracy of execution" parameter that describes how often the deterministic threshold model is followed (Guan and Lee, 2014). The threshold model is completed by the observed data being distributed according to the choice probabilities, so that

$$y_{ij}^{m} \sim \text{Categorical}\big(\theta_{ij1}^{m},\ldots,\theta_{ij5}^{m}\big). \tag{2.3}$$

Figure 2.2 shows how the shape of threshold functions changes with different values of β and γ, shown along with the optimal decision threshold for a problem of length four in a neutral environment. The optimal threshold corresponds to the case with $\beta=0$ and $\gamma=0$, and is shown in bold. The β parameter represents a shifting bias from this optimal curve, with positive values resulting in thresholds that are above optimal, and negative values resulting in thresholds that are below optimal. The γ parameter represents how quickly thresholds are reduced throughout the problem sequence, relative to the optimal rate of reduction. Positive values of γ produce thresholds that drop too slowly, while negative values of γ produce thresholds that drop too quickly. Priors are placed on the two risk parameters and consistency parameter for each participant: $\gamma,\beta \sim \text{Normal}\big(0,1\big), \alpha \sim \text{Uniform}\big(0,1\big)$.

The BFO model is implemented as a graphical model using JAGS (Plummer, 2003), software that facilitates MCMC-based computational Bayesian inference (Lee and Wagenmakers,

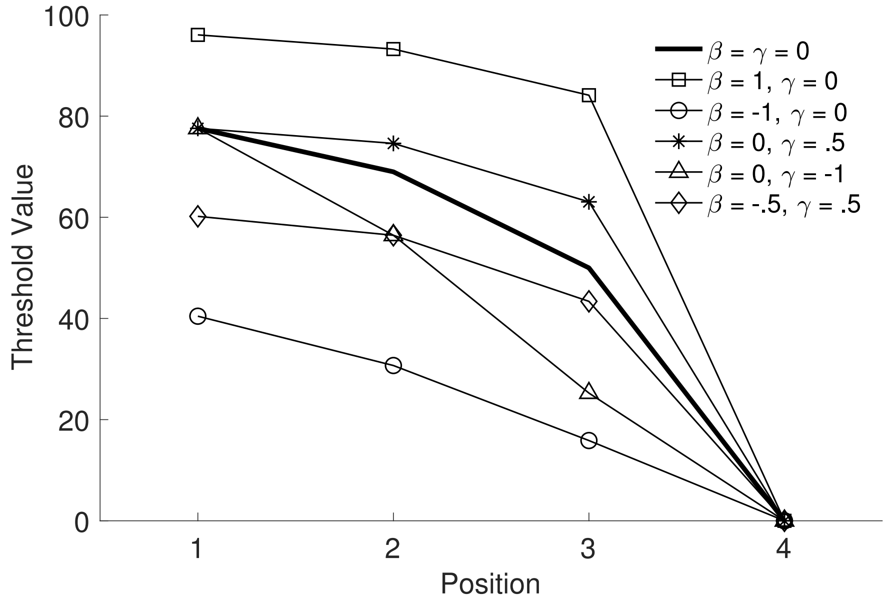

Figure 2.2: The behavior of the bias-from-optimal threshold model under different parameterizations.

2013). All modeling results are based on four chains of 1000 samples each, collected after 2000 discarded burn-in samples. The chains were verified for convergence using the standard $\hat{R}$ statistic (Brooks and Gelman, 1997). Figure 2.3 shows the BFO model in the graphical modeling formalism. The latent parameters corresponding to the thresholds $\boldsymbol{\tau}_i^m$ and accuracy of execution α_i^m are represented by unshaded and circular nodes, since they are unobserved and continuous. The values $\boldsymbol{v}_j^m$ presented on the jth problem, standardized to lie between 0 and 1, instead of the 0 to 100 scale used in the experiment, are shown as a shaded node, since they are observed and continuous. Together, the parameters and problem values determine the probabilities $\boldsymbol{\theta}_{ij}^m$ for each possible decision, shown as a double-bordered node since it is a deterministic function of its parents in the graphical model. The decision y_{ij}^m is shown as a shaded and square node, since it is observed and discrete. Encompassing plates for participants and problems indicate independent replications of the graph structure in the model.

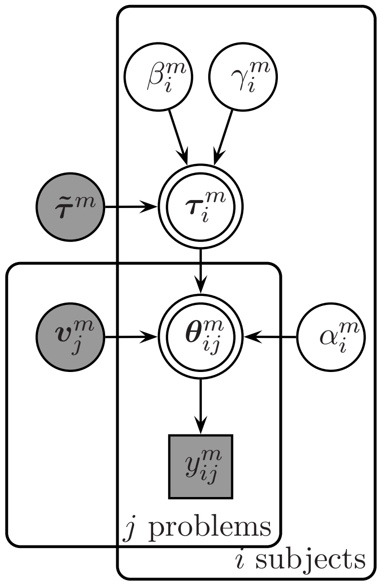

$$
\begin{aligned}
\beta_i^m &\sim \text{Gaussian}(0,1)\\
\gamma_i &\sim \text{Gaussian}(0,1)\\
\alpha_i &\sim \text{Uniform}(0,1)\\
\tau_{ik}^m &= \Phi\left(\Phi^{-1}\left(\tilde{\tau}_k^m\right) + \beta_i + \tfrac{k}{m}\gamma_i\right)\\
\tau_{im}^m &= 0\\
\theta_{ijk}^m &= \begin{cases} \alpha_i & \text{if } v_{ijk}^m > \tau_{ik}^m \ \&\ v_{ijk}^m = \max\left\{v_{ij1}^m, \ldots, v_{ijk}^m\right\} \\ \frac{1-\alpha_i}{m} & \text{otherwise} \end{cases}\\
\theta_{ijm}^m &= 1 - \textstyle\sum_{k=1}^{m-1} \theta_{ijk}^m\\
y_{ij}^m &\sim \text{Categorical}(\boldsymbol{\theta}_{ij}^m)
\end{aligned}
$$

Figure 2.3: Graphical model for the BFO threshold model applied to all four conditions of the optimal stopping task.

2.3 Results

2.3.1 Removing Contaminants

A contaminant model was implemented in order to identify a potential subset of participants who are not performing the task properly for any reason. For example, some participants may not have understood the task properly, or they may have wanted to rush through the experiment to complete the task as quickly as possible. A method to allow for the possibility of contaminant participants is to provide an alternative account for how they made decisions in the task, and use model selection techniques to determine the evidence for each of these from the behavioral data.

To determine whether participants were behaving at random in the optimal stopping problem, a simple *guessing* model was implemented as a baseline. The guessing model does not have any free parameters. According to the guessing model, the probability that the ith participant will choose the value they are presented with in the kth position on their jth problem for a problem of length m is

$$\theta_{ijk}^{m} = \frac{1}{m} \tag{2.4}$$

This guessing model was compared to the BFO model using Bayes factors (Kass and Raftery, 1995), a standard Bayesian approach to model comparison that simultaneously controls for goodness-of-fit and model complexity. Bayes factors for each participant were estimated using a latent mixture procedure based on model-indicator variables (Wagenmakers et al., 2010). For all participants in the optimal stopping problem, there was overwhelming evidence for the BFO model over the guessing model.

Furthermore, participants were also considered contaminants if they simply picked the first number in the sequence repeatedly across the majority of the problems just to complete the task faster (no matter the value of the number). Three contaminant participants were subsequently removed from the remainder of the analyses.

2.3.2 Performance

Next, the performance of participants in terms of the proportion of problems where they correctly chose the maximum was examined across blocks of problems. Figure 2.4 shows the results of measuring performance across problems. The problems were split into 4 blocks of 10 problems each. In the two length 4 conditions, the mean participant performance was between about 0.5 and 0.6. In the two length 8 conditions, the mean participant performance was between about 0.3 and 0.5. Consistent with previous research (Guan and Lee, 2014, 2017; Lee, 2006), there does not appear to be any learning effects. Participant performance for all four conditions do improve across problems. Within the same length problems, there also

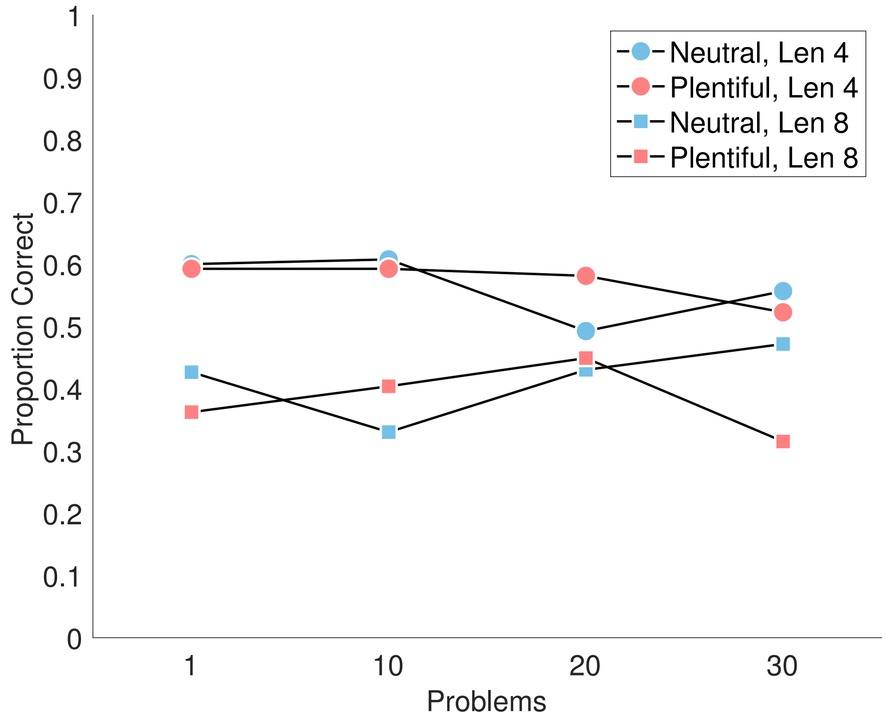

Figure 2.4: Learning curves of participants across blocks of problems.

doesn't seem to be differences in performance between neutral and plentiful conditions.

2.3.3 Inferred Thresholds

The BFO model allows us to quantify people's risk propensity by inferring the thresholds that they use in the optimal stopping task. Someone who is risk seeking might have a threshold that is higher and drops slower than optimal, while someone who is risk-averse might have a threshold that is lower and drops faster than the optimal. Figure 2.5 shows the marginal posterior expectations for all the inferred thresholds under all four conditions for all of the participants. The optimal decision threshold in each condition is also shown as a solid black line.

It is clear that participants are sensitive to both length of the problem and the underlying environment of the stimuli. This is evident in that thresholds in the plentiful environments

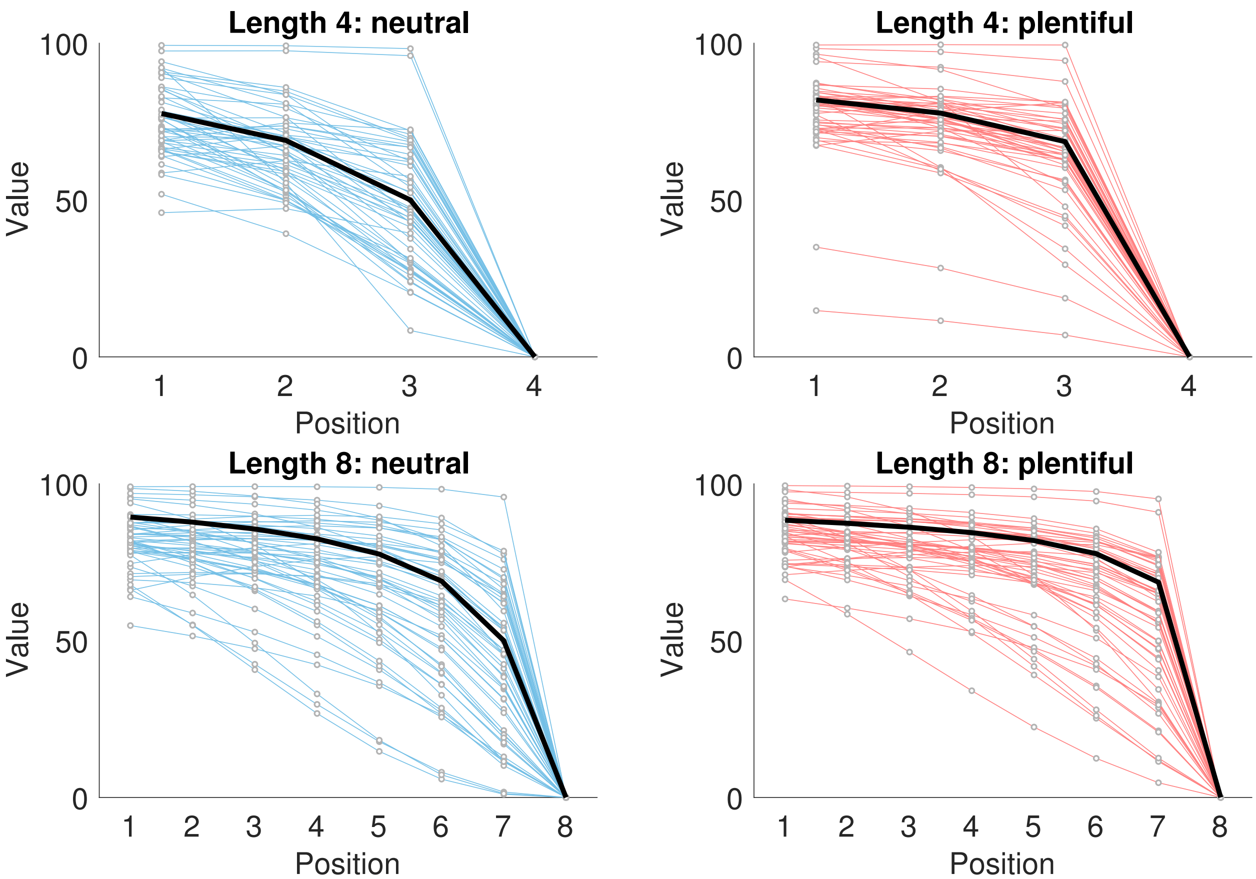

Figure 2.5: The inferred thresholds for all participants across the four conditions in the optimal stopping task. The top-left panel shows the inferred thresholds in the length 4 neutral environment, the top-right panel shows the thresholds in the length 4 plentiful environment, the bottom-left panel shows the inferred thresholds in the length 8 environment, and the bottom-right panel shows the inferred thresholds in the length 8 plentiful environment.

(right panels) are relatively lower than the thresholds in the neutral environments (left panels). Furthermore, thresholds in the length 8 conditions (bottom panels) hold out longer than thresholds in the length 4 conditions (top panels). However, it is also important to note that there are individual differences in the thresholds that people use across participants within all four conditions. Lastly, it appears that participants in the length 8 conditions tend to use thresholds that are lower than optimal, especially in the plentiful environment.

The way that participants' thresholds are generated in each of the conditions are through the two risk parameters: β, a measure of how far above or below the threshold is shifted from

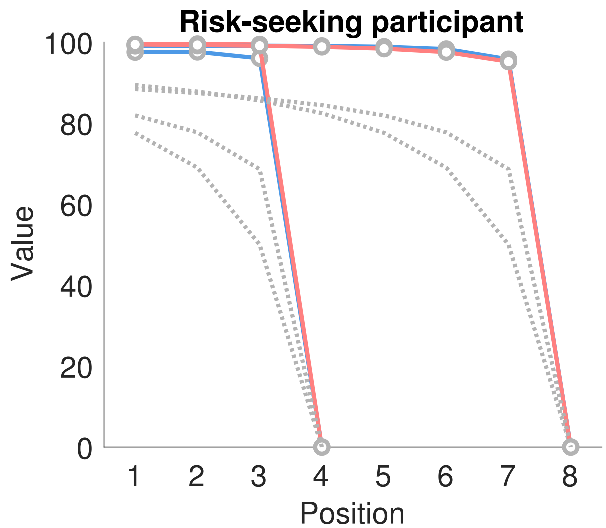

Figure 2.6: The inferred thresholds for two participants across the four conditions in the optimal stopping task. The left panel shows the inferred thresholds for the relatively risk-seeking participant, and the right panel shows the inferred thresholds for the relatively risk-averse participant. The thresholds in blue are the thresholds in the neutral condition, while the thresholds in red are the thresholds in the plentiful condition.

optimal, and γ a measure of how quickly the threshold drops relative to the optimal rate of dropping throughout the sequence. To illustrate how these parameters capture the individual differences of participants in the optimal stopping problem, Figure 2.6 shows the inferred thresholds under all four conditions for two representative participants. These participants were chosen because they are on the opposite ends of the risk propensity spectrum seen across all participants in the data set.

The participant in the left panel can be seen as risk seeking, because their thresholds for all four conditions are much higher than optimal. No matter whether they are choosing from four or eight alternatives, or whether they are in a neutral or plentiful environment, this participant is not willing to take low values at any point in the sequence. This participant starts their threshold high and is not willing to lower them no matter what. This risk-seeking behavior is quantified in the risk parameter values, which are positive with relatively large magnitude. Conversely, the participant in the right panel can be seen as risk averse, because their thresholds are much lower than optimal in all four conditions. Regardless of the length of the problem and the environment they are in, this participant starts their threshold low and drops it very quickly throughout the sequence. This risk-averse behavior is quantified in the risk parameter values, which are negative with relatively large magnitude.

Figure 2.7 summarizes the individual differences across all participants, using the means of the posterior expectations of β and γ across all four conditions, for each participant. The posterior expectations of the β and γ risk parameters are shown jointly as a scatterplot in the center panel, with their marginal distributions shown as histograms on the bottom and left panels. The dotted lines represent where β and γ are equal to 0. Where the dotted lines meet in the center represents the optimal threshold, where there is no bias in either direction for both of the risk parameters. It is clear that there is a wide range of individual differences in risk propensity, because all four quadrants around optimality are populated. β and γ values for all participants range from positive to negative biases.

The risk-seeking participant from Figure 2.6 is labeled in the scatterplot as Participant 1. This participant is in the quadrant where both β and γ are positive. The risk-averse participant from Figure 2.6 is labeled as Participant 2. This participant is in the quadrant where both β and γ are negative. From inferring the thresholds that people use in the optimal stopping problem, the BFO model is able to capture their risk propensity with these two risk parameters.

2.3.4 Descriptive Adequacy

To examine the ability of the model to fit the behavioral data, a standard Bayesian approach based on the posterior predictive distribution was used (Gelman et al., 2004). The posterior predictive is the distribution of choices that the model expects, based on the inferred joint posterior distribution over the model parameters α, β, and γ. The mode of the posterior predictive distribution for each participant on each problem was taken as the decision the model expects the participant to have made. The BFO model was able to predict about 77% of the decisions that participants made, given that the base-rate or chance level of agreement is 25% for length 4 problems and 12.5% for length 8 problems, these results suggest that it provides a reasonable account of people's behavior.

