A cognitive modeling analysis of risk in sequential choice tasks

Maime Guan*   Ryan Stokes *   Joachim Vandekerckhove*   Michael D. Lee*†

Abstract

There are many ways to measure how people manage risk when they make decisions. A standard approach is to measure risk propensity using self-report questionnaires. An alternative approach is to use decision-making tasks that involve risk and uncertainty, and apply cognitive models of task behavior to infer parameters that measure people's risk propensity. We report the results of a within-participants experiment that used three questionnaires and four decision-making tasks. The questionnaires are the Risk Propensity Scale, the Risk Taking Index, and the Domain Specific Risk Taking Scale. The decision-making tasks are the Balloon Analogue Risk Task, the preferential choice gambling task, the optimal stopping problem, and the bandit problem. We analyze the relationships between the risk measures and cognitive parameters using Bayesian inferences about the patterns of correlation, and using a novel cognitive latent variable modeling approach. The results show that people's risk propensity is generally consistent within different conditions for each of the decision-making tasks. There is, however, little evidence that the way people manage risk generalizes across the tasks, or that it corresponds to the questionnaire measures.

Keywords: risky decision making, sequential choice tasks, optimal stopping problems, bandit problems, Balloon Analogue Risk Task, cognitive latent variable modeling

1 Introduction

From financial investments to choosing dating partners, people regularly encounter risky decision-making situations. We are constantly evaluating the potential gains and losses, and the probabilities of each occurring. An individual's intrinsic tendency to be risk seeking, known as their *risk propensity*, has been argued to be a meaningful latent construct that can be interpreted as a dominant influence on people's behavior in risky situations (Dunlop & Romer, 2010; Frey et al., 2017; Josef et al., 2016; Lejuez et al., 2004; Mishra et al., 2010; Pedroni et al., 2017; Sitkin & Weingart, 1995; Stewart Jr & Roth, 2001). Frey et al. (2017) suggest an analogy with the general intelligence construct 'g' from psychometrics (Deary, 2020), raising the possibility of a similar latent construct that guides the balance between risk-seeking and risk-avoiding behavior in uncertain situations.

There are many ways to assess risk propensity. One approach relies on self-report questionnaires, usually in the form of responses to questions using Likert-type scales. Another involves measuring the frequency and type of real-world behaviors related to risk that people engage in. A third approach uses decision-making tasks that involve uncertainty, so that different patterns of decisions can be associated with different risk propensities. If risk is a stable trait, there should be clear relationships between these three types of measures. Accordingly, there is a body of research that examines the relationship between risk questionnaires

and decision-making tasks that aim to measure risk (e.g., Frey et al., 2017; Josef et al., 2016; Szrek et al., 2012). The most commonly used tasks are ones that require choices between gambles (De Martino et al., 2006; Russo & Dosher, 1983; Rieskamp et al., 2006), but other cognitive tasks are also considered. For example, Frey et al. (2017) use the Balloon Analogue Risk Task (BART: Lejuez et al., 2002, 2003a), the Columbia Card Sorting task (Figner et al., 2009), as well as various decision-from-description and decision-from-experience tasks, lotteries, and other tasks. Berg et al. (2005) use a variety of different forms of auctions. Typically, measuring risk propensity using decision-making tasks relies on simple experimental measures. For example, Frey et al. (2017, Table 1) rely entirely on direct behavioral measures of risk, such as counting the number of pump decisions in the BART.

The findings from this literature have been mixed. There is some evidence of risk propensity having a trait-like breadth of influence and stability over time when measured by questionnaires about attitudes and patterns of real-world behavior (e.g., Josef et al., 2016; Mata et al., 2018). The link to behavioral measures in decision-making tasks, however, is far less clear (e.g., Berg et al., 2005; Frey et al., 2017). The motivation for the current research is the possibility that the relationship between cognitive task behavior and risk propensity can be better assessed using cognitive models than simple experimental measures. Our approach is to apply cognitive models of the decision-making process to infer latent psychological parameters that represent risk propensities. In the model-based approach, risk propensity is inferred

*Department of Cognitive Sciences, University of California, Irvine
†Email: mdlee@uci.edu.

from its influence on observed task behavior. Potentially, the model-based approach offers an opportunity to measure an individual's risk propensity in a way that is less open to manipulation, and is more precisely assessed than through simple experimental measures.