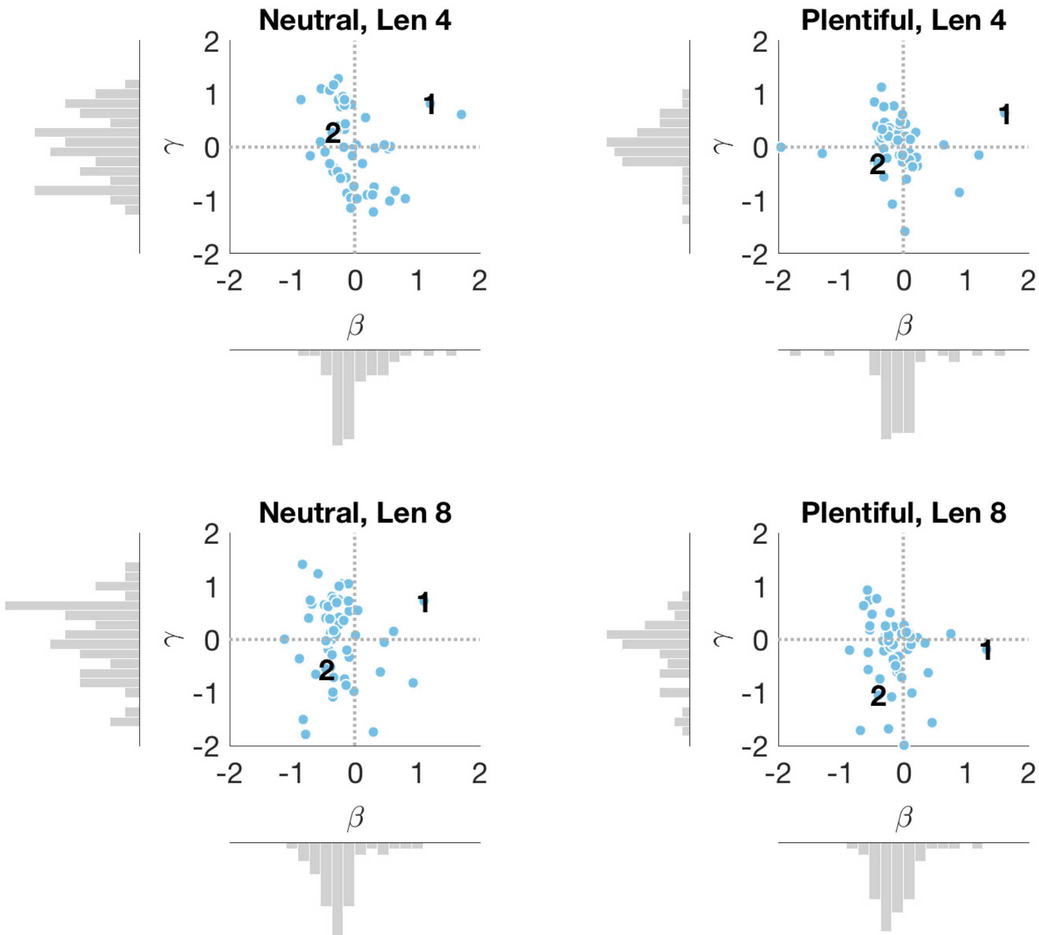

Figure 2.7: Joint and marginal distributions of β and γ posterior expectations across the four conditions for each participant. Participant 1 labeled in this figure is the risk-seeking participant from Figure 2.6, while Participant 2 is the risk-averse participant.

Chapter 3

The Bandit Problem

Bandit problems are also commonly used to study human decision making under risk and uncertainty in cognitive science (Daw et al., 2006; Steyvers et al., 2009; Lee et al., 2011). In bandit problems, people must choose repeatedly between a set of alternatives. Each alternative has a fixed reward rate that is unknown to the decision maker, but they receive feedback after every decision in the form of either a reward or a failure. The goal of this cognitive task is to maximize the total number of rewards over all problems. There are 'infinite horizon' bandit problems, where the total number of trials (or decisions to be made) within a problem are unknown in advance. There is some probability that the problem stops at any trial. There are also "finite horizon" bandit problems, where the total number of trials within a problem are fixed and known to the decision maker in advance. Furthermore, the number of alternatives to choose from can also vary in multi-arm bandit problems. This work focuses on a version of the multi-arm bandit problem where people are asked to choose between just two alternatives.

Bandit problems are interesting because they provide a sequential decision-making environment where exploration and exploitation must be balanced. People must balance between exploring the different alternatives available, while trying to exploit the alternative with the highest reward rate in the set. There are many different heuristics and models for human decision making in bandit problems, including the ϵ-greedy heuristic (Sutton and Barto, 1998), ϵ-decreasing heuristic (Sutton and Barto, 1998), τ-first model (Sutton and Barto, 1998), and the τ-switch model (Lee et al., 2011).

One of the most widely studied models for human decision making in bandit problems is the Win-Stay Lose-Shift (WSLS) heuristic (Robbins, 1952; Sutton and Barto, 1998). In the most simple and deterministic version of the WSLS heuristic, the decision maker is assumed to stay with an alternative if a reward was given, but shifts to another alternative if there was a failure to reward. In the stochastic version of the WSLS heuristic, there is a probability of staying after a reward, and a probability of shifting after a failure, which are both parameterized by a single probability parameter γ. This work focuses on the extended WSLS model, where the probability of staying after a reward (γ_w) and the probability of shifting after a failure (γ_l) are both free parameters and therefore are allowed to differ (Zhang and Lee, 2010b,a). The extended WSLS model allows there to be a psychological difference between successes and failures in the decision making process.

The extended WSLS model does not require memory of previous actions and outcomes, except for the immediately preceding trial. This heuristic is also insensitive to whether the horizon is an infinite or a finite number of trials, since the decision process is identical for all trials and only depends on whether the previous trial resulted in a reward or failure. One can imagine that the WSLS model can capture the risk propensity of an individual doing the cognitive task. Someone who is risk-seeking might shift to another alternative with a high probability following a failure, in order to explore the other available options. However, someone who is risk-averse might shift to another alternative with a relatively lower probability following a failure. They might be inclined to stick around and continue to exploit the current alternative, even after a decision leads to a failure.

The following section describes the bandit problem task in detail. The next section after gives the details of the extended WSLS model, and how it is applied to measure risk propensity in individuals doing two-armed finite-horizon bandit problems. Lastly, this chapter concludes with empirical and modeling results for bandit problems.

3.1 Experiment

In the bandit task, all Amazon Mechanical Turk participants completed four sets of bandit problems: 1) length 8 problems in the neutral environment, 2) length 8 problems in the plentiful environment, 3) length 16 problems in the neutral environment, and 4) length 16 problems in the plentiful environment. In the length 8 problems, participants were asked to maximize their reward between two alternatives with 8 trials within each problem. In the length 16 problems, participants were asked to maximize their reward between two alternatives with 16 trials within each problem. In the neutral environments, reward probabilities were generated from a $\text{Beta}(1, 1)$ distribution. In the plentiful environments, reward probabilities were generated from a $\text{Beta}(4, 2)$ distribution. Consequently, the plentiful environments contained alternatives that had relatively higher reward rates, while the neutral environments contained alternatives that had relatively lower reward rates. All participants completed the same 40 problems within each set, while the order of problems within each set was randomized across participants.

Participants were instructed to maximize the number of rewards they get by pulling the arms of two slot machines. Before beginning each block of problems, they were told that the reward probabilities for each machine are different for each problem in the block, as well as how many pulls they are allowed for each problem. However, they are not told the underlying distribution of the reward probabilities. Figure 3.1 shows a screenshot of the bandit experiment online.

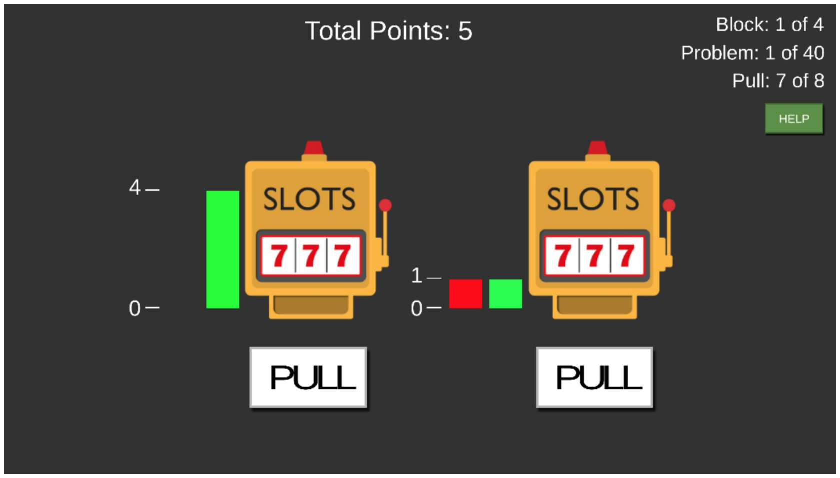

Figure 3.1: Screenshot of the online experiment in the bandit task.

The top-right corner shows the current block out of four total blocks, the current problem, and current trial pull they are on. Participants indicate their selection for each pull of the slot machine by clicking the "pull" button under the slot machine that they wish to choose. Feedback of either reward or failure is given after each pull within a problem, in the form of a green or red bar. If a pull resulted in a reward, a green bar would be added to the left side of the chosen slot machine. If a pull resulted in a failure, a red bar would be added to the left side of the chosen slot machine. In this screenshot, the participant is currently on the 7th pull in a length 8 bandit problem. The total reward points earned on the current problem is shown on the top of the screen. Alternatively, the participant can also count the total points by adding up the height of the green bars. A green bar represents reward and is worth one point, while a red bar means failure which is worth zero points. A problem is completed once the participant completes all pulls in a problem. The participant is reminded with a pop-up message that the next two slot machines will have different reward probabilities, along with a "next" button to move on to the next problem.

3.2 Extended Win-Stay Lose-Shift Model

The extended Win-Stay Lose-Shift model (e-WSLS) is taken from Zhang and Lee (2010b). The e-WSLS is a natural extension to the original WSLS model, except that it allows the probabilities to differ depending on a reward or a failure. In the original WSLS model, the probability of staying following a reward and shifting following a failure are the same. However, the e-WSLS model has two free parameters to quantify these two probabilities. The probability of staying on the current alternative following a reward is denoted as γ_w, and the probability of shifting to the other alternative following a failure is denoted as γ_l.

The probability of choosing the ath alternative ($a = 1, 2$) on the kth trial in the jth problem is:

$$\theta_{jk}^{a} = \begin{cases} 1/2 & \text{, if } k = 1 \\ \gamma_w & \text{, if chose } a\text{, and } r_{k-1} = 1 \\ 1 - \gamma_l & \text{, if chose } a\text{, and } r_{k-1} = 0 \\ \gamma_l & \text{, if chose } \overline{a}\text{, and } r_{k-1} = 0 \\ 1 - \gamma_w & \text{, if chose } \overline{a}\text{, and } r_{k-1} = 1 \end{cases} \tag{3.1}$$

where r_{k-1} is 1 if the previously selected alternative resulted in a reward, and 0 if the previously selected alternative resulted in a failure. The observed rewards and failures on each trial r_k are generated by:

$$r_k \sim \text{Bernoulli}\big(p_i\big) \tag{3.2}$$

where p_1 and p_2 are the reward rates for the two alternatives. These reward rates are generated by two different distributions depending on whether the current block is in the *neutral* or *plentiful* environment:

$$\text{Neutral: } p_i \sim \text{Beta}\big(1,1\big) \tag{3.3}$$

$$\text{Plentiful: } p_i \sim \text{Beta}\big(4,2\big) \tag{3.4}$$

Finally, priors are placed on both probabilities $\gamma_w, \gamma_l \sim \text{Uniform}\big(0,1\big)$. The e-WSLS model for the bandit task is implemented as a Bayesian graphical model in JAGS. All modeling results are based on four chains of 1000 samples each, collected after 2000 discarded burn-in samples. The chains were verified for convergence using the standard $\hat{R}$ statistic (Brooks and Gelman, 1997).

3.3 Results

3.3.1 Removing Contaminants

First, a contaminant model was implemented to identify participants who may not have performed the task properly for any reason, and should therefore be excluded from the analyses. This contaminant model provides an alternative account for how participants made decisions in the task, and then model selection techniques can be used to determine the evidence for each of these models from the behavioral data.

To determine whether participants were behaving at random in the bandit task, a *guessing* model was implemented as a baseline (as in the optimal stopping task). According to the guessing model, the probability that any participant chooses the ath alternative ($a = 1, 2$) on the kth trial in the jth problem is:

$$\theta^a_{jk} = \frac{1}{2} \tag{3.5}$$

This means that on every trial within a problem, the participant chooses to pull a slot machine at random. The guessing model was compared to the e-WSLS model using Bayes factors (Kass and Raftery, 1995). The Bayes factors for each participant were again estimated using a latent mixture procedure based on model-indicator variables. In the bandit task, all participants showed overwhelming evidence for the e-WSLS model over the guessing model. Consequently, no contaminant participants were removed and a total of 56 participants were included in the bandit task analyses.

3.3.2 Empirical Results

Before going into modeling results, the empirical data was examined to see how people's behavior change in the bandit problems across the four experimental conditions varying in environment and length. Figure 3.2 shows the total number of switches (i.e., shifting to the other alternative) following a reward versus a failure. First, it is clear that the total number of switches is greater in the length 16 conditions compared to the length 8 conditions. This is expected because there are simply more trials and therefore opportunities to switch in the length 16 condition. More interestingly, the number of switches following a failure (i.e., "lose-shift") is much higher than the number of switches following a reward, but only for the *neutral* conditions. Intuitively, the number of switches following a failure is expected to be higher than the number of switches following a reward (hence "win-stay lose-shift"). However, due to the plentiful environment being saturated with rewards, participants have less of a reason to switch to the other alternative even after the occasional failure.

It is also interesting to look at the switching behavior across positions in the sequence of a problem. Figure 3.3 shows the number of switches across positions for every participant in each of the four conditions. First, it is clear that there are individual differences in how often participants switch according to position in the sequence. However, the general pattern, as

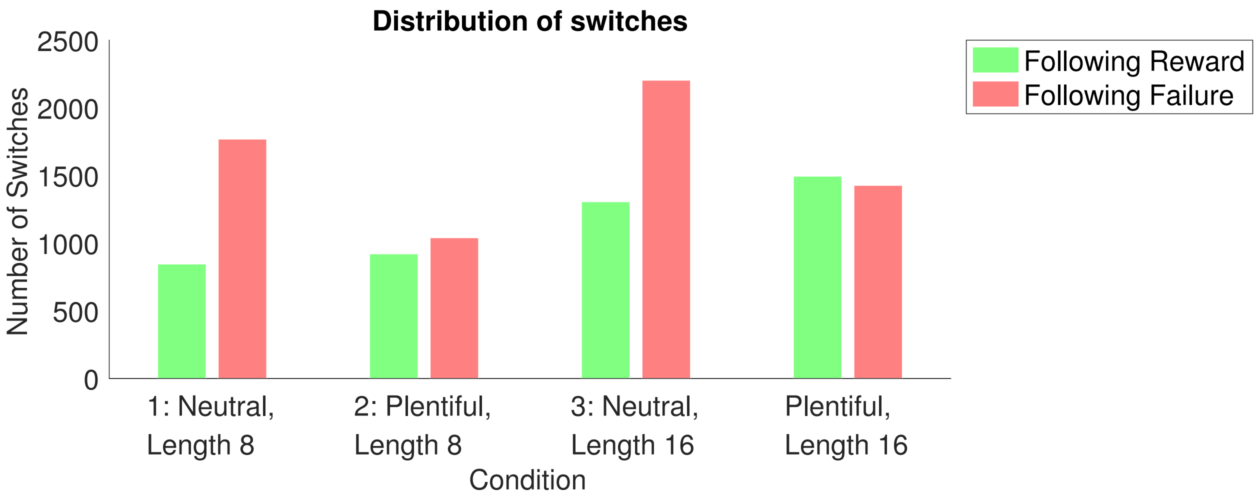

Figure 3.2: Number of switches following rewards vs. failure by condition. The number of switches are much higher following failure than reward in the two neutral conditions, but this effect is not seen in the two plentiful conditions.

expected, is that the number of switches earlier on in the sequence is higher than the number of switches later in the sequence. For all four conditions, it appears that those participants who switch often in earlier positions tend to switch much less when they are about halfway through the sequence. It also appears that participants in the plentiful conditions switch slightly less in general compared to the neutral conditions of the same length. Furthermore, it is also evident that some participants show spikes in the number of switches in certain positions. It appears that some participants tend to switch on the odd number positions, which suggests that they might select a particular alternative twice before moving on to explore the other alternative. The thick red line shows the mean number of switches across participants for each condition. Although there is a wide range of individual differences, the mean number of switches across all positions are small with a slight decline across the sequence.

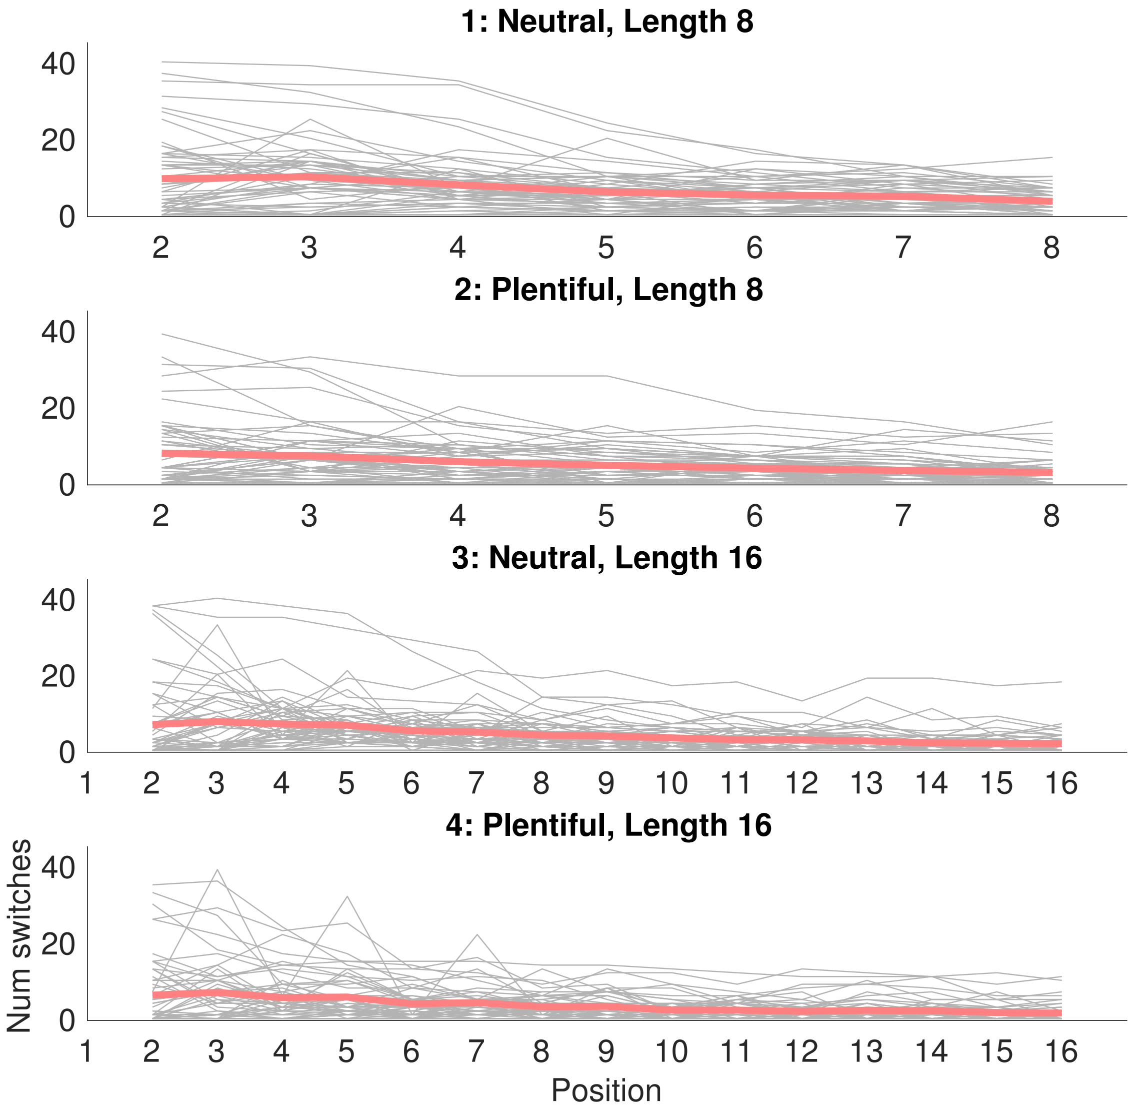

Figure 3.3: Number of switches across positions for all participants by experimental condition. Each panel shows the number of switches to the other alternative by position in the sequence for a particular condition. The thicker red line shows the mean number of switches at each position within the condition.

3.3.3 Inferred Win-Stay Lose-Shift Probabilities

The goal is to capture these individual differences in decision making on bandit problems using the e-WSLS model parameters in terms of risk propensity. Recall that the e-WSLS model has two parameters: γ^w for the probability of staying on the current alternative following a reward, and γ^l for the probability of shifting to the other alternative following a failure. One can imagine that a risk-averse individual might have a low probability of shifting to the other alternative even after a failure, because they are less willing to explore and want to stick with the better-known option. On the other hand, a risk-seeking individual might have a relatively higher probability of shifting after a failure, in order to explore the other alternative in hopes of finding a better reward payout. Therefore, in bandit problems, it might be reasonable to consider the γ^{lose} parameter as a measure of risk propensity. Similarly, someone who behaves consistently might have a relatively high probability of staying on the current alternative following reward, while someone who is relatively less consistent might have a lower probability of staying following reward. Therefore, the γ_{lose} parameter can be interpreted as a measure of consistency.

Figure 3.4 shows the number of switches following reward versus failure across positions for four representative participants. These participants were picked because they span the sorts of individual differences seen in the data set. The left panels show the length 8 conditions while the right panels show the length 16 conditions. Neutral conditions are solid lines while plentiful conditions are dotted lines. The number of switches following reward are shown in green, and the number of switches following failure are shown in red. Participant 1 switches relatively often after a failure but very rarely after a reward, in all four conditions. This is captured by high values of γ^{win} and high values of γ^{lose} inferred by the model. Participant 2 almost does not switch at all whether it is following a reward or a failure, in all four conditions. This is reflected in the very high values of γ^{win} and very low values of γ^{lose}. Participant 3 switches relatively more following failure than Participant 2, but also switches

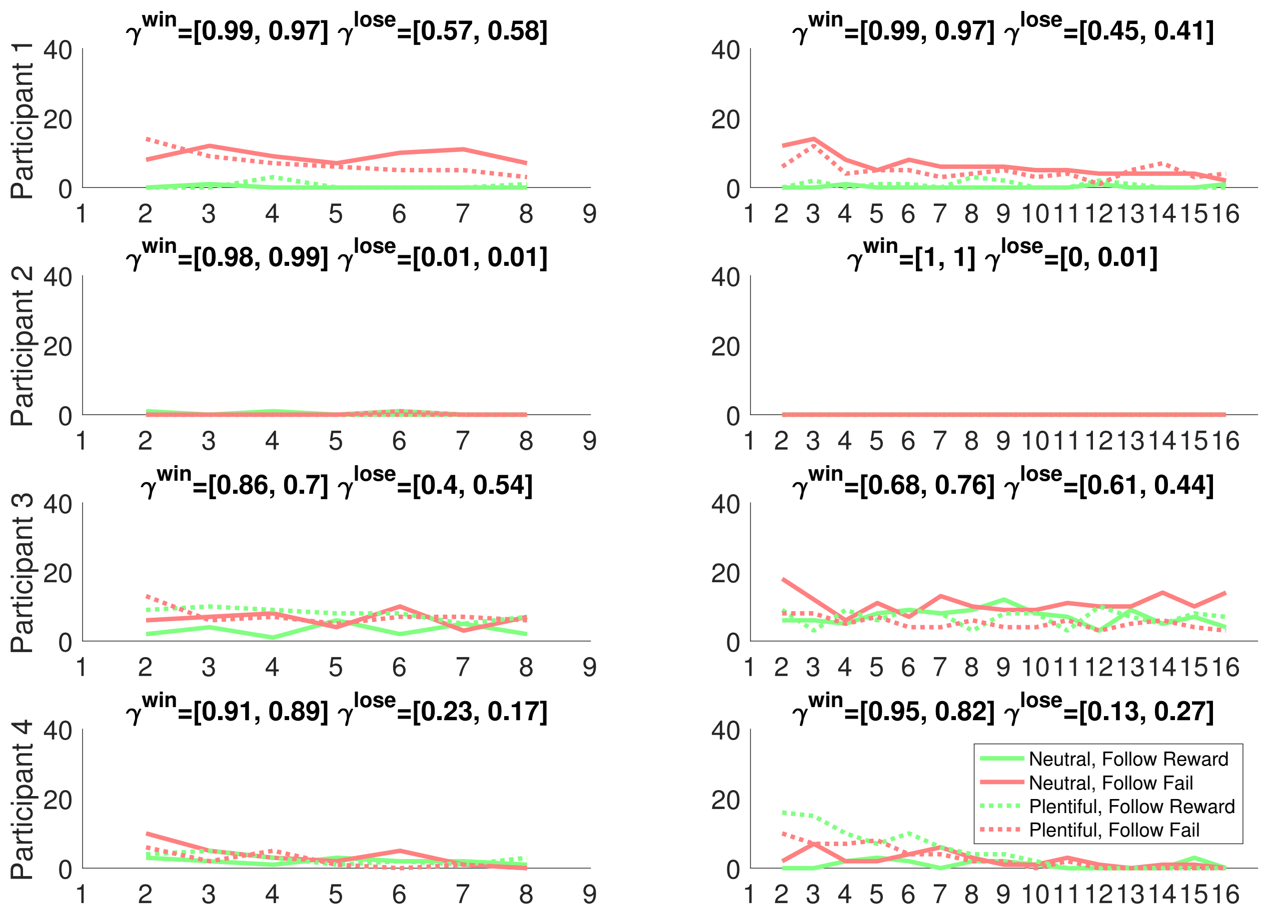

Figure 3.4: Number of switches following reward versus failure for four representative participants in each of the four conditions. The left panels show the length 8 conditions while the right panels show the length 16 conditions. Neutral conditions are solid lines while plentiful conditions are dotted lines. The number of switches following reward are shown in green, while the number of switches following failure are shown in red.

sometimes even following reward. Consequently, they have relatively lower values of γ^{win} with relatively high values of γ^{lose} in all four conditions. Lastly, Participant 4 is somewhere in between with the number of switches following reward and failure to be about the same across the four conditions. The inferred γ^{win} and γ^{lose} values are in between the more extreme Participant 2 and Participant 3 values.

Figure 3.5 shows the joint and marginal distributions of γ^{win} and γ^{lose} posterior expectations across the four conditions for each participant. The four representative participants from Figure 3.4 are labeled in the scatterplot. It is clear that there is a range of individual differences in both win-stay probabilities and lose-shift probabilities. However, there is a negative relationship between the two parameters. Participants who tend to stay following a reward also tend to stay following a failure. Participants who shift relatively more even after a reward also tend to explore the other alternative after a failure.

3.3.4 Descriptive Adequacy

The ability of the model to fit the behavioral data was assessed using a standard Bayesian approach based on the posterior predictive distribution (Gelman et al., 2004). The posterior predictive is the distribution of choices that the model expects, based on the inferred joint posterior distribution over the model parameters γ^{win} and γ^{lose}. The mode of the posterior predictive for each participant on each problem was used as the decision that the model expects the participant to have made. The e-WSLS model was able to predict about 84% of the decisions that the participants made. Given that the chance level of agreement for selecting either of the two alternatives is 50% on each trial within all problems, this result suggests that the e-WSLS model provides a reasonable account of people's behavior.

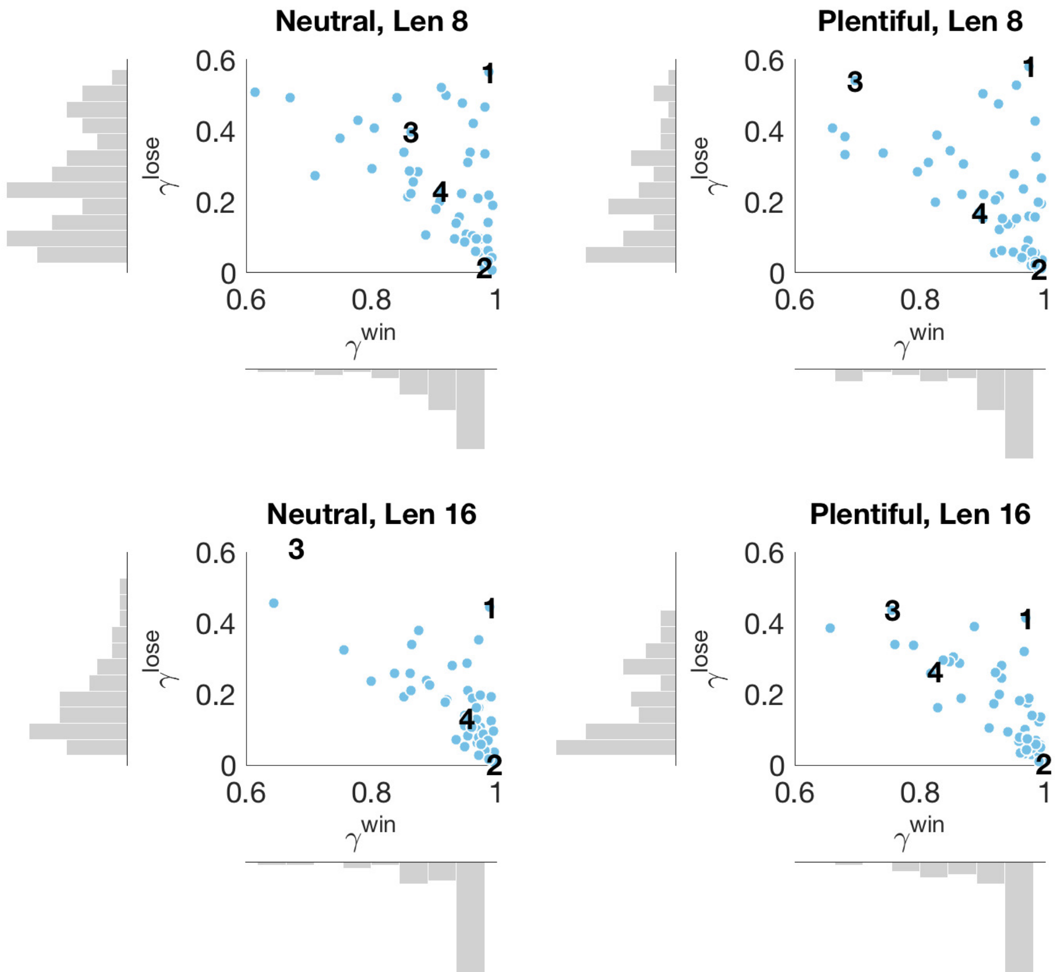

Figure 3.5: Joint and marginal distributions of the means of the γ^{win} and γ^{lose} posterior expectations across the four conditions for each participant. The four representative participants shown from Figure 3.4 are labeled here in the scatterplot.

Chapter 4

Balloon Analogue Risk Task

The Balloon Analogue Risk Task, also known as the BART, is a decision-making task that aims to measure risk-taking behavior (Lejuez et al., 2002). In the BART, a balloon is displayed on a computer screen that represents some monetary value at the beginning of each problem. The participant is then given a choice to either *play it safe* and cash out the current value of the balloon to a virtual bank, or *take a risk* and pump the balloon to add some small amount of air into the balloon. However, there is some probability the balloon will burst each time it is pumped, in which case the participant will receive nothing. On the other hand, if the balloon does not burst after a pump, it grows in size and is worth more money. If this happens, the participant is again given the choice to either bank or pump. The problem ends when the participant banks or when the balloon bursts.

Although the BART is a simple laboratory task, it captures people's tendency to take risks. When the participant decides to pump the balloon, they are pursuing a larger reward while risking the entire amount that the balloon is worth. Consequently, the BART requires the decision maker to balance the choice between a risk-taking action of pumping and a more secure, risk-averse action of banking.

The BART has been used in a variety of contexts in the decision-making literature to study human risk-taking behavior (Lejuez et al., 2003; Lighthall et al., 2009; Rao et al., 2008; Aklin et al., 2005). Human performance on the BART is traditionally quantified by the mean number of pumps across problems, excluding those problems where the balloons bursted. One can imagine that an individual with high risk propensity might pump the balloon relatively more times across problems, while an individual who tends to be risk-averse might pump relatively fewer times across problems. This simple mean number of pumps measure has been shown to correlate with risk taking behaviors such as smoking, alcohol abuse, and drug abuse(Lejuez et al., 2003, 2002; Kathleen Holmes et al., 2009; Hopko et al., 2006) as well as psychological traits such as impulsivity, anxiety, and psychopathy (Hunt et al., 2005; Lauriola et al., 2014).

However, the goal of this work is to model the cognitive processes that determine how humans make decisions in the BART. Rather than simply measuring the performance on the task using a simple metric such as mean number of pumps, the goal is to learn *why* an individual makes the decisions they do. To examine the unobserved psychological processes that are involved in the BART, I implement a cognitive process model that includes parameters that quantify risk taking and behavioral consistency in the task (van Ravenzwaaij et al., 2011). By using a cognitive model, people's performance on the BART can be interpreted separately in terms of both their tendency to take risks and their consistency in the behaviors they make across problems.

The following section gives the details of the BART experiment. Then, the 2-parameter model developed by van Ravenzwaaij et al. from the original 4-parameter model (Wallsten et al., 2005) will be explained in detail. Finally, empirical and modeling results will be shown for the BART.

4.1 Experiment

In the BART, all Amazon Mechanical Turk participants completed two sets of the BART: 1) where the probability of the balloon bursting at each trial in a problem is set to 0.1 and 2) where that bursting probability is set to 0.2. The participants were told at the beginning of the task that they will be pumping balloons from two different bags of balloons, and that balloons from the same bag have the same probability of bursting. However, they were not told the probabilities of bursting are 0.1 and 0.2. All participants completed the same 50 problems within each of the 2 sets, while the order of problems within each set was randomized across participants.

Participants were instructed to collect as much monetary reward as they can. At the beginning of the experiment, they receive a bank with \$0. They were told that each balloon at the beginning of each problem is worth \$1. With each pump, the balloon's worth increases by \$1. At any point within a problem, they can *bank* the amount that the balloon is currently worth and move on to the next problem with a new balloon. They can also *pump* the balloon, but there is some probability that it will burst and the value of that balloon will be lost.

Figure 4.1 shows screenshots of the online experiment. The left panel shows the balloon at the beginning of a problem, and the right panel shows a trial within the problem at which the balloon bursted. The balloon and its current monetary value is shown in the middle at the beginning of each problem. The participants were instructed to click "Pump" if they wish to pump the balloon, or "Bank" if they wish to secure the amount of the balloon into their bank. If the participant chooses to pump, the monetary value of the balloon is added to the Bank at the top of the screen, and the next problem begins with a new balloon. However, if the balloon bursts on the screen, the participant no longer has the option to pump or bank, and must click "Next" to move on the next problem with a new balloon. Furthermore, the

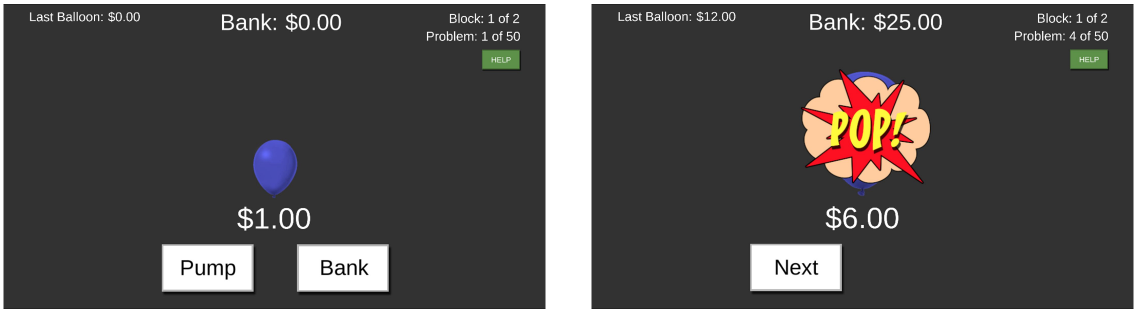

Figure 4.1: Screenshots of the online experiment in the BART.

top-corner shows how much money they were rewarded in the last balloon.

The following section details the cognitive model for the BART, which is a 2-parameter simplification developed by van Ravenzwaaij et al. from the original 4-parameter model (Wallsten et al., 2005).

4.2 2-Parameter Model

The 2-parameter model, taken from van Ravenzwaaij et al. (2011), assumes that the decision maker believes that there is a single constant probability that a pump will make a balloon burst, p^{belief}, which is fixed over all problems. The next assumption is that the participant decides on a number of pumps prior to the first pump in a problem, and does not adjust this number during pumping. This number of pumps that the participant considers to be optimal, denoted by ω, depends on their propensity for risk taking, γ^+, and their belief about the bursting probability of the balloon when it is pumped:

$$\omega = \frac{-\gamma^{+}}{\ln\left(1 - p^{belief}\right)} \tag{4.1}$$

with $\gamma^{+} \geq 0$. Relatively higher values of γ^{+} correspond to more number of pumps that the participant deems optimal, or more risk-seeking behavior. Relatively lower values of γ^{+} correspond to fewer number of pumps that the participant considers optimal, or more risk-averse behavior. Then, the probability that the participant will pump on the kth trial on a particular problem, p_k^{pump}, depends on the number of pumps that the participant considers to be optimal, ω, and their behavioral consistency β:

$$p_k^{pump} = \frac{1}{1 + e^{\beta(k-\omega)}} \tag{4.2}$$

with $\beta \geq 0$. This logistic relationship means that relatively higher values for β corresponds to more consistency in the participant's responses. When $\beta = 0$, $p_k^{pump} = 0.5$ and the participant's decision to pump or bank is random. When $\beta \rightarrow \infty$, the participant's behavior is completely determined by whether or not the pump opportunity k is greater than ω. The version of the 2-parameter model used in this work adds an additional censoring function to the probability of pumping at a given trial. The probability of pumping on a trial within a problem now also depends on whether the current trial exceeds the trial at which the balloon bursts on the problem:

$$p_{j,k}^{pump} = \begin{cases} \frac{1}{1+e^{\beta(k-\omega)}} & \text{, if } k < b_j \\ 0 & \text{, if } k \geq b_j \end{cases} \tag{4.3}$$

where b_j is the trial at which the balloon bursted on problem j. Finally, the probability of the observed data, $p\left(D|\gamma^{+}, \beta\right)$ for all j problems and all pumps within each problem $n_{(j)}$, depends on whether the jth problem resulted in a bank ($d_j = 1$) or a balloon burst ($d_j = 0$). The likelihood function that connects the two parameters to the data is:

$$p\left(D|\gamma^{+},\beta\right)=\prod_{j=1}^{J}\prod_{i=1}^{K_{(j)}}p_{ji}^{pump}\left(1-p_{j,k_{(j)}+1}^{pump}\right)^{d_j} \tag{4.4}$$

where $d_k = 1$ if the participant banked on trial k and $d_k = 0$ if the balloon burst. This likelihood is the product of all probabilities that the participant will pump times one minus this probability on those problems where the participant banked. Finally, priors are placed on the risk and consistency parameters for each participant: $\gamma_{+}, \beta \sim \text{Uniform}\big(0, 10\big)$. The 2-parameter BART model implemented as a Bayesian graphical model in JAGS. All modeling results are based on four chains of 1000 samples each, collected after 2000 discarded burn-in samples. The chains were verified for convergence using the standard $\hat{R}$ statistic (Brooks and Gelman, 1997).

The two parameters to be inferred per participant in this model are risk propensity γ^{+} and behavioral consistency β. The parameter of interest for the goal of measuring risk propensity across tasks is the γ^{+} parameter.