The questionnaires we consider are the Risk Propensity Scale (RPS: Meertens & Lion, 2008), the Risk Taking Index (RTI: Nicholson et al., 2005), and the Domain Specific Risk Taking scale (DOSPERT: Blais & Weber, 2006). These three questionnaires have been used in a variety of contexts and have been found to be reliable in measuring people's risk propensity (Harrison et al., 2005).

The decision-making tasks we consider are the BART, the preferential choice gambling task, the optimal stopping problem (Goldstein et al., 2020; Guan et al., 2015; Guan & Lee, 2018; Lee, 2006; Seale & Rapoport, 2000), and the bandit problem (Lee et al., 2011; Steyvers et al., 2009; Zhang & Lee, 2010b). All four of these decision-making tasks involve risk and uncertainty, and have corresponding cognitive models with parameters that can be interpreted as measuring some form of risk propensity. As mentioned earlier, the BART and gambling tasks have previously been considered as natural measures of risk propensity. Our inclusion of the optimal stopping and bandit tasks is relatively novel and exploratory, although optimal stopping tasks are sometimes considered in the related literature on measuring cognitive styles like impulsivity (e.g. Baron et al., 1986).

The structure of this article is as follows: In the next section, we provide an overview of the within-participants experiment involving all of the questionnaires and decision-making tasks. We then present analyses of each of the decision-making tasks separately, describing the experimental procedure and conditions, providing basic empirical results, and describing and applying a cognitive model that makes inferences about risk propensity. Once all four decision-making tasks have been examined, we present results for the questionnaires. Finally, we bring the results together, by presenting first a correlation analysis and then a cognitive latent variable analysis that compare all of the measures of risk propensity. We conclude by discussing the implications of our findings for understanding whether and how risk propensity varies across individuals and generalizes across different tasks and contexts.

2 Overview of Experiment

2.1 Participants

A total of 56 participants were recruited through Amazon Mechanical Turk. Each participant was paid USD$8.00 for completing the experiment. There were 37 male participants and 19 female participants, with ages ranging from 20 to 61 ($M = 36.4$ years, $SD = 11.6$ years).

2.2 Procedure

Each of the four cognitive tasks took about 20–30 minutes to complete. The RPS and RTI took about 5 minutes each, while the DOSPERT took about 10–15 minutes. Each participant completed all of the questionnaires and decision-making tasks. Because the entire experiment took about two hours to complete, the experiment was split into two parts of about one hour each. Each part included two decision-making tasks and either the RPS and RTI or the DOSPERT. The RPS and RTI were completed in the same part because these two questionnaires are much shorter than the DOSPERT. The order of questionnaires and decision-making tasks was randomized across participants.

Upon completing Part 1 of the experiment, each participant was given a unique code. This code allowed them to complete Part 2 and receive compensation. All participants who completed Part 1 returned and completed Part 2. Participants were also encouraged to take a break between Part 1 and Part 2, subject to the requirement that they complete both parts within six days.

3 Balloon Analogue Risk Task

The Balloon Analogue Risk Task is a well-established and widely-used decision-making task for measuring risk propensity (Lejuez et al., 2003a; Lighthall et al., 2009; Rao et al., 2008; Aklin et al., 2005). In the BART, the level of inflation of a balloon corresponds to monetary value. People are repeatedly given the choice either to bank the current value of the balloon, or to take a risk and pump the balloon to add some small amount of air and corresponding monetary value to the balloon. There is some probability the balloon will burst each time it is pumped, in which case the value of the balloon is lost. Usually, the probability of the balloon bursting increases with each successive pump, but a simpler version in which this probability is fixed has been used by some authors (e.g., Cavanagh et al., 2012; Van Ravenzwaaij et al., 2011). A BART problem involves a sequence of bank or pump choices, and finishes when either the value is banked or the balloon bursts.

Individual risk propensity is most often quantified by the mean number of pumps made across problems, excluding those problems where the balloon burst (Schmitz et al., 2016). An individual who is risk seeking is likely to pump the balloon more times across problems than an individual who is risk averse. The mean number of pumps has been shown to correlate with risk taking behaviors such as smoking, alcohol and drug abuse, and financial decision making (Hopko et al., 2006; Holmes et al., 2009; Lejuez et al., 2002, 2003a; Schonberg et al., 2011), as well as psychological traits such as impulsivity, anxiety, and psychopathy (Hunt et al., 2005; Lauriola et al., 2014).