4.3 Results

4.3.1 Removing Contaminants

First, a subset of contaminant participants were identified in the data set. In order to identify those participants who either did not understand the task properly or simply rushed through the task to finish the experiment as fast as possible, the consistency parameter β was used to form a cutoff for contaminant behavior. If a participant's behavior was extremely inconsistent across all problems within the two conditions, they were considered contaminants. A cutoff of 0.2 was used for β, removing 11 participants due to inconsistency. Furthermore, if a participant simply banked on all of the problems to finish the experiment as soon as possible, they were also considered contaminants. A total of 14 contaminant participants were removed from the BART in the remainder of the analyses.

4.3.2 Inferred Risk Propensity and Consistency

The 2-parameter model quantifies people's risk propensity with γ^+, and consistency with β. The risk parameter γ^+ along with the probability of a balloon popping determine the number of pumps that a participant considers to be optimal, ω, in a given condition. If someone has a relatively large value of γ^+, then their optimal number of pumps is also large, which corresponds to being risk-seeking.

To show how these two parameters capture the risk propensity and consistency of individuals, Figure 4.2 shows the inferred γ^+ and β parameter values for four representative participants along with their empirical data. These participants were chosen because they span the range of individual differences that we see in risk propensity and consistency in the data set. The histograms show the distribution of the number of pumps that each participant made, excluding problems where the balloon bursted. The inferred values for each of the parameters are shown at the top of each panel, where ω is the number of pumps that the participant considers to be optimal. The four panels on the left are from condition 1, where $p^{burst} = 0.1$, while the four panels on the right are from condition 2, where $p^{burst} = 0.2$.

Participant 1 can be seen as risk seeking and consistent. They tend to pump a relatively large number of times across both conditions, and is quite consistent in doing so. This pattern of behavior is captured by the risk and consistency parameters, with relatively high values of γ^+ and relatively high values of β. Participant 2 is also risk seeking in that they pump a relatively large number of times across both conditions, but they are very inconsistent. The number of times they pump in both conditions vary widely from about 3 to more than 15 pumps. This behavior is also reflected in their inferred parameter values, with relatively

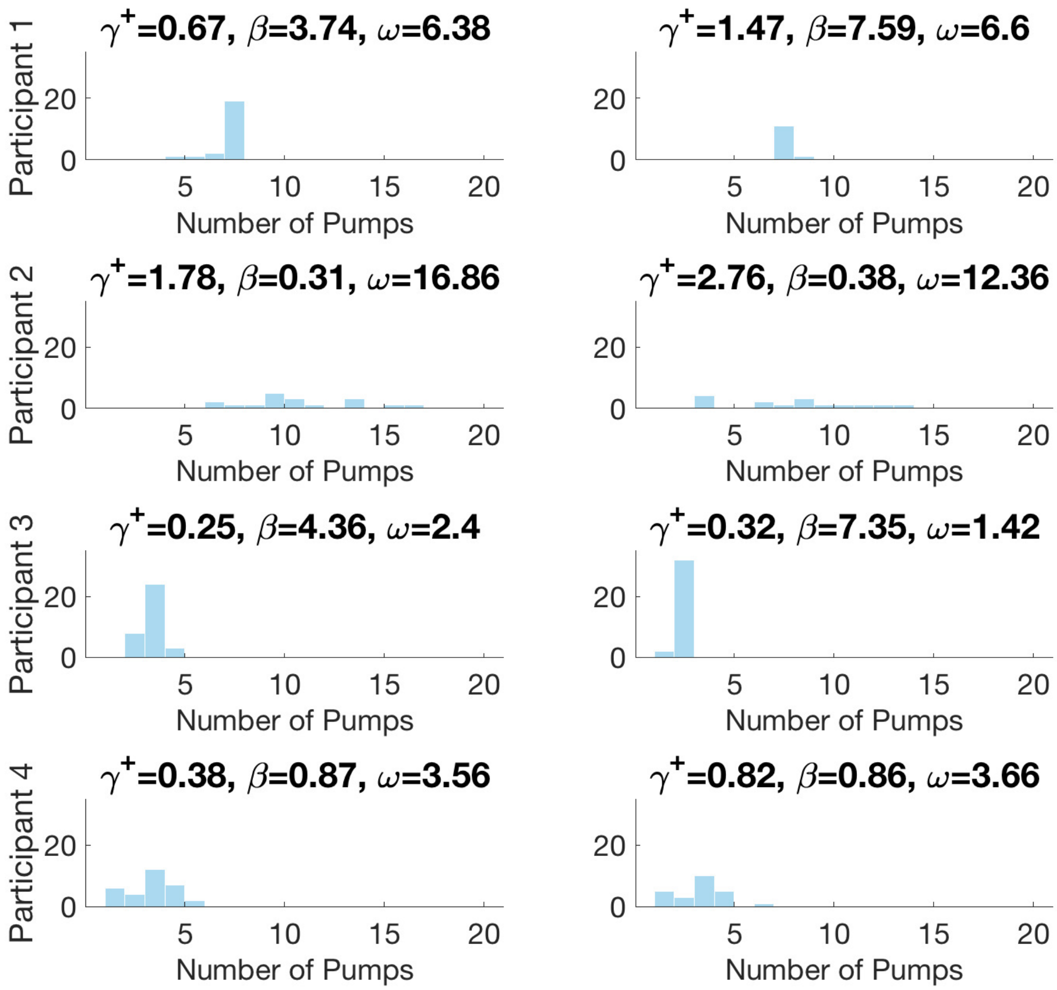

Figure 4.2: Empirical data of decisions made and inferred parameter values for four representative participants. The subplots in the left column are from condition 1 ($p^{burst} = 0.1$) and the subplots in the right column are from condition 2 ($p^{burst} = 0.2$). The histograms show the distribution of the number of pumps each participant made before banking, excluding problems where the balloon bursted. The inferred values of γ^+, β, and ω are shown above.)

high values of γ^+ but relatively low values of β.

On the other hand, participant 3 can be seen as consistently risk averse. They pump a relatively small number of times across both conditions, and are very consistent in doing so. This is reflected in a relatively low γ^+ and high β. Lastly, participant 4 is risk averse but more inconsistent than participant 3. They pump a relatively small number of times across both conditions as well, but are less consistent in doing so. This is captured with relatively low values of both γ^+ and β.

Figure 4.3 shows the joint and marginal distributions of γ^+ and β posterior expectations across the two conditions for each participant. The four representative participants shown from Figure 4.2 are labeled in the scatterplot. First, it is evident that there is a wide range of individual differences in both risk propensity and consistency parameters. However, there appears to be a negative, nonlinear relationship between the two parameters in both conditions. Participants with relatively high values of γ^+ also tend to have low values of β, and vice versa. Participants near the origin have low values of both γ^+ and β, and are consequently both risk averse and inconsistent. However, as participants move further away from the origin and closer to the top-left corner, they become more risk seeking. As participants move from the origin closer to the lower-right corner, they become more consistent.

Participant 1, who is relatively consistent and risk seeking, can be seen towards the high end of β as well as the upper end of γ^+. Participant 3, who is relatively consistent but risk averse can be seen almost directly below participant 1. Participant 4, who is inconsistent and risk averse can be seen close to the origin. Participant 2, who is risk seeking but very inconsistent can be seen close to the y-axis but with a high value of γ^+ compared to everyone else.

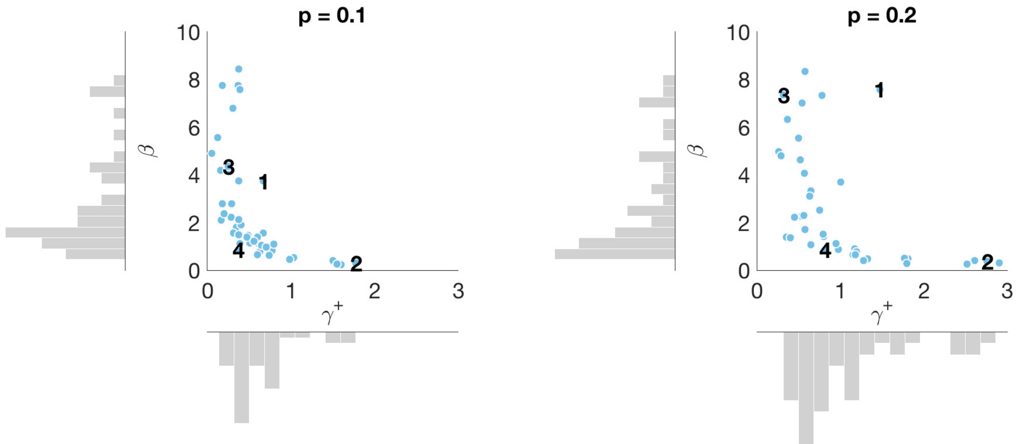

Figure 4.3: Joint and marginal distributions of β and γ^+ posterior expectations across the two conditions for each participant. The four representative participants shown from Figure 4.2 are labeled here in the scatterplot

4.3.3 Descriptive Adequacy

A standard Bayesian approach based on the posterior predictive distribution was used to assess the ability of the model to fit the behavioral data (Gelman et al., 2004). In this analysis, the posterior predictives were generated by sampling 1000 values of our risk parameter γ^+ and consistency parameter β from the joint posterior distribution for each participant in each condition. From each set of γ^+ and β values, a trial of BART data was generated. The density across the sampled number of pumps for each participant and each condition is shown in Figure 4.4 as gray boxes of varying sizes. The observed data are shown to the left of the posterior predictive distributions. The dots represent the median number of pumps, the thicker solid lines represent the 0.25 and 0.75 quantiles, and the thin solid lines span the minimum and maximum observed number of pumps. The top panel shows this for condition 1, with probability of bursting $p = 0.1$, and the bottom panel shows this for condition 2, with probability of bursting $p = 0.2$. If the model predicted data closely resembles the observed data, then the model provides an adequate account of how behavior was generated in the BART. From this figure, it is clear that the posterior predictives tend to follow the observed

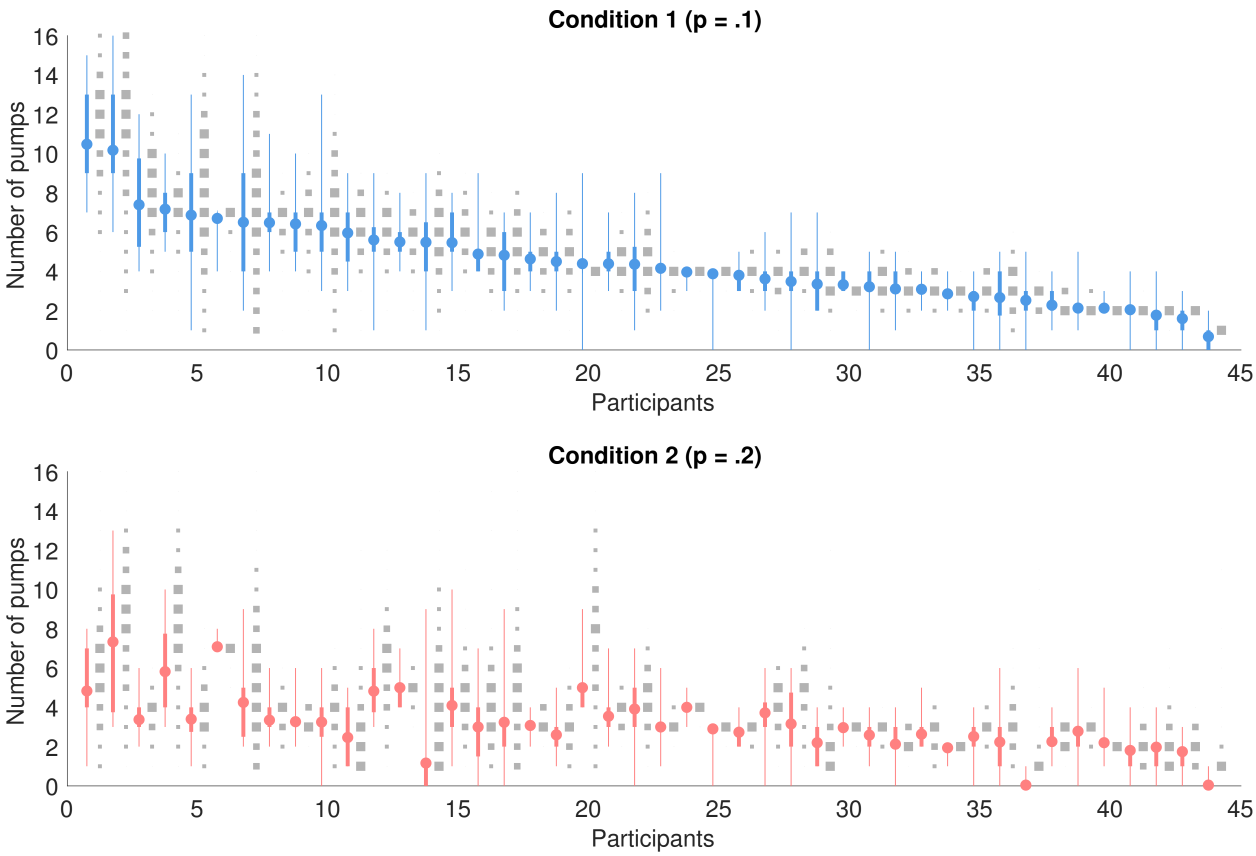

Figure 4.4: Posterior predictive distributions of the number of pumps for each participant in the two conditions, sorted by the mean number of pumps per participant in condition 1. The top panel shows the posterior predictives for condition one, and the bottom panel shows the posterior predictives for condition two. Posterior predictives shown as gray boxes of varying size are based on parameter estimates of γ^+ and β. The minimum and maximum, 0.25 and 0.75 quantiles, and median of the observed behavioral data are shown to the immediate left of the posterior predictives in blue for condition 1 and red for condition 2.

data, suggesting that the model provides a reasonable account of people's behavior.

Chapter 5

Gambling Task

Perhaps the most commonly-used task for studying how people make decisions under risk and uncertainty is letting people choose between pairs of gambles (De Martino et al., 2006; Russo and Dosher, 1983; Rieskamp, 2008). Each gamble has some probability and monetary outcome associated with it, which yields a different monetary outcome that is associated with some probability. For instance, a participant might be asked to choose between Gamble A, which leads to \$50 with probability 0.6 and -\$50 with probability 0.4 (denoted as [\$50, 0.6; -\$50, 0.4], and Gamble B, which leads to \$100 with probability 0.65 and -\$100 with probability 0.35 ([\$100, 0.65; -\$100, 0.35]). The goal is to examine how people make decisions when presented with two gambles such as the pair above.

The economic view assumes that people make decisions rationally under risk and uncertainty, and simply choose the gamble that maximizes expected utility (Von Neumann and Morgenstern, 1947). This view suggests that people can calculate the expected value of the outcome of Gamble A and Gamble B, and choose the alternative that yields the larger expected payoff. In the example above, the expected payoff is \$10 for Gamble A and \$30 for Gamble B. Hence, according to the expected utility view, the decision maker should find

Gamble A more attractive than Gamble B and pick Gamble A.

However, the expected utility view has been challenged widely in the field of psychology. There are many alternative theories to expected utility theory developed by psychologists over the years including regret theory (Loomes and Sugden, 1982), decision field theory (Busemeyer and Townsend, 1993), the priority heuristic (Brandstätter et al., 2006), anticipated utility theory (Quiggin, 1982). This thesis focuses on one of the most well-known alternatives, prospected theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1981). Under the perspective of prospect theory, people have subjective weights on the value of outcomes and probabilities. Specifically, people also weigh outcomes differently depending on whether they are framed in terms of gains or losses. Prospect theory is able to account for people's tendency to be risk-averse for gains of large probability and risk-seeking for gains of small probability. Furthermore, people also tend to be risk-averse for losses of small probability while risk-seeking for losses of large probability.

One can imagine that if an individual exhibits strong loss aversion, this can be seen as being risk-averse. That is, if someone is more strongly impacted by a *loss* of a large magnitude than a *gain* of the same magnitude, they are risk-averse. Consequently, a model with a loss aversion parameter can be used to measure people's risk propensity. This thesis implements the Cumulative Prospect Theory model taken from Nilsson et al..

In the next section, the gambling task will be described in detail. The next section gives the details of the Cumulative Prospect Theory model. Lastly, this chapter concludes with modeling results for the gambling task.

5.1 Experiment

In the version of the gambling task used in this experiment, all Amazon Mechanical Turk participants completed two sets of gambling tasks: 1) the gambles were framed in terms of *gains* and 2) the gambles were framed in terms of *losses*. In the gains environment, participants were instructed to maximize their monetary reward, while in the losses environment, participants were instructed to minimize their monetary losses over the entire set of problems. All participants completed the same 40 problems within each set, while the order of problems within each set was randomized across participants.

Participants were instructed at the beginning of each set of tasks to either maximize reward or minimize losses. The pairs of gambles were presented as pie charts labeled with their respective payouts and probabilities. The expected value of the outcomes are not provided for the participants. They were to indicate their selection of which gamble they prefer, either Gamble A or Gamble B, by clicking either the pie chart on the left or the right. There was no feedback given after each problem. After indicating their preference of either Gamble A or Gamble B, the experiment moves on to the next problem with a new pair of gambles.

Figure 5.1 shows a screenshot of the online experiment in the losses environment. Gamble A is always on the left side, while Gamble B is always on the right side. The respective probabilities and payouts are color-coded and shown above the pie chart. Participants indicate their preference by clicking anywhere on the pie chart for the gamble they wish to select. The top-right corner shows the problem number and block number that they are currently in.

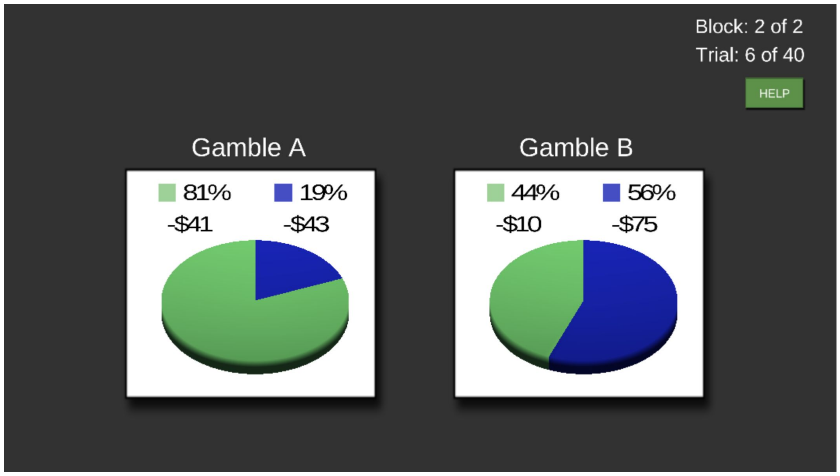

Figure 5.1: Screenshot of the online experiment in the gambling task.

5.2 Cumulative Prospect Theory Model

The Cumulative Prospect Theory (CPT) model is taken from (Nilsson et al., 2011). According to prospect theory, people do not always choose the alternative that maximizes their expected utility. The expected utility (EU) of an alternative O is defined as:

$$EU(O) = \sum_i p_i u(x_i), \tag{5.1}$$

where $u(\cdot)$ is a utility function that defines the *subjective utility* of x_i. The subjective utility of x_i is weighted by the probability p_i that outcome i occurs.

Prospect theory follows expected utility theory in that it assumes that each alternative has a subjective value that quantifies its desirability to the decision maker. However, the way that this subjective value is computed is different from expected utility theory. Prospect theory assumes that the outcomes of risky alternatives are evaluated relative to a reference point, so that outcomes can be framed in terms of losses and gains. Prospect theory also assumes that the absolute value of a loss has a larger impact on the decision than the same absolute value in terms of a gain. This phenomenon is known as loss aversion (Kahneman and Tversky, 1979). Another assumption of prospect theory is that people also have a subjective representation of probabilities. People are assumed to overestimate small probabilities and underestimate large probabilities. CPT (Tversky and Kahneman, 1992) is a refined version of the original prospect theory. According to CPT, if alternative O has two possible outcomes (such as the alternatives in the gambling task of this thesis), then the subjective value V of O is defined as:

$$V(O) = \sum_i \pi(p_i) v(x_i), \tag{5.2}$$

where $\pi(\cdot)$ is a weighting function of the objective probabilities and $v(\cdot)$ is a function defining the subjective value of outcome i. The version of the CPT model used in this thesis assumes that the probability weighting function and the value function differ for gains and losses. The subjective value of payoff x is defined as:

$$v(x) = \begin{cases} x^{\alpha} & \text{, if } x \geq 0 \\ -\lambda(-x)^{\alpha} & \text{, if } x < 0 \end{cases} \tag{5.3}$$

where $0 < \alpha < 1$ is a free parameter that modulates the curvature of the value function. In the original Nilsson et al. (2011) paper, the value function differs for gains and losses. However, this thesis uses a simplification of the model where the shape of the value function, determined by α is the same for gains and losses. The $\lambda > 0$ parameter specifies loss aversion, with larger values of λ reflecting larger loss aversion. When $0 < \lambda < 1$, gains have larger impact on the decision than losses, which is the opposite of loss aversion. Although CPT assumes that losses carry more weight than gains, which is the case when $\lambda > 1$, this model tests the loss aversion assumption by allowing λ to be less than 1. If an individual has a λ value less than 1, that person gives larger weight to gains than to losses of the same

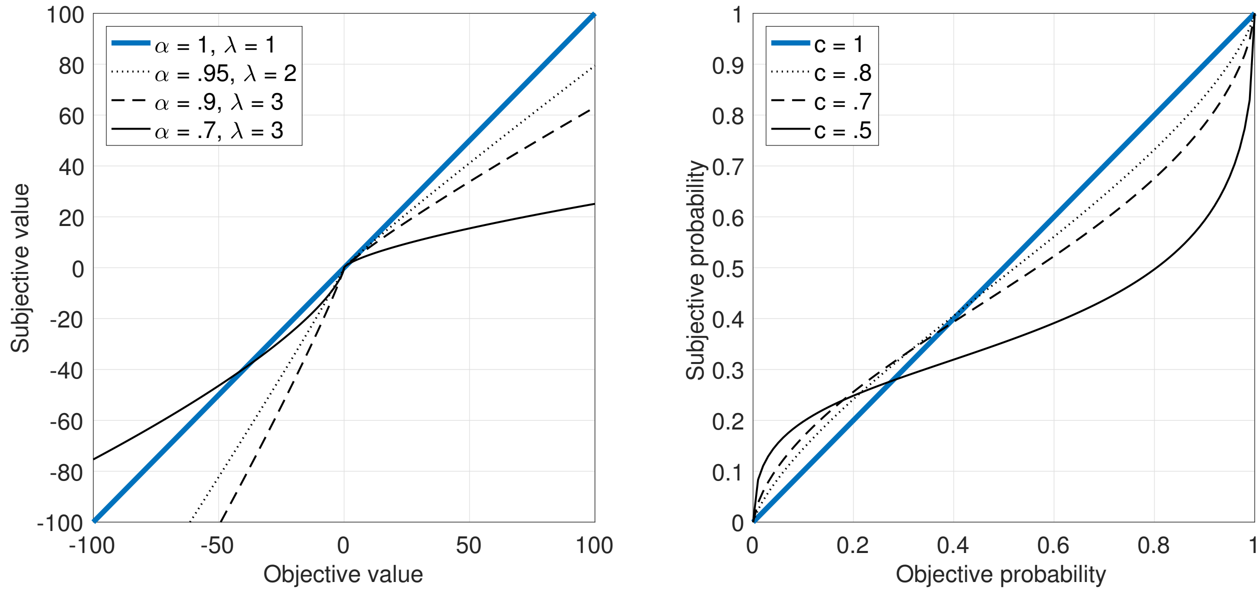

Figure 5.2: Different parameterizations of α, λ, and c to generate various shapes of the value functions and probability weighting functions. The left panel shows three value functions compared to the case where $\alpha = \lambda = 1$ (no loss aversion, and subjective utility is equal to expected utility). The right panel shows three probability weighting functions compared to the case where $c = 1$ (subjective probability is equal to objective probability).

magnitude.

Probabilities are transformed to subjective probabilities by a weighting function. The weighting function for two possible outcomes (as for all alternatives in this task) can be defined as:

$$\pi\left(p_i\right) = \frac{p_i^c}{\left(p_i^c + (1 - p_i^c)^c\right)^{1/c}} \tag{5.4}$$

where $c = \gamma$ for gains and $c = \delta$ for losses. Parameter $0 < c < 1$ determines the inverse s-shape transformation of the weighting function.

Figure 5.2 shows different parameterizations of α, λ, and c to generate various shapes of the value functions and probability weighting functions. The left panel shows three value functions compared to the special case of $\alpha = \lambda = 1$, where there is no loss aversion and subjective utility is equal to expected utility. As α moves away from 1 towards 0, the subjective value of larger magnitudes for both gains and losses are under-valued. As λ increases, loss aversion becomes greater and the subjective value of the magnitude of losses become greater than the subjective value of the magnitude of gains. The dotted line shows a value function with $\alpha = 0.95, \lambda = 2$, the dashed line shows a value function with $\alpha = 0.9, \lambda = 3$, and the thinner solid black line shows a value function with $\alpha = 0.7, \lambda = 3$.

The right panel shows three weighting functions compared to the special case of $c = 1$, where subjective probability is equal to objective probability. As long as $c < 1$, small probabilities are overestimated and large probabilities are underestimated. As c moves away from 1 towards 0, the magnitude of the overestimation and underestimation increases. The dotted line shows a weighting function with $c = 0.8$, the dashed line shows a weighting function with $c = 0.7$, and the thinner solid black line shows a value function with $c = 0.5$.

To allow for the probabilistic nature of human decision making, the version of the CPT implemented uses a choice rule where the choice probabilities are assumed to be a monotonic function of the differences of the gambles' subjective values. Specifically, the exponential Luce choice rule (rewritten as a logistic choice rule) states that the probability of choosing alternative A over alternative B, denoted by $\theta_{A,B}$, is defined as:

$$\theta_{A,B} = \frac{1}{1 + e^{\phi[V(B) - V(A)]}} \tag{5.5}$$

The parameter ϕ can be seen as a consistency measure of choice behavior. When ϕ equals 0, the probability of choosing alternative A over B becomes 0.5, and choice behavior is random. As ϕ increases, choice behavior is determined increasingly more by the difference in subjective value between alternative A and alternative B. As ϕ approaches infinity, the individual will always be consistent with their preference, and the probabilistic version of the CPT becomes identical to the deterministic version.

Finally, priors are placed on all five free parameters for each participant: $\alpha \sim \text{Uniform}\big(0, 1\big), \lambda \sim$

$\text{Uniform}(0, 10), \gamma \sim \text{Uniform}(0, 1), \delta \sim \text{Uniform}(0, 1), \phi \sim \text{Gamma}(2, 1)$. The CPT model is implemented as a Bayesian graphical model in JAGS. All modeling results are based on four chains of 1000 samples each, collected after 2000 discarded burn-in samples. The chains were verified for convergence using the standard $\hat{R}$ statistic (Brooks and Gelman, 1997).

Although there are a total of five free parameters $(\alpha, \lambda, \gamma, \delta\, \phi)$, this thesis focuses on the loss aversion parameter λ as a measure of risk propensity, and the consistency parameter ϕ which can be compared to the the other consistency parameters in other risk tasks.

5.3 Results

5.3.1 Removing Contaminants

A contaminant model was implemented in order to identify a potential subset of participants who did not perform the task properly for any reason, and should therefore be excluded from the analysis. A method to allow for the possibility of contaminant participants is to provide an alternative account for how they made decisions in the task, and use model selection techniques to determine the evidence for each model from the behavioral data.

To determine whether participants were behaving at random in the gambling task, a simple *guessing* model was implemented as a baseline. The guessing model does not have any free parameters. According to the simple guessing model, the probability that the any participant will choose alternative A over B is

$$\theta_{A,B} = \frac{1}{2} \tag{5.6}$$

This guessing model was compared to the CPT model using Bayes factors (Kass and Raftery,

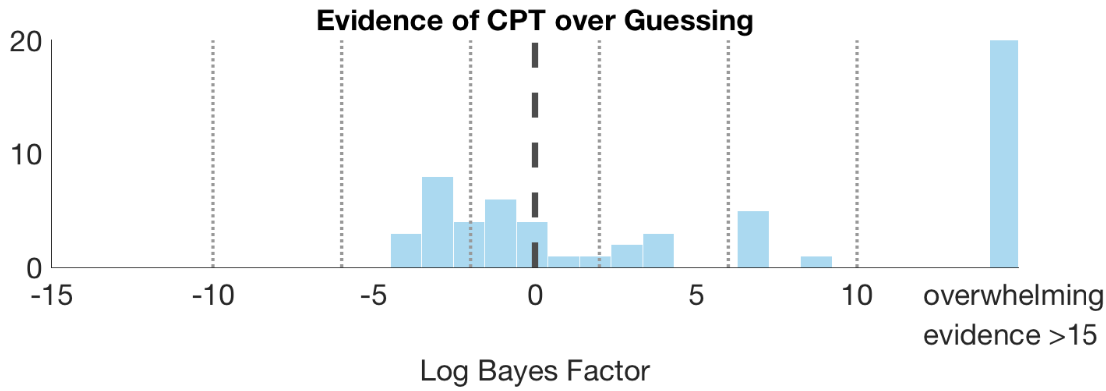

Figure 5.3: Distribution of log Bayes factors for the Cumulative Prospect Theory model over the simple guessing model.

1995), just as in the optimal stopping task. Bayes factors for each participant were estimated using a latent mixture procedure based on model-indicator variables. Unlike in the optimal stopping task, there were a number of participants who showed evidence for the guessing model. Figure 5.3 shows the distribution of log Bayes factors for each participant in the gambling task. The darker dashed line indicates a log Bayes factor of 0, or no evidence towards either model. The standard interpretative boundaries at log-odds of 2, 6, and 10 corresponding to "moderate", "strong", and "very strong" evidence (Kass and Raftery, 1995) are also shown as lighter dotted lines. The participants where there was overwhelming evidence (log Bayes factor too large to appear on this figure) are shown on the right, where the x-axis is marked as "overwhelming evidence".

Participants who had negative log Bayes factors, or showed evidence towards the guessing model, were considered contaminants. Consequently, 24 participants were removed from the gambling task for remainder of the analyses, with a total of 34 participants remaining for the gambling task.

5.3.2 Inferred Subjective Value Functions

The CPT model infers the shapes of the subjective value function and the probability weighting function that people use. Specifically, the values of α and λ determine the shape of the subjective value function for both gains and losses, and the values of γ and δ determine the shapes of the probability weighting functions for gains and losses, respectively. Lastly, the ϕ parameter captures the consistency of people's choice behavior.

There are wide individual differences in the subjective value functions and probability weighting functions that people use, which is reflected in the parameter estimates of the CPT parameters. To illustrate how the CPT model captures the individual differences of participants in the gambling task, Figure 5.4 shows the inferred subjective value functions and probability weighting functions for a few representative participants. These participants were chosen because they span the sorts of individual differences that we see in the data set.

The left panel shows the subjective value functions for three participants. The first participant shown by the dotted line has a relatively high value of λ but a low value of α. Consequently, their subjective value curve significantly undervalues the magnitude of both gains and losses, but still shows loss aversion in that the magnitude of losses weigh more than gains. The second participant shown by the dashed line has a relatively high value of both α and λ. This participant's subjective value curve also undervalues the magnitude of both gains and losses, but shows strong loss aversion. The subjective magnitude of losses are much larger than gains. Lastly, the third participant shown by the solid blue line has a relatively high value of α but λ is close to 1. Consequently, the effect of undervaluing the magnitude of both gains and losses are weakest compared to the other two participants. Furthermore, this effect is the same for both gains and losses because λ is close to 1.

The right panel shows the probability weighting functions for three different participants. The weighting functions for gains are determined by γ, and the weighting functions for losses

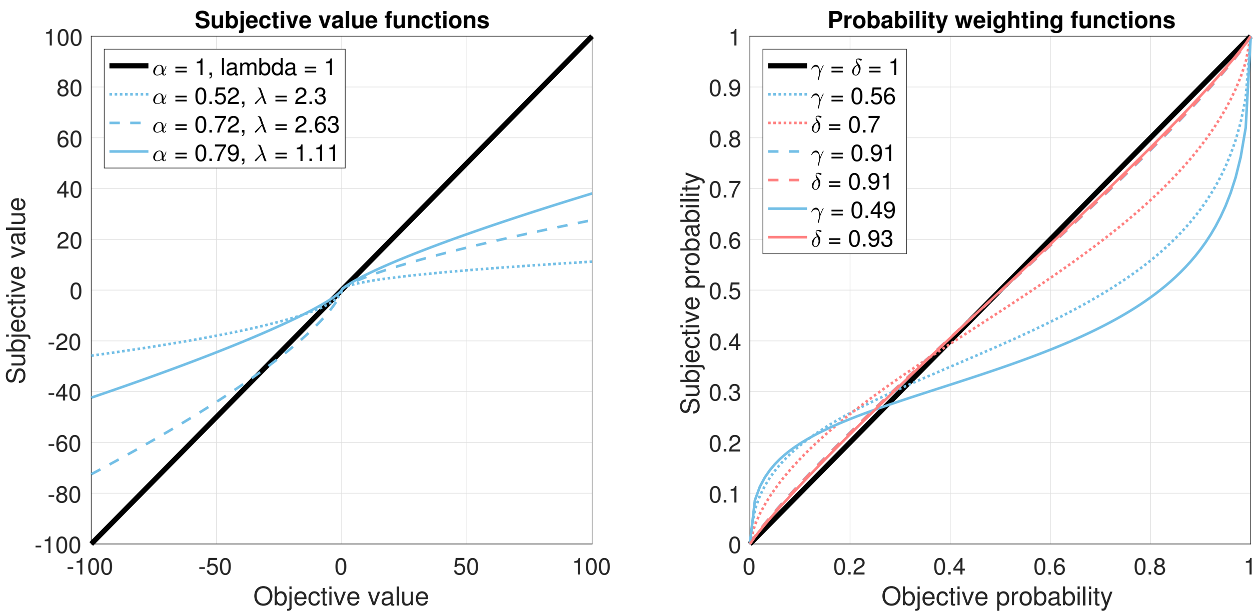

Figure 5.4: Inferred subjective value function curves and probability weighting function curves for a few representative participants. The three representative participants in the left panel span the sorts of individual differences we see in α and λ in the data set. The three representative participants in the right panel span the sorts of individual differences we see in γ and δ in the data set. Note: the three participants chosen in the left panel are not the same participants as the ones in the right panel.

are determined by δ. The same participant's weighting functions for gains and losses are shown in the same line type, while weighting functions for gains are shown in blue and weighting functions for losses are shown in red. The first participant shown by the dotted lines has relatively lower values of both γ and δ. Consequently, their weighting functions for both conditions overestimate smaller probabilities and underestimate larger probabilities. The second participant shown by the dashed lines has relatively high values of both γ and δ. Their probability weighting functions are extremely close to the $y = x$ line. Lastly, the third participant chosen is shown by the solid lines has a relatively low value of γ but high value of δ. They have very different weighting functions for gains versus losses. Their probability weighting function for gains significantly underestimates small probabilities and overestimates larger probabilities, but their weight function for losses is very close to the $y = x$ line.

However, the two parameters of interest from this model are the loss aversion λ and consistency ϕ per participant. Figure 5.5 shows the joint and marginal distributions of the λ and ϕ posterior expectations for each participant. First, it is clear that there is a wide range of individual differences in both loss aversion and consistency, with values ranging from about 0.5 to 2 for both parameters. Second, about one third of the participants exhibit the opposite of loss aversion, with λ values below 1. The remainder of the participants all show loss aversion with values of λ greater than 1. About one third of the participants exhibit relatively strong loss aversion with values of λ over 1.5. Furthermore, all values of phi are well above 0, suggesting that participants are generally consistent in their behavior (or at least more so than guessing). Lastly, there appears to be a negative relationship between the loss aversion and consistency parameters. Participants who tend to show stronger effects of loss aversion are also less consistent, while participants who show little to no loss aversion tend to be more consistent.

The chosen representative participants shown from the left panel of Figure 5.4 are labeled

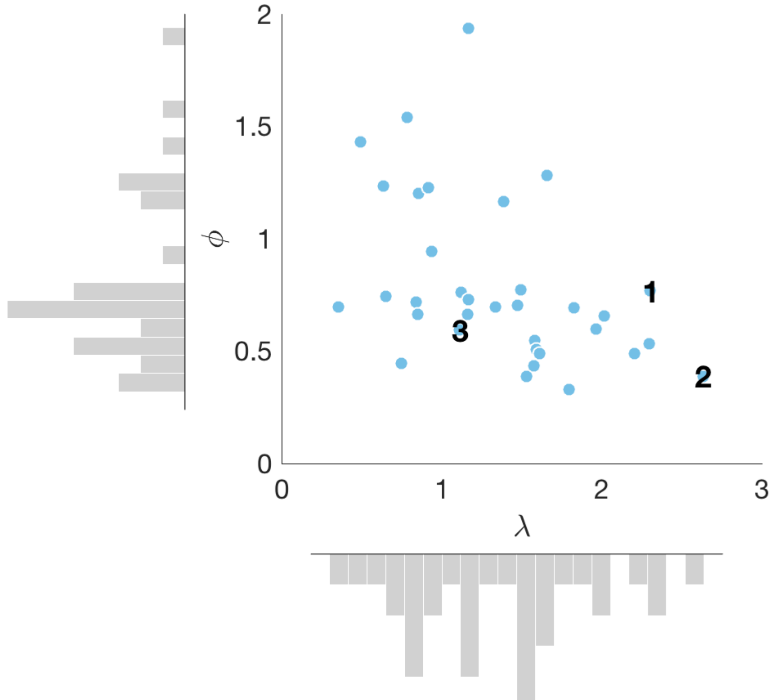

Figure 5.5: Joint and marginal distributions of the λ and ϕ posterior expectations for each participant. The chosen representative participants from 5.4 are labeled in the scatterplot.

in the scatterplot for reference. As seen before, the first two participants have high values of the loss aversion parameter λ. From this figure, it is evident that they are also relatively less consistent than the average participant. The third participant has a λ value around 1, but this figure reveals that they are also relatively inconsistent.

These results show that there are wide individual differences in both risk propensity, in the form of loss aversion, and consistency that are reflected in two parameters of the CPT model. In this way, the risk propensity and consistency of participants can be measured by how they make decisions in a gambling task involving gains and losses.

5.3.3 Descriptive Adequacy

A standard Bayesian approach based on the posterior predictive distribution was used to assess the ability of the model to fit the behavioral data (Gelman et al., 2004). The posterior predictive is the distribution of choices that the model expects, based on the inferred joint posterior distribution over the model parameters α, λ, γ, δ, and ϕ. The mode of the posterior predictive distribution for each participant on each problem was taken as the decision that the model predicts the participant to have made. The CPT model was able to predict about 77% of the decisions that participants made. Given that the chance level of agreement for choosing between two gambles is 50%, these results suggest that the CPT model provides a reasonable account of people's behavior in the gambling task.

It is important to note here that although the model is taken from Nilsson et al. (2011), there were a few differences in the current work. One key difference is that the original model was applied to the behavioral data from Rieskamp (2008), where there were three conditions instead of two. Rieskamp (2008) included an extra third "mixed" condition, in which participants were asked to choose between gambles that had gains and losses. Another key difference is that Rieskamp (2008) had a total of 180 problems per participants, with

60 problems in each of the three conditions. Participants in this study only completed 40 problems in each of the two conditions, for a total of 80 problems per participant. Perhaps if there were more data, the modeling results and performance might be affected. However, fewer conditions and problems were used in this study to reduce the effects of fatigue and boredom, since these data were collected online and not within a controlled laboratory setting.