3.1 Method

Participants completed two BART conditions, differing in the fixed probability of the balloon bursting at each trial. These probabilities were $p = 0.1$ and $p = 0.2$. 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. At the beginning of the experiment, they received a virtual bank with $0 and a balloon that was worth $1. At the bottom of the screen there was a "pump" button and a "bank" button. With each pump, the balloon's worth increased by $1. Participants were instructed to maximize their monetary reward. All participants completed the same 50 problems within each of the two conditions. The order of problems within each condition was randomized across participants.

3.2 Two-Parameter BART Model

Wallsten et al. (2005) pioneered the development of cognitive models for the BART that are capable of inferring latent parameters measuring risk propensity. Their modeling approach was further developed by Pleskac (2008) and Zhou et al. (2019). We use the two-parameter BART model developed by Pleskac (2008, see also Van Ravenzwaaij et al. 2011) as a simplification of one of the original Wallsten et al. (2005) models. The two-parameter model assumes that a decision maker believes that there is a single constant probability that a pump will make a balloon burst p^{belief} that is fixed over all problems. It also assumes that they decide on a number of pumps prior to the first pump in a problem, and do 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. It is defined as

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

where $\gamma^+ \sim \text{uniform}(0, 10)$.

Our implementation of the two-parameter BART model naturally incorporates censoring by modeling the probability of each of participant pumping or banking on each trial they completed. Thus, the behavioral data are represented as $y_{ijk} = 1$ if the ith participant pumped on the kth trial of the jth problem, and $y_{ijk} = 0$ if they banked.

In the two-parameter BART model, the probability that the ith participant will pump on the kth trial of the jth problem, p_{ijk}^{pump} depends on both ω_i and a behavioral consistency parameter β_i, in terms of a logistic function

$$p_{ijk}^{\text{pump}} = \frac{1}{1 + \exp\left(\beta_i\left(k - \omega_i\right)\right)},$$

with $\beta_i \sim \text{uniform}(0, 10)$. Given this pumping probability, the observed data are simply modeled as $y_{ijk} \sim$

Bernoulli$\left(p_{ijk}^{\text{pump}}\right)$ over all observed trials, finishing on the trial for each problem at which the participant banked or the balloon burst.

The logistic relationship that defines the pumping probabilities means that relatively higher values for β_i correspond to more consistency in decision making. If $\beta_i = 0$ then $p_{ijk}^{\text{pump}} = 0.5$, and the participant's decision to pump or bank is random. As $\beta_i \to \infty$, the participant's behavior becomes completely determined by whether or not the number of pumps k is greater than ω_i.

The γ_i^+ parameter provides a measure of risk propensity, since it controls the number of pumps attempted. Larger values of γ_i^+ correspond to more pumps and greater risk seeking. Smaller values of γ_i^+ correspond to fewer pumps, and more risk-averse behavior.

We implemented the two-parameter model model, and all of the other cognitive models considered in this article, as graphical models using JAGS (Plummer, 2003). JAGS is software that facilitates MCMC-based computational Bayesian inference (Lee & Wagenmakers, 2013). All of our modeling results are based on four chains of 1,000 samples each, collected after 2,000 discarded burn-in samples. The chains were verified for convergence using visual inspection and the standard $\hat{R}$ statistic (Brooks & Gelman, 1997).

3.3 Modeling Results

For all of the cognitive modeling in this article, we apply the model to the data in three steps. First, we define task-specific contaminant models, identifying those participants who did not understand the task, or did not complete it in a motivated way. These contaminant participants are removed from the subsequent analysis. Secondly, we examine the descriptive adequacy of the model for the remaining participants, using the standard Bayesian approach of posterior-predictive checking (Gelman et al., 2004). Finally, we report the inferences for the model parameters, usually starting with a few illustrative participants who demonstrate the range of interpretable individual differences, before showing the inferences for all non-contaminant participants.