Chapter 6

Self-Report Risk Surveys

The underlying idea behind this work is that if an individual has a fundamental propensity to take risks, this trait should be reflected in various tasks and surveys that are designed to measure risk propensity. Although not all of the four cognitive tasks are designed to measure risk propensity, the previous four chapters presented cognitive models for how people make decisions in these tasks that include psychological parameters interpretable as risk propensity. In addition to comparing those measures from the cognitive tasks to one another other, the goal is also to compare those risk parameters to the measures obtained from widely-used surveys designed to measure risk propensity.

Therefore, all 56 participants also completed 3 surveys: the Risk Propensity Scale (Meertens and Lion, 2008), Risk Taking Index (Nicholson et al., 2005), and Domain-Specific Risk-Taking scale Blais and Weber (2006). Each of these surveys were designed to assess an individual's general tendency to take risks but use various scales to do so. The next section describes in detail each of these surveys and empirical results.

6.1 Risk Propensity Scale

The Risk Propensity Scale (RPS), taken from Meertens and Lion (2008), was designed to be a short and easily administered test for measuring general risk-taking tendencies. The RPS originally consisted of only 9 items, from which 2 items were later removed. The version of the RPS used in this work consists of the 7 remaining items. All statements were rated on a 9-point scale ranging from 1 (*totally disagree*) to 9 (*totally agree*), except for the last statement, which was rated from 1 (*risk avoider*) to 9 (*risk seeking*). Items 1, 2, 3, and 4 were reverse-scored so that high scores represent high risk propensity. Meertens and Lion (2008) reported an internal reliability coefficient (Cronbach's α) of .77. The RPS survey is provided in Appendix A.1.

Participants indicated their selection by checking the appropriate box under each number. To obtain an overall RPS score for each participant, the mean of the seven items were taken (after reverse scoring 1, 2, 3, and 4). Figure 6.1 shows the distribution of RPS scores across all 56 participants. The RPS scores are right-skewed ranging from 1 to 8.14, with $M = 3.61$ and $SD = 1.86$. These results are different from Meertens and Lion (2008); the original study reported a mean score of 4.63 and standard deviation of 1.23.

6.2 Risk Taking Index

The Risk Taking Index (RTI), taken from Nicholson et al. (2005), was also designed to be a short and simple measure for assessing overall risk propensity in six domains: recreation, health, career, finance, safety, and social. There is only one item for each of the six domains. Each item was answered using a 5-point Likert scale ranging from 1 (*strongly disagree*) to 5 (*strong agree*). The RTI also takes into account past risk taking, so participants were asked to rate their risk taking in each of the six domains for both now and in the past. The RTI

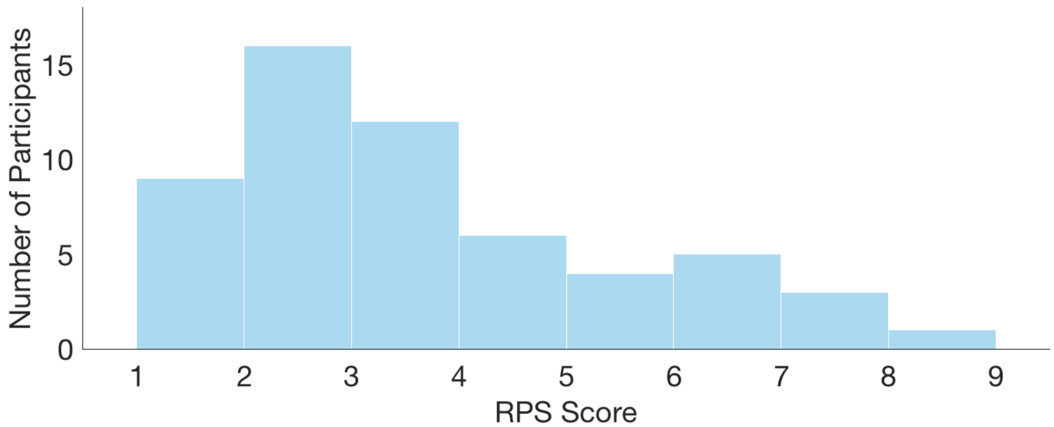

Figure 6.1: Distribution of RPS scores across participants ($M = 3.61$, $SD = 1.86$, min = 1, max = 8.14)

survey is provided in Appendix A.2.

Participants indicated their selection by checking the appropriate box under each number. To obtain an overall RTI score for each participant, the sum of each domain's response was taken across now and in the past. Then, the sum of each domain was taken as the overall RTI score. Therefore, RTI scores can potentially range from 12 to 60, where higher scores indicate higher risk propensity. Figure 6.2 shows the distribution of RTI scores across all 56 participants. The RTI scores appear to be bimodal, ranging from 12 to 42. There is a larger group of participants with a peak around 20, and another smaller group of participants with a peak around 35. These results are similar to Nicholson et al. (2005); the original study reported a mean score of 27.54 and standard deviation of 7.65.

6.3 Domain-Specific Risk-Taking Scale

The Domain-Specific Risk-Taking scale (DOSPERT) was originally developed by Weber et al. (2002) and later revised by Blais and Weber (2006) to be shorter and more applicable to a broader range of people. The original version was revised from 40 items down to 30 items

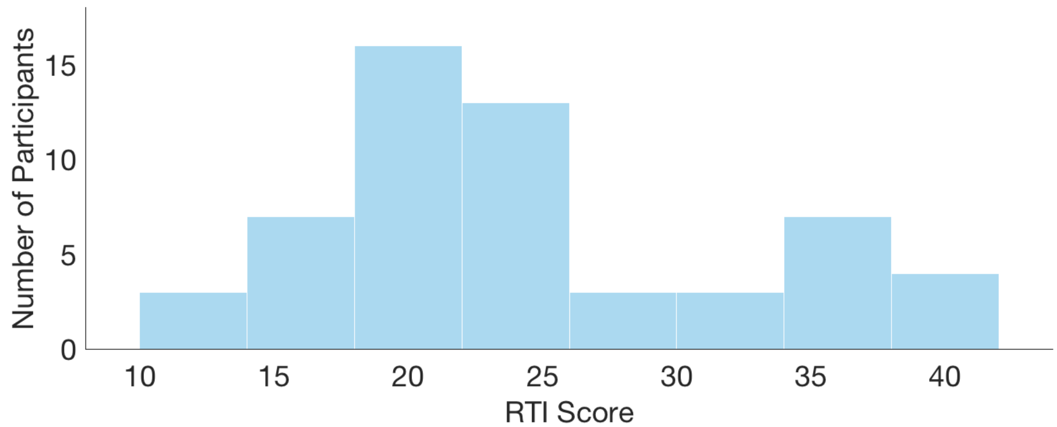

Figure 6.2: Distribution of RTI scores across participants (M = 24.25, SD = 8.08, min = 12, max = 42)

that evaluate risky behavioral intentions originating from five domains (ethical, financial, health/safety, social, and recreational risks). Each domain has six items.

The DOSPERT differs from the RPS and RTI in that it distinguishes people's tendency to be *risk seeking* from people's *risk perception.* Blais and Weber (2006) found a negative relationship between the two; people who tend to engage in more risk seeking behavior also tend to perceive situations as less risky, and vice versa. Therefore, the DOSPERT is split into two assessments: "risk taking" and "risk perception". Participants rated each of the 30 statements in terms of self-reported likelihood of engaging in risky behaviors as "risk taking", and in terms of their gut level assessment of the riskiness of these behaviors as "risk perception". In the risk taking assessment, they were asked to rate these 30 statements using a 7-point rating scale ranging from 1 (*Extremely Unlikely*) to 7 (*Extremely Likely*). In the risk perception assessment, they were asked to evaluate how risky each behavior is on a 7-point rating scale ranging from 1 (*Not at all*) to 7 (*Extremely Risky*). The DOSPERT survey is provided in Appendix A.3.

Participants indicated their selection by checking the appropriate box under each number. Ratings are added across all items of each domain to obtain five subscale scores for risk taking and five subscale scores for risk perception. The overall DOSPERT risk taking score is the mean of each subscale score for the risk taking assessment. Similarly, the overall DOSPERT risk perception score is the mean of each subscale score for the risk perception assessment. Therefore, each of the scores can potentially range from 6 to 42, where higher scores indicate higher risk propensity.

Figure 6.3 shows the scatterplot of risking taking and risk perception scores from the DOSPERT across all 56 participants, along with the marginal distributions of each as histograms on the side. The risk taking scores also appear to be slightly bimodal, with a large group of people centered around about 16-18 and then a smaller group near 30. However, risk perception scores are normally distributed, centered around 27. These results are consistent with the findings from Blais and Weber (2006), in that there is a negative relationship between risk taking and risk perception scores ($r = -0.22$).

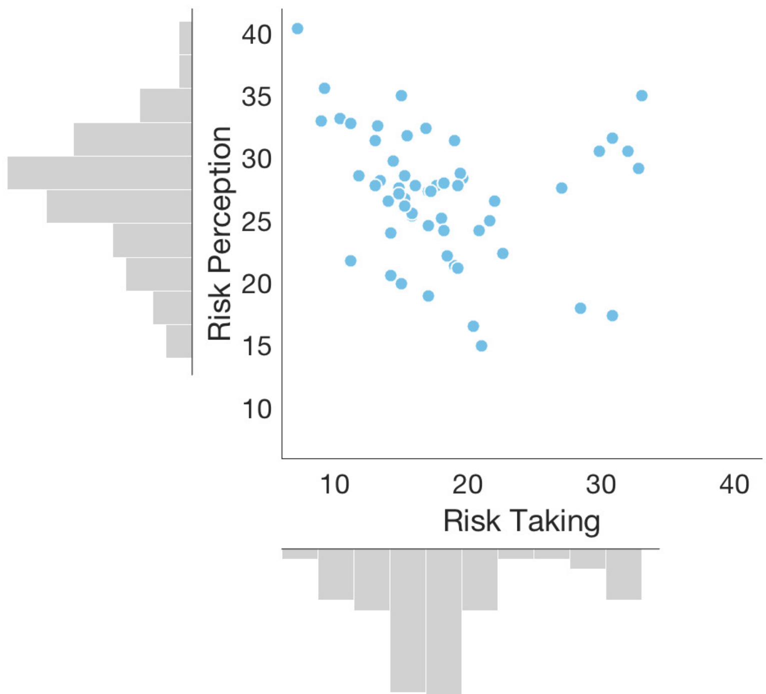

Figure 6.3: Scatterplot of DOSPERT risk taking versus risk perception scores, with marginal distributions (risk taking: M = 18.06, SD = 6.15, min = 7.2 max = 33; risk perception: M = 27.06, SD = 5.13, min = 15, max = 40.4). There is a negative correlation between risk taking and risk perception scores ($r = -0.22$).

Chapter 7

Correlations Analysis Across Tasks and Surveys

The goal of this work is to measure the risk propensity of individuals using risky decision-making tasks, as well as to examine the relationship between these tasks and cognitive models. Each of the cognitive models for the decision-making tasks have at least one parameter that relates to risk propensity, and at least one parameter that relates to consistency of behavior. This last section examines the relationship between the risk propensity and consistency parameters within and across tasks, through a Bayesian approach to estimating correlations.

Before looking at correlations of model parameters across tasks, the raw performance of participants within and across each cognitive task was compared. Figure 7.1 shows the correlations of participant performance across each condition in all of the cognitive tasks. The size of the circles represent the magnitude of Pearson's correlation r, while red circles represent negative correlations and blue circles represent positive correlations. These empirical results suggest that participant performance is highly correlated within tasks, and is also correlated across tasks although the correlations are not as strong.

7.0.1 Cognitive Task Overview

Table 7.1 shows an overview of the four cognitive tasks, models, and relevant parameters. The optimal stopping task has 8 total risk parameters ($\beta_{1:4}$, $\gamma_{1:4}$) and 4 total consistency parameters ($\alpha_{1:4}$). The BART has 2 total risk parameters ($\gamma^{+}_{1:2}$) and 2 total consistency parameters ($\beta_{1:2}$). The bandit task has 4 total risk parameters ($\gamma^{lose}_{1:4}$) and 4 total consistency parameters ($\gamma^{win}_{1:4}$). The gambling task has 1 risk parameter (λ) and 1 consistency parameter (ϕ). Consequently, there are 26 total relevant parameters from the cognitive tasks to be compared within and across tasks for each individual.

Task	**Conditions**	**Model**	**Risk parameter(s)**	**Consistency parameter**
Optimal Stopping	4	BFO	$\beta_1, \ldots, \beta_4, \gamma_1, \ldots, \gamma_4$	$\alpha_1, \ldots, \alpha_4$
BART	2	2-parameter	γ_1^+, γ_2^+	β_1, β_2
Bandit	4	e-WSLS	$\gamma_1^{lose}, \ldots, \gamma_4^{lose}$	$\gamma_1^{win}, \ldots, \gamma_4^{win}$
Gambling	2	CPT	λ	ϕ

Table 7.1: Overview of tasks and parameters.

Although each of the cognitive models provide a suitable account for how the behavioral data were generated in the risk tasks, as shown in the previous chapters, there is uncertainty around the measures of risk propensity and consistency. Since the measures were estimated in a Bayesian framework, the full posterior distribution of each risk and consistency parameter is available. The goal is to estimate the correlations between pairs of measures using Pearson's correlation coefficient r in a Bayesian framework, while taking into account the uncertainty of the risk and consistency measures inferred from each task.

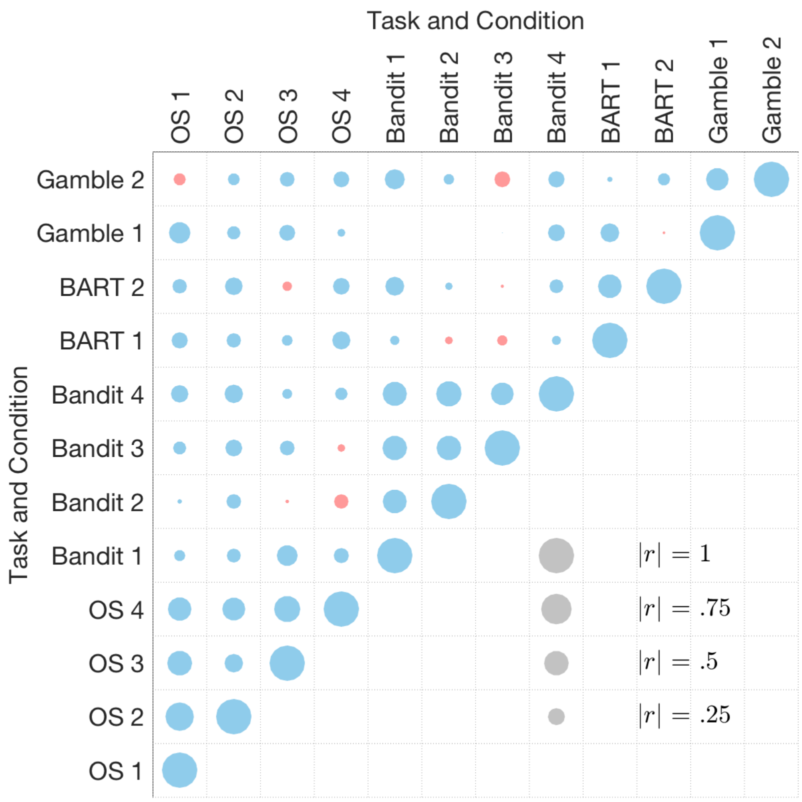

Figure 7.1: Correlations of performance across each condition in all of the cognitive tasks. Performance in the optimal stopping problem is computed as the proportion of problems where the participant correctly chose the maximum. Performance in the bandit task is computed as the average proportion of trials that resulted in reward,. Performance in the BART is computed as the average dollar amount collected on each problem. Lastly, performance in the gambling task is computed as the proportion of problems where the participant chose the better gamble. The red circles represent negative correlations, while the blue circles represent positive correlations. The area of the circles correspond to the magnitude of Pearson's correlation r.

7.1 Estimating Correlations with Uncertainty

The correlations between each risk and consistency parameter from all of the four cognitive tasks are estimated in a Bayesian approach, based on Lee and Wagenmakers (2013, chap. 5). This correlation model assumes that although we have risk and consistency measures from each of the cognitive models, there is uncertainty about the exact values of these parameters. For example, it is likely that response times are measured very accurately because it is a physical quantity with measurement tools, but risk propensity measures from these cognitive models are likely to be less accurate. Therefore, when estimating correlations between variables that may not be accurately measured, it is important to account for the uncertainty in measurement.

The correlation model assumes that the observed data for each participant are the model estimates of, say, a risk propensity and a consistency parameter, denoted by $\boldsymbol{x_i} = (x_{i1}, x_{i2})$. However, these observations are now sampled from a Gaussian distribution, centered on some latent true risk propensity and consistency of that individual $\boldsymbol{y_i} = (y_{i1}, y_{i2})$:

$$x_{ij} \sim \text{Gaussian}\big(y_{ij}, \lambda_j^e\big) \tag{7.1}$$

where the precision of the measurements is captured by $\boldsymbol{\lambda^e} = (\lambda_1^e, \lambda_2^e)$. In this version of the correlation model, these measure precisions are assumed to be known and are simply taken to be the precisions of the posterior distributions of the risk and consistency parameters. The latent true values of these cognitive measures, $\boldsymbol{y_i}$, are then modeled as a draw from a multivariate Gaussian distribution:

$$\boldsymbol{y_i} \sim \text{Gaussian}\left((\mu_1, \mu_2), \begin{bmatrix} \sigma_1^2 & r\sigma_1\sigma_2 \\ r\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}^{-1} \right) \tag{7.2}$$

The remainder of the Bayesian correlation model is completed by setting priors on r, σ_1, σ_2,

μ_1, and μ_2:

$$r \sim \text{Uniform}\big(-1, 1\big) \tag{7.3}$$

$$\sigma_1^2, \sigma_2^2 \sim \text{InvGamma}\big(0.001, 0.001\big) \tag{7.4}$$

$$\mu_1, \mu_2 \sim \begin{cases} \text{Uniform}\big(0, 1\big) & \text{, for OS } \alpha \text{, Bandit } \gamma^{win}, \gamma^{lose} \\ \text{Uniform}\big(0, 10\big) & \text{, for BART } \beta, \gamma^{+} \\ \text{Gaussian}\big(0, 0.001\big) & \text{, for OS } \beta, \gamma \\ \text{Uniform}\big(1, 9\big) & \text{, for RPS} \\ \text{Uniform}\big(10, 50\big) & \text{, for RTI, DOSPERT} \end{cases} \tag{7.5}$$

The correlation model was implemented as a Bayesian graphical model in JAGS. All modeling results are based on four chains of 1000 samples each, collected after 2000 discarded burn-in samples. The chains were verified for convergence using the standard $\hat{R}$ statistic (Brooks and Gelman, 1997).

In addition to looking the inferred correlations between each of the pairwise comparisons, each comparison also fundamentally involves model selection. It is important to know if there exists a correlation between two risk measures, or if there is no such correlation in either direction. This means that for each pairwise comparison, we should also ask if a model that assumes no correlation is better than a model that allows for some magnitude of correlation in either direction. To address this model selection component, a simple approximation for computing Bayes factors was used, based on the Savage-Dickey method (Wetzels et al., 2010). The Savage-Dickey method relies only on the prior and posterior distributions of Pearson's r from the correlation model. For each comparison, the ratio of the posterior to prior density of r at the critical point of zero is calculated. This critical point of zero corresponds to the value at which the model that assumes a correlation exists reduces to the simpler nested model that assumes no correlation. The Savage-Dickey result is that the Bayes factor between nested models is given by the ratio of the posterior to prior densities at the critical point. The Bayes factor for the null model M_0, in which there is no correlation between the two measures of comparison, versus an alternative model M_1, in which there is a non-zero correlation between the two measures, is given by:

$$b_{01} = \frac{p\left(D \mid M_0\right)}{p\left(D \mid M_1\right)} = \frac{p\left(\delta_k = 0 \mid D, M_1\right)}{p\left(\delta_k = 0 \mid M_1\right)} \tag{7.6}$$

Although in this formulation, a negative log Bayes factor corresponds to there being evidence for a nonzero correlation, the remainder of the analyses take the opposite sign of the log Bayes factor so that a positive log Bayes factor now corresponds to evidence for a nonzero correlation between two measures.

7.2 Correlation Results

As an example of some pairwise comparisons from the correlation analysis, Figure 7.2 looks at four handpicked parameters. These four parameters are the risk parameters from the first condition in each of the behavioral tasks. The lower triangle shows scatterplots of the posterior expectations of each parameter against one another, while the upper triangle shows the posterior distribution of Pearson's r from the correlation analysis for each comparison. As an example, the inferred correlation between OS β and BART γ^+ is around 0, while the inferred correlation between Gamble λ and Bandit γ^{lose} is relatively uncertain but trending towards a positive relationship. An inferred correlation coefficient is obtained for each pairwise comparison between risk and consistency measures, as illustrated in this figure.

Across the 4 cognitive tasks and 3 risk questionnaires, there are a total of 30 risk and consistency measures to be compared, for a total of 435 pairwise correlations. These 30 risk and consistency measures are:

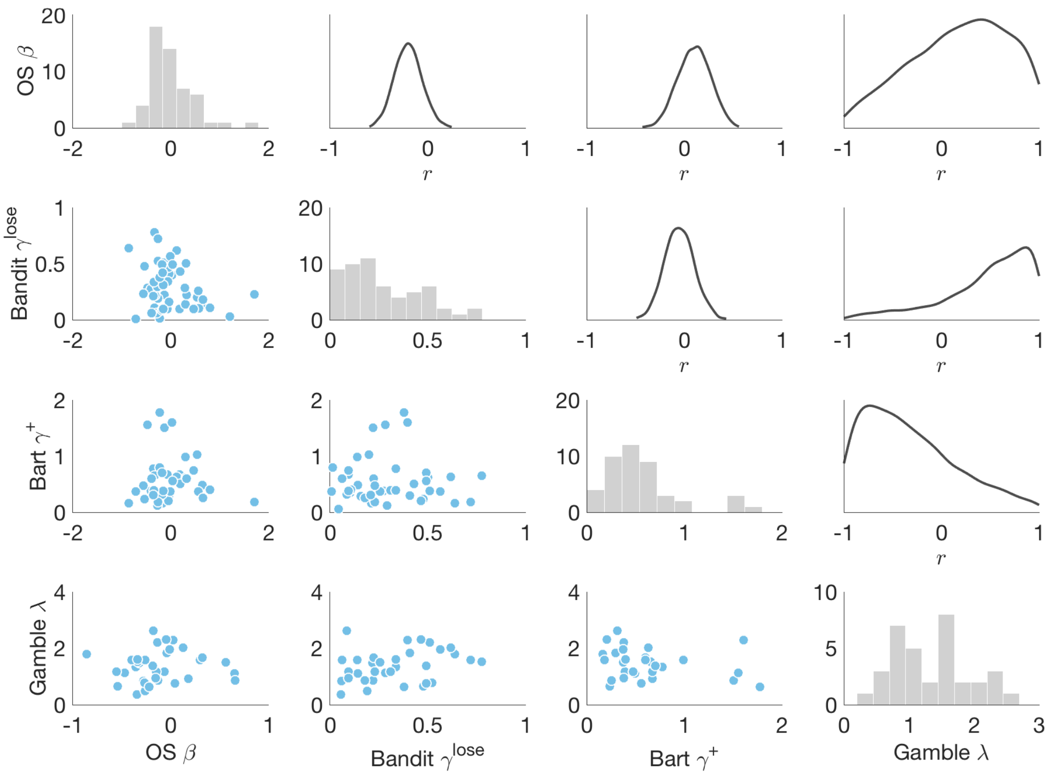

Figure 7.2: Pairwise comparisons for four example risk parameters. The lower triangle shows scatterplots of posterior expectations for each risk parameter. The distributions of posterior expectations for each participant are shown along the diagonal. The upper triangle shows the posterior distribution of Pearson's r from the correlation model for the corresponding pairwise comparison.

- Risk measures:
  - Optimal Stopping $\beta_1, \ldots, \beta_4$ (4)
  - Optimal Stopping $\gamma_1, \ldots, \gamma_4$ (4)
  - Bandit $\gamma_1^{lose}, \ldots, \gamma_4^{lose}$ (4)
  - BART γ_1^+, γ_2^+ (2)
  - Gamble λ (1)
  - RPS score (1)
  - RTI score (1)
  - DOSPERT Risk Taking (RT) score (1)
  - DOSPERT Risk Perception (RP) score (1)
- Consistency measures:
  - Optimal Stopping $\alpha_1, \ldots, \alpha_4$ (4)
  - Bandit $\gamma_1^{win}, \ldots, \gamma_4^{win}$ (4)
  - BART β_1, β_2 (2)
  - Gamble ϕ (1)

If the psychological construct of risk propensity, or an individual's tendency to take risks, is measured by these parameters in the tasks and questionnaires, then these risk parameters and scores are expected to correlate with one another. Figure 7.3 shows the results from the correlation analysis looking at the 435 pairwise comparisons. The area of the circles correspond to the magnitude of the correlation, taken as the posterior expectation of Pearson's r. The red circles indicate that the correlation is negative, while the blue circles indicate that the correlation is positive. The empty squares without any circles correspond to those comparisons where the log Bayes factor was negative (showing evidence for there being no correlation). The gray dashed lines divide the grid into the four cognitive tasks and risk questionnaires.

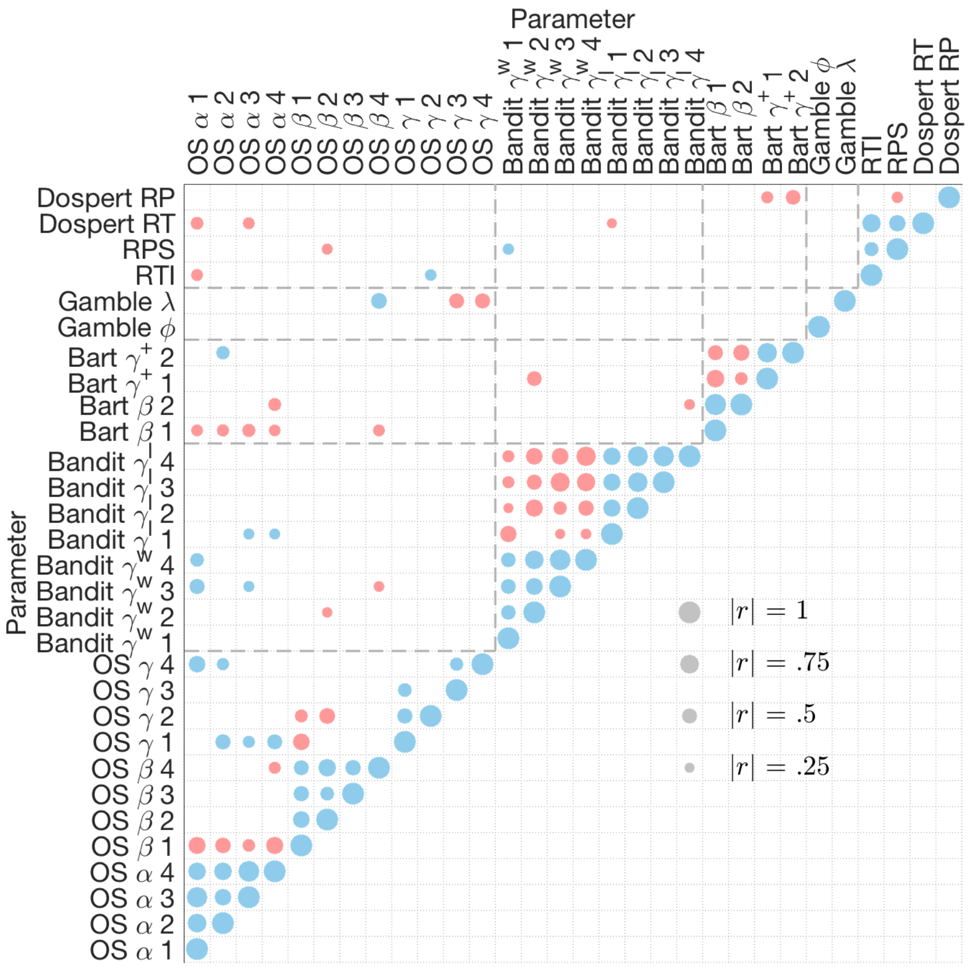

Figure 7.3: Correlation matrix of the risk and consistency parameters. The red circles represent negative correlations, while the blue circles represent positive correlations. The area of the circles correspond to the absolute value of the posterior expectation of Pearson's correlation r. Empty boxes indicate that the log Bayes factor was negative, in favor of the null model where there is no correlation. The parameters within tasks are boxed out by the dashed gray lines. OS α represents consistency in the optimal stopping problem and OS β and γ represent risk in the optimal stopping problem. Bandit γ^w represents consistency in the bandit task and Bandit γ^l represents risk. Bart γ^+ represents risk in the BART task and Bart β represents consistency. Gamble ϕ represents consistency and Gamble λ represents risk.

First, it is evident that there are positive correlations between the same parameters within tasks. For example, all of the consistency parameters across conditions from optimal stopping are highly correlated, as well as the risk parameters within the BART. This is evident by the triangles of blue circles along the $y = x$ axis. The positive correlations across conditions within the same task are expected, since the tasks are still similar (but just varying in length or environment, for example) so the psychological constructs measured across the different conditions are more likely to be stable. Furthermore, the RTI, RPS, and DOSPERT RT are also positively correlated with each other, replicating findings in the literature.

There also appears to be some negative correlations between different parameters within tasks. For example, the γ^{win} and γ^{lose} parameters in the bandit task are negatively correlated with each other, and the risk and consistency parameters in the BART are also negatively correlated with each other. There is some tradeoff between parameters in some of these tasks, which was also evident in the independent task analyses from the previous chapters.

However, there appears to be less evidence for systematic correlations in either direction across tasks. If risk propensity is a psychological construct that can be captured in these cognitive models by the risk parameters, then these parameters should correlate with one another and the risk questionnaire scores. There does not appear to be any systematic positive correlations between the risk parameters or between the consistency parameters. Figure 7.4 shows the log Bayes factors of r against zero for just the risk measures from tasks and questionnaires. The hollow circles indicate comparisons where the log Bayes factor was less than -1, in favor of the null model. The filled-in gray circles indicate comparisons where the log Bayes factor was greater than 1, in favor of the model where a nonzero correlation exists. The empty boxes indicate inconclusive evidence in either direction ($-1 < |r| < 1$). It is clear here that within tasks and parameters, there is evidence for a nonzero correlation. However, across tasks and parameters (the areas boxed in by the dashed gray lines) there is either evidence in favor of the null model or inconclusive evidence. This further supports the results seen in Figure 7.3, suggesting that there is no systematic relationships between the risk measures across tasks and questionnaires.

To take a closer look at those comparisons where there is evidence for nonzero correlations, Figure 7.5 shows results only from those comparisons where the log Bayes factors were greater than 6, in favor of the alternative model where there is a nonzero correlation. The cutoff of a log Bayes factor of 6 is used because it is a standard interpretative boundary corresponding to "strong" evidence (Kass and Raftery, 1995, p. 777). The left panel shows the 95% Bayesian credible intervals of r for each comparison. The right panel shows the log Bayes factors for the corresponding comparisons. The parameters with strong positive correlations are the same parameters across conditions within tasks. There are also some tradeoffs between different parameters within tasks, shown by the negative correlations. Furthermore, two of the three risk questionnaires appear here as well and are positive correlated with each other, with the strongest correlation between RTI and DOSPERT Risk Taking.

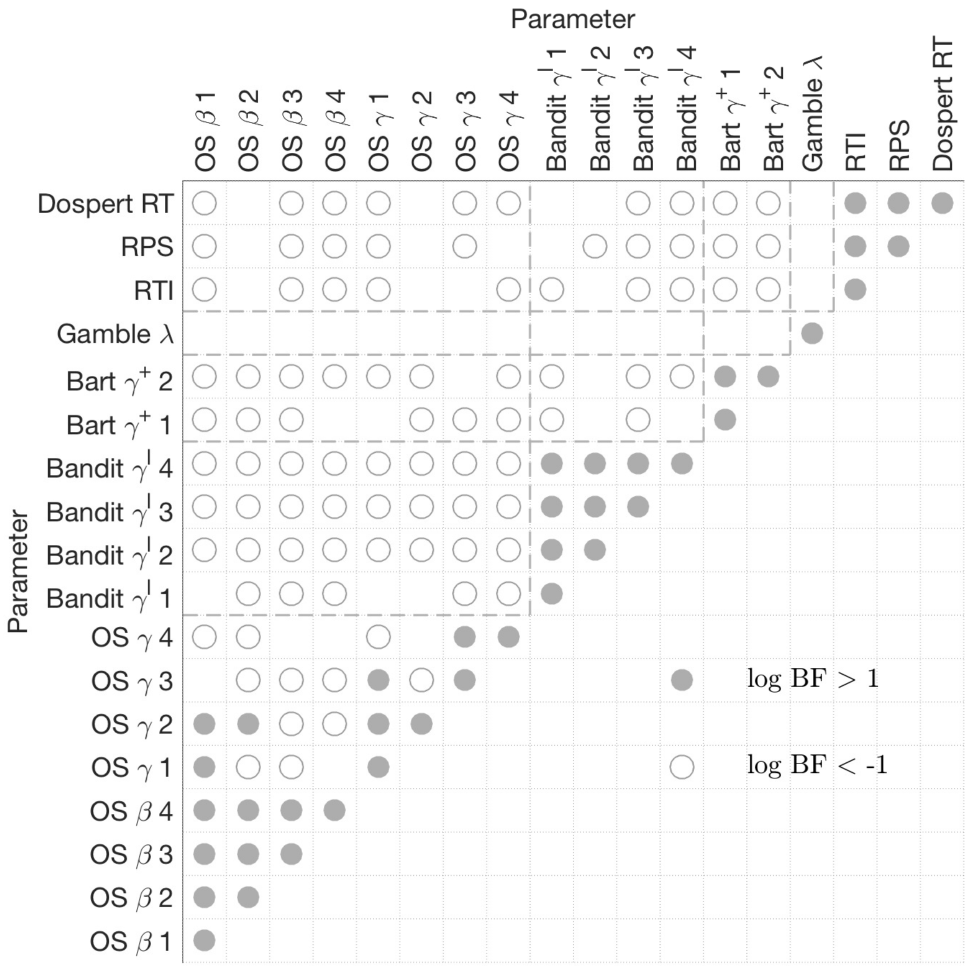

Figure 7.4: Log Bayes factors of r against zero for the risk measures. The hollow circles indicate comparisons where the log Bayes factor was less than -1, in favor of the null model. The filled-in gray circles indicate comparisons where the log Bayes factor was greater than 1, in favor of the model where a nonzero correlation exists. The empty boxes indicate inconclusive evidence in either direction ($-1 < |r| < 1$). The parameters within tasks are boxed out by the dashed gray lines. OS α represents consistency in the optimal stopping problem and OS β and γ represent risk in the optimal stopping problem. Bandit γ^w represents consistency in the bandit task and Bandit γ^l represents risk. Bart γ^+ represents risk in the BART task and Bart β represents consistency. Gamble ϕ represents consistency and Gamble λ represents risk.

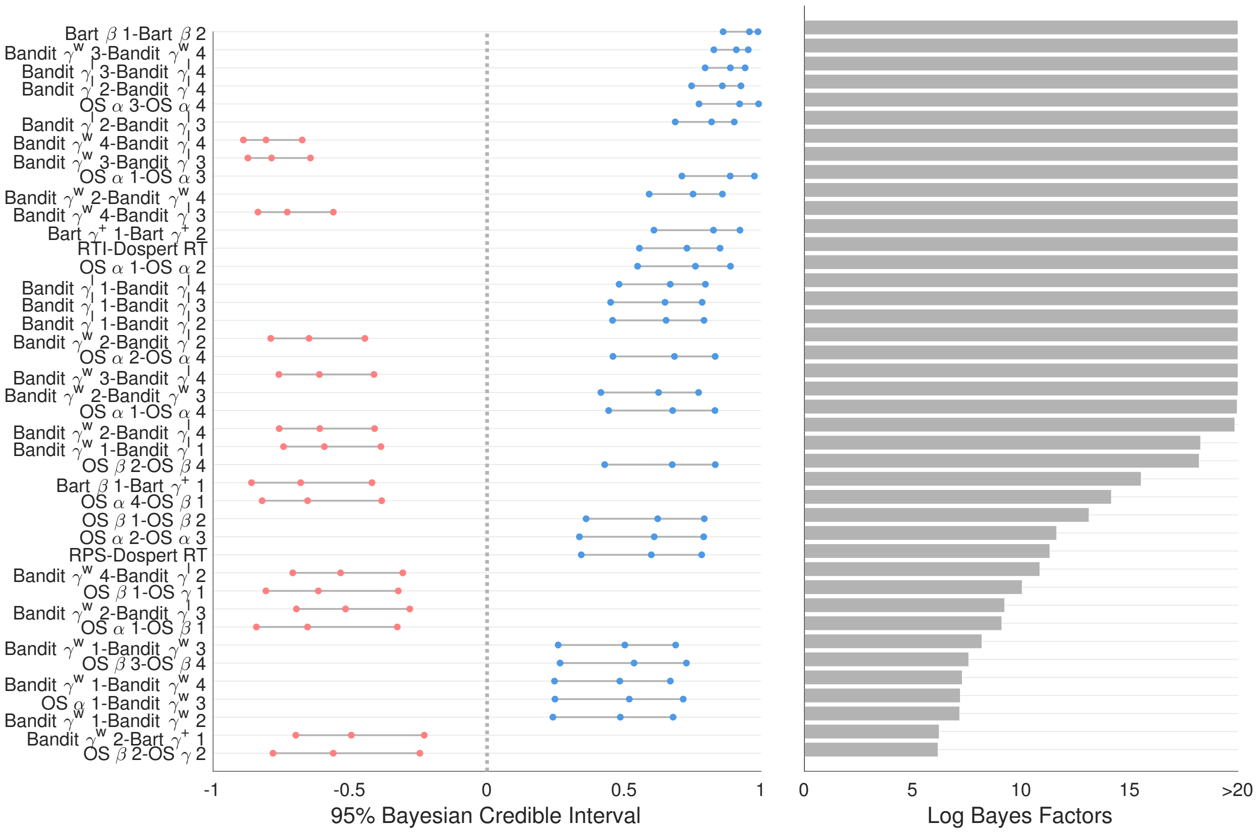

Figure 7.5: 95% Bayesian credible intervals and log Bayes factors for Pearson's correlation coefficient r. Out of the 435 pairs of correlations, only the 41 correlations where log BFs were greater than 6 in favor of the alternative nonzero correlation model are shown here. The pairwise comparisons between parameters are sorted in descending order from highest to lowest log Bayes factor.

Chapter 8

Discussion

This thesis looks at the psychological construct of risk propensity, and compares measures of risk across sequential decision-making tasks and more traditional self-report risk surveys. The first four chapters detailed how risk propensity can potentially be measured in optimal stopping problems, bandit problems, the BART, and a preferential choice gamble task through cognitive models of human behavior. In each of the independent task analyses, the cognitive models provided adequate accounts of participants' behavior and captured stable individual differences across participants in parameters that can be interpreted as risk propensity and consistency. In the RPS, RTI, and DOSPERT independent analyses, the results were consistent with previous studies with similar means and standard deviations of the scores. There was also a negative correlation found between DOSPERT Risk Taking and Risk Perception, supporting the findings from the original study (Blais and Weber, 2006).

If risk propensity is a stable psychological construct that can be measured by these cognitive tasks and risk questionnaires, then the risk parameters and scores are expected to correlate across tasks and surveys. In line with previous research, the correlation analysis found evidence for positive correlations between the risk scores obtained from the three self-report risk surveys. The RTI, RPS, and DOSPERT Risk Taking scores were positively correlated with each other. The results also found stable individual differences in the psychological variables within tasks, with risk and consistency parameters being highly correlated across experimental conditions within cognitive tasks. However, the correlation analysis failed to find evidence for systematic relationships across tasks and surveys.

Similar findings of little empirical relationship between measures from behavioral tasks and self-report surveys has been found in psychological research on individual differences in other domains such as self-regulation (Eisenberg et al., 2018) and intelligence (Friedman et al., 2006). In the case of this thesis, the cognitive tasks are similar to each other in that they are sequential decision-making tasks involving risk and uncertainty. However, there are still major differences between them, and the psychological variables in these cognitive process models may not directly capture risk propensity and consistency in the same way. The optimal stopping problem involves holding out until a desirable option comes along, but the value of each option is presented to the decision maker explicitly. In contrast, the preferential choice gambling task requires the decision maker to make judgments based on both the value of each option and probabilities associated with those values, without explicitly stating the expected reward from each gamble. The bandit problem gives feedback after each decision is made, explicitly showing the number of rewards and failures of each slot machine on the screen. However, the BART does not provide any feedback to the decision maker throughout a problem while they are pumping (or even after they bank the amount) unless the balloon bursts. Each of these decision-making tasks require related but different cognitive processes. It is entirely plausible that suboptimality in the optimal stopping problem does not translate directly to loss aversion in the gambling task. Similarly, the tendency to pump a balloon more with the risk of losing it all in the BART might not be the same thing psychologically as balancing explore and exploitation in the bandit task.

While the results in this thesis failed to find evidence for correlations across tasks, the self-report survey scores correlated with one another. However, this result may be due to the nature of the self-report surveys themselves. The survey questions all rely on self reports of risk-taking behavior and risk perceptions, and use similar scales for assessment. In addition, self-report surveys are susceptible to biases and dependent on an individual's knowledge and perception of their own behavior (Randall and Fernandes, 1991; Furnham and Henderson, 1982), in a way that cognitive tasks avoid. Therefore, the survey scores can simply be seen as measurements of people's perceived risk-taking behavior, rather than measurements of the latent psychological construct risk propensity that influences their risk-taking behavior.

8.0.1 Limitations and Future Directions

One limitation is the relatively small sample size collected. Generally, studies on individual differences across a number of behavioral tasks and surveys use larger sample sizes, most over 100 participants with some studies recruiting over 1000 participants (Eisenberg et al., 2018). Furthermore, the preferential choice gambling task can be extended to include more problems per participant. The current experiment only had participants complete 2 conditions of 40 problems each with fatigue and boredom effects in mind, but is much fewer than the original experiment (Nilsson et al., 2011). Consequently, the posterior distributions for risk and consistency in the gambling task were relatively uncertain and the correlation analyses involving the gambling task were inconclusive.

Another limitation is the correlational analysis framework, because it only explore linear relationships between pairwise comparisons of model parameters and survey scores. Frey et al. (2017) looked at risk preference in a psychometric framework, and also found correlations between risk propensity and behavioral measures to be weak. However, they found a general factor of risk preference that generalized to real-world risky activities. The current correlation analysis does not take into account how each risk parameter from a cognitive model contributes to the total risk propensity in a particular task, and is unable to infer higher-order latent traits such as a general factor of risk propensity. Risk propensity is a complex psychological construct that may not be stable across all situations, contexts, and environments. Future work should focus on simultaneously analyzing the behavioral and psychometric data from tasks and surveys through cognitive latent variable models (Vandekerckhove, 2014; Pe et al., 2013). Cognitive latent variable models (CLVMs) are a broad category of models that involves a latent variable structure built on top of cognitive process models, to allow inference of latent variables that have higher-order cognitive interpretations. By using a CLVM framework, conclusions can be drawn about latent psychological traits that affect the risk propensity and consistency in certain behavioral tasks and risk surveys. It is possible to find a higher-order general risk propensity trait that perhaps affects the risk parameters in certain tasks more than others. This same general risk propensity might also affect some or all of the self-report survey scores. It is also possible to see if there is some level of risk propensity that is measured in the risk surveys, but not captured by the risk parameters from cognitive models.

8.0.2 Conclusion

This thesis explored four cognitive models for four sequential decision-making tasks involving risk and uncertainty, and found stable individual differences in risk and consistency psychological variables within tasks. However, when compared to widely-used self-report survey scores and across tasks, there is less evidence for any meaningful relationships between tasks and surveys. These results contribute to the discussion about how cognitive process models of sequential decision-making tasks can be used to measure risk, and whether risk propensity is a stable psychological construct that can be measured by these behavioral tasks and surveys. Although the correlation analysis does not find evidence for relationships between risk measures across tasks and surveys, these results challenge the idea that risk propensity is a coherent construct and provide a basis for future research on the measurement of risk propensity in behavioral tasks and self-report surveys.

Bibliography

Aklin, W. M., Lejuez, C., Zvolensky, M. J., Kahler, C. W., and Gwadz, M. (2005). Evaluation of behavioral measures of risk taking propensity with inner city adolescents. *Behaviour research and therapy*, 43(2):215–228.

Bearden, J. N., Rapoport, A., and Murphy, R. O. (2006). Sequential observation and selection with rank-dependent payoffs: An experimental study. *Management Science*, 52:1437–1449.

Blais, A.-R. and Weber, E. U. (2006). A domain-specific risk-taking (dospert) scale for adult populations.

Brandstätter, E., Gigerenzer, G., and Hertwig, R. (2006). The priority heuristic: making choices without trade-offs. *Psychological review*, 113(2):409.

Brooks, S. P. and Gelman, A. (1997). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7:434–455.

Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, 100(3):432.

Campbell, J. and Lee, M. D. (2006). The effect of feedback and financial reward on human performance solving 'secretary' problems. In Sun, R., editor, *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, pages 1068–1073, Mahwah, NJ. Erlbaum.

Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., and Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095):876.

De Martino, B., Kumaran, D., Seymour, B., and Dolan, R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, 313(5787):684–687.

Dunlop, S. M. and Romer, D. (2010). Adolescent and young adult crash risk: Sensation seeking, substance use propensity and substance use behaviors. *Journal of Adolescent Health*, 46(1):90–92.

Eisenberg, I. W., Bissett, P. G., Enkavi, A. Z., Li, J., MacKinnon, D., Marsch, L., and Poldrack, R. (2018). Uncovering mental structure through data-driven ontology discovery. *PsyArXiv*.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, 4:282–296.

Frey, R., Pedroni, A., Mata, R., Rieskamp, J., and Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science advances*, 3(10):e1701381.

Friedman, N. P., Miyake, A., Corley, R. P., Young, S. E., DeFries, J. C., and Hewitt, J. K. (2006). Not all executive functions are related to intelligence. *Psychological science*, 17(2):172–179.

Furnham, A. and Henderson, M. (1982). The good, the bad and the mad: Response bias in self-report measures. *Personality and Individual Differences*, 3(3):311–320.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2004). *Bayesian Data Analysis*. Chapman & Hall/CRC, Boca Raton, FL, second edition.

Gilbert, J. P. and Mosteller, F. (1966). Recognizing the maximum of a sequence. *American Statistical Association Journal*, 61:35–73.

Goldstein, D. G., McAfee, R. P., Suri, S., and Wright, J. R. (2007). Learning when to stop searching. *Management Science*.

Guan, M. and Lee, M. D. (2014). Threshold models of human decision making on optimal stopping problems in different environments. In Bello, P., Mcshane, M., and Scassellati, B., editors, *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, pages 553–558, Austin, TX: Cognitive Science Society.

Guan, M. and Lee, M. D. (2017). The effect of goals and environments on human performance in optimal stopping problems.

Guan, M., Lee, M. D., and Silva, A. (2014). Threshold models of human decision making on optimal stopping problems in different environments. In *Proceedings of the 36th annual meeting of the cognitive science society. Austin, TX: Cognitive Science Society*.

Guan, M., Lee, M. D., and Vandekerckhove, J. (2015). A hierarchical cognitive threshold model of human decision making on different length optimal stopping problems. In *Proceedings of the 37th annual meeting of the cognitive science society. Austin, TX: Cognitive Science Society*.

Harrison, J. D., Young, J. M., Butow, P., Salkeld, G., and Solomon, M. J. (2005). Is it worth the risk? a systematic review of instruments that measure risk propensity for use in the health setting. *Social science & medicine*, 60(6):1385–1396.

Hopko, D. R., Lejuez, C., Daughters, S. B., Aklin, W. M., Osborne, A., Simmons, B. L., and Strong, D. R. (2006). Construct validity of the balloon analogue risk task (bart): Relationship with mdma use by inner-city drug users in residential treatment. *Journal of Psychopathology and Behavioral Assessment*, 28(2):95–101.

Hunt, M. K., Hopko, D. R., Bare, R., Lejuez, C., and Robinson, E. (2005). Construct validity of the balloon analog risk task (bart) associations with psychopathy and impulsivity. *Assessment*, 12(4):416–428.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pages 263–291.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90:377–395.

Kathleen Holmes, M., Bearden, C. E., Barguil, M., Fonseca, M., Serap Monkul, E., Nery, F. G., Soares, J. C., Mintz, J., and Glahn, D. C. (2009). Conceptualizing impulsivity and risk taking in bipolar disorder: importance of history of alcohol abuse. *Bipolar disorders*, 11(1):33–40.

Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior and Organization*, 14:381–392.

Lauriola, M., Panno, A., Levin, I. P., and Lejuez, C. W. (2014). Individual differences in risky decision making: A meta-analysis of sensation seeking and impulsivity with the balloon analogue risk task. *Journal of Behavioral Decision Making*, 27(1):20–36.

Lee, M. D. (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, 30:555–580.

Lee, M. D. and Wagenmakers, E.-J. (2013). *Bayesian cognitive modeling: A practical course.* Cambridge University Press.

Lee, M. D., Zhang, S., Munro, M. N., and Steyvers, M. (2011). Psychological models of human and optimal performance on bandit problems. *Cognitive Systems Research*, 12:164–174.

Lejuez, C., Aklin, W. M., Jones, H. A., Richards, J. B., Strong, D. R., Kahler, C. W., and Read, J. P. (2003). The balloon analogue risk task (bart) differentiates smokers and nonsmokers. *Experimental and clinical psychopharmacology*, 11(1):26.

Lejuez, C., Simmons, B. L., Aklin, W. M., Daughters, S. B., and Dvir, S. (2004). Risk-taking propensity and risky sexual behavior of individuals in residential substance use treatment. *Addictive behaviors*, 29(8):1643–1647.

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., Strong, D. R., and Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: the balloon analogue risk task (bart). *Journal of Experimental Psychology: Applied*, 8(2):75.

Lighthall, N. R., Mather, M., and Gorlick, M. A. (2009). Acute stress increases sex differences in risk seeking in the balloon analogue risk task. *PLoS One*, 4(7):e6002.

Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The economic journal*, 92(368):805–824.

Meertens, R. M. and Lion, R. (2008). Measuring an individual's tendency to take risks: The risk propensity scale 1. *Journal of Applied Social Psychology*, 38(6):1506–1520.

Mishra, S., Lalumière, M. L., and Williams, R. J. (2010). Gambling as a form of risk-taking: Individual differences in personality, risk-accepting attitudes, and behavioral preferences for risk. *Personality and Individual Differences*, 49(6):616–621.

Nicholson, N., Soane, E., Fenton-O'Creevy, M., and Willman, P. (2005). Personality and domain-specific risk taking. *Journal of Risk Research*, 8(2):157–176.

Nilsson, H., Rieskamp, J., and Wagenmakers, E.-J. (2011). Hierarchical bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, 55(1):84–93.

Pe, M., Vandekerckhove, J., and Kuppens, P. (2013). A diffusion model account of the relationship between the emotional flanker task and depression and rumination. *Emotion*, 13:739–747.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In Hornik, K., Leisch, F., and Zeileis, A., editors, *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. Vienna, Austria.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323–343.

Randall, D. M. and Fernandes, M. F. (1991). The social desirability response bias in ethics research. *Journal of business ethics*, 10(11):805–817.

Rao, H., Korczykowski, M., Pluta, J., Hoang, A., and Detre, J. A. (2008). Neural correlates of voluntary and involuntary risk taking in the human brain: an fmri study of the balloon analog risk task (bart). *Neuroimage*, 42(2):902–910.

Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34:1446–1465.

Rieskamp, J., Busemeyer, J. R., and Mellers, B. A. (2006). Extending the bounds of rationality: Evidence and theories of preferential choice. *Journal of Economic Literature*, 44(3):631–661.

Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535.

Russo, J. E. and Dosher, B. A. (1983). Strategies for multiattribute binary choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(4):676.

Seale, D. A. and Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the "Secretary Problem". *Organizational Behavior and Human Decision Processes*, 69:221–236.

Seale, D. A. and Rapoport, A. (2000). Optimal stopping behavior with relative ranks. *Journal of Behavioral Decision Making*, 13:391–411.

Shu, S. B. (2008). Future-biased search: the quest for the ideal. *Journal of Behavioral Decision Making*, 21:352–377.

Sitkin, S. B. and Weingart, L. R. (1995). Determinants of risky decision-making behavior: A test of the mediating role of risk perceptions and propensity. *Academy of management Journal*, 38(6):1573–1592.

Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior and Organization*, 35:309–332.

Stewart Jr, W. H. and Roth, P. L. (2001). Risk propensity differences between entrepreneurs and managers: A meta-analytic review. *Journal of applied psychology*, 86(1):145.

Steyvers, M., Lee, M. D., and Wagenmakers, E.-J. (2009). A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*. The MIT press.

Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211:453–458.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323.

van Ravenzwaaij, D., Dutilh, G., and Wagenmakers, E.-J. (2011). Cognitive model decomposition of the bart: Assessment and application. *Journal of Mathematical Psychology*, 55(1):94–105.

Vandekerckhove, J. (2014). A cognitive latent variable model for the simultaneous analysis of behavioral and personality data. *Journal of Mathematical Psychology*, 60:58–71.

Von Neumann, J. and Morgenstern, O. (1947). Theory of games and economic behavior, 2nd rev.

Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., and Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the savage–dickey method. *Cognitive psychology*, 60(3):158–189.

Wallsten, T. S., Pleskac, T. J., and Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychological review*, 112(4):862.

Weber, E. U., Blais, A.-R., and Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of behavioral decision making*, 15(4):263–290.

Wetzels, R., Grasman, R. P. P. P., and Wagenmakers, E. (2010). An encompassing prior generalization of the Savage-Dickey density ratio test. *Computational Statistics and Data Analysis*, 54:2094–2102.

Zhang, S. and Lee, M. (2010a). Cognitive models and the wisdom of crowds: A case study using the bandit problem. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 32.

Zhang, S. and Lee, M. D. (2010b). Optimal experimental design for a class of bandit problems. *Journal of Mathematical Psychology*, 54(6):499–508.

Appendix A

Risk Questionnaires

The risk questionnaires are provided here. The Risk Propensity Scale is taken from Meertens and Lion (2008), the Risk Taking Index is taken from Nicholson et al. (2005), and the Domain-Specific Risk Taking scale is taken from Blais and Weber (2006).

A.1 Risk Propensity Scale

Please indicate the extent to which you agree or disagree with the following statement by putting a circle around the option you prefer. Please do not think too long before answering; usually your first inclination is also the best one.

1. Safety first.

   totally disagree 1 2 3 4 5 6 7 8 9 totally agree

2. I do not take risks with my health.

   totally disagree 1 2 3 4 5 6 7 8 9 totally agree

3. I prefer to avoid risks.

   totally disagree 1 2 3 4 5 6 7 8 9 totally agree

4. I take risks regularly

   totally disagree 1 2 3 4 5 6 7 8 9 totally agree

5. I really dislike not knowing what is going to happen

   totally disagree 1 2 3 4 5 6 7 8 9 totally agree

6. I usually view risks as a challenge.

   totally disagree 1 2 3 4 5 6 7 8 9 totally agree

7. I view myself as a ...

   risk avoider 1 2 3 4 5 6 7 8 9 risk seeker

A.2 Risk Taking Index

We are interested in everyday risk-taking. Please could you tell us if any of the following have ever applied to you, *now* or in your adult *past*?

Please use the scales as follows:

1=never, 2=rarely, 3=quite often, 4=often, 5=very often

	Now	*In the Past*
a) recreational risks (*e.g. rock-climbing, scuba diving*)	1 2 3 4 5	1 2 3 4 5
b) health risks (*e.g. smoking, poor diet, high alcohol consumption*)	1 2 3 4 5	1 2 3 4 5
c) career risks (*e.g. quitting a job without another to go to*)	1 2 3 4 5	1 2 3 4 5
d) financial risks (*e.g. gambling, risky investments*)	1 2 3 4 5	1 2 3 4 5
e) safety risks (*e.g. fast driving, city cycling without a helmet*)	1 2 3 4 5	1 2 3 4 5
f) social risks (*e.g. standing for election, publicly challenging a rule or decision*)	1 2 3 4 5	1 2 3 4 5

A.3 Domain-Specific Risk-Taking Scale

Instructions for DOSPERT Risk Taking:

For each of the following statements, please indicate the likelihood that you would engage in the described activity or behavior if you were to find yourself in that situation. Provide a rating from *Extremely Unlikely* to *Extremely Likely*, using the following scale:

1	2	3	4	5	6	7
Extremely Unlikely	Moderately Unlikely	Somewhat Unlikely	Not Sure	Somewhat Likely	Moderately Likely	Extremely Likely

Instructions for DOSPERT Risk Perception:

People often see some risk in situations that contain uncertainty about what the outcome or consequences will be and for which there is the possibility of negative consequences. However, riskiness is a very personal and intuitive notion, and we are interested in your gut level assessment of how risky each situation or behavior is. For each of the following statements, please indicate how risky you perceive each situation. Provide a rating from *Not at all Risky* to *Extremely Risky*, using the following scale:

1	2	3	4	5	6	7
Not at all Risky	Slightly Risky	Somewhat Risky	Moderately Risky	Risky	Very Risky	Extremely Risky

30 items to be rated for both Risk Taking and Risk Perception assessments:

1. Admitting that your tastes are different from those of a friend. (S)
2. Going camping in the wilderness. (R)
3. Betting a day's income at the horse races. (F)
4. Investing 10% of your annual income in a moderate growth mutual fund. (F)
5. Drinking heavily at a social function. (H/S)
6. Taking some questionable deductions on your income tax return. (E)

7. Disagreeing with an authority figure on a major issue. (S)
8. Betting a day's income at a high-stake poker game. (F)
9. Having an affair with a married man/woman. (E)
10. Passing off somebody else's work as your own. (E)
11. Going down a ski run that is beyond your ability. (R)
12. Investing 5% of your annual income in a very speculative stock. (F)
13. Going whitewater rafting at high water in the spring. (R)
14. Betting a day's income on the outcome of a sporting event (F)
15. Engaging in unprotected sex. (H/S)
16. Revealing a friend's secret to someone else. (E)
17. Driving a car without wearing a seat belt. (H/S)
18. Investing 10% of your annual income in a new business venture. (F)
19. Taking a skydiving class. (R)
20. Riding a motorcycle without a helmet. (H/S)
21. Choosing a career that you truly enjoy over a more secure one. (S)
22. Speaking your mind about an unpopular issue in a meeting at work. (S)
23. Sunbathing without sunscreen. (H/S)
24. Bungee jumping off a tall bridge. (R)
25. Piloting a small plane. (R)
26. Walking home alone at night in an unsafe area of town. (H/S)
27. Moving to a city far away from your extended family. (S)
28. Starting a new career in your mid-thirties. (S)
29. Leaving your young children alone at home while running an errand. (E)
30. Not returning a wallet you found that contains $200. (E)

Note: E = Ethical, F = Financial, H/S = Health/Safety, R = Recreational, and S = Social