3.3.1 Removing Contaminants

We developed two contaminant models for BART behavior. The first was based on a cutoff for the β consistency parameter. If a participant's behavior was extremely inconsistent across the problems they completed, they were considered contaminants. We used a cutoff of 0.2, which removed 11 participants. The second contaminant model was developed to capture behavior motivated by wanting to finish the experiment as quickly as possible. If a participant banked on all of the problems they were also considered contaminants. Three participants were associated with this form of contaminant
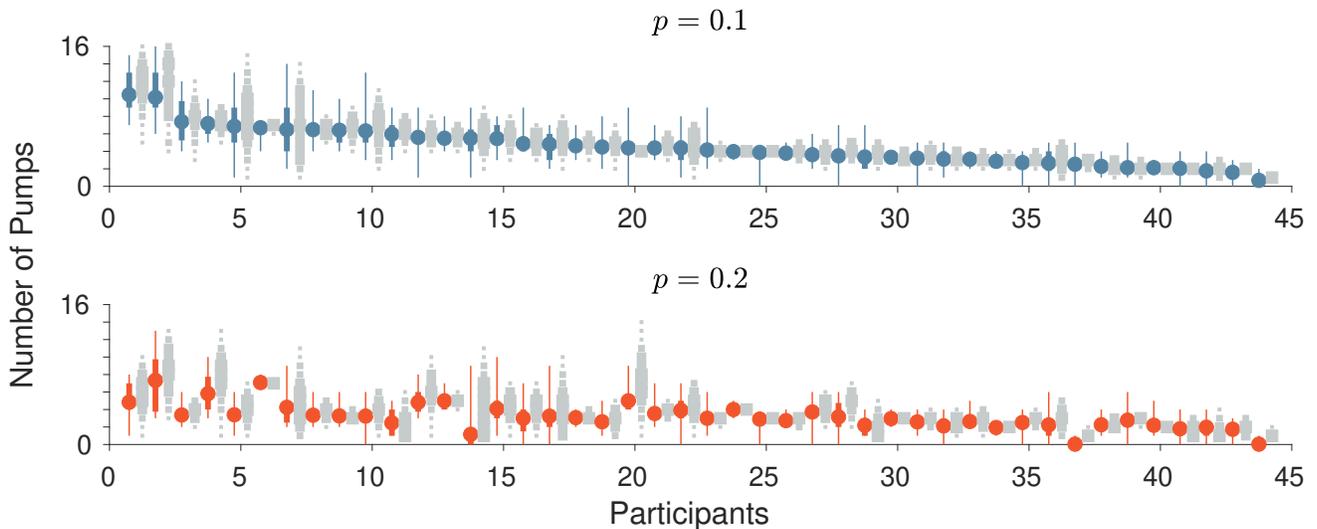
FIGURE 1: Posterior predictive distributions of the number of pumps for each participant in each condition, sorted by the mean number of pumps per participant in the $p = 0.1$ condition. The top panel corresponds to the condition with bursting probability $p = 0.1$, and the bottom panel corresponds to the condition with bursting probability $p = 0.2$. The posterior predictive distributions are shown as gray squares. The minimum and maximum, as well as the 0.25 and 0.75 quantiles, and the median of the behavioral data are shown to the immediate left in blue for the $p = 0.1$ condition and red for the $p = 0.2$ condition.

behavior. Thus, overall, a total of 14 contaminant participants were removed, and a total of 42 participants were used in the modeling analysis.

3.3.2 Descriptive Adequacy

Figure 1 summarizes a posterior predictive check of the descriptive adequacy of the two-parameter BART model. The distributions of the number of pumps are shown as gray squares, with areas proportional to the posterior predictive mass. The observed data are shown to the left, with dots representing the median number of pumps, thicker solid lines representing the 0.25 and 0.75 quantiles, and thinner lines spanning the minimum and maximum observed number of pumps. The posterior predictive distributions generally match the observed data, suggesting that the model provides a reasonable account of people's behavior.

3.3.3 Inferred Risk Propensity and Consistency

Figure 2 shows the inferred γ^+ and β parameter values for four representative participants, together with a summary of their observed behavior. Each panel shows the distribution of the number of pumps that the participant made, excluding problems on which the balloon burst. The left column shows the condition with $p = 0.1$ and the right column shows the condition with $p = 0.2$.

Participant 1 can be seen to be consistently risk seeking. They choose to pump a relatively large number of times in both conditions. This pattern of behavior is captured by their

risk and consistency parameters, with relatively high values of γ^+ and β parameter values. Participant 2 is also risk seeking, in the sense that they generally pump a relatively large number of times across both conditions, but they do so less consistently. The number of times they pump in both conditions varies widely from 3 to more than 15 pumps. This behavior is quantified by their inferred parameter values, with relatively high values of γ^+ but relatively low values of β. Participant 3 is 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 parameter values. Participant 4 is also risk averse, but is more inconsistent than Participant 3. This is captured with relatively low values of both γ^+ and β parameter values .

Figure 3 shows the joint and marginal distributions of the posterior expectations of γ^+ and β for all participants and for both conditions. The four representative participants shown from Figure 2 are labeled. It is evident that there is a wide range of individual differences in both risk propensity and consistency parameters. There appears to be a negative and 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 from the origin closer to the lower-right corner, they become more risk seeking but continue to lack consistency. As participants move further away from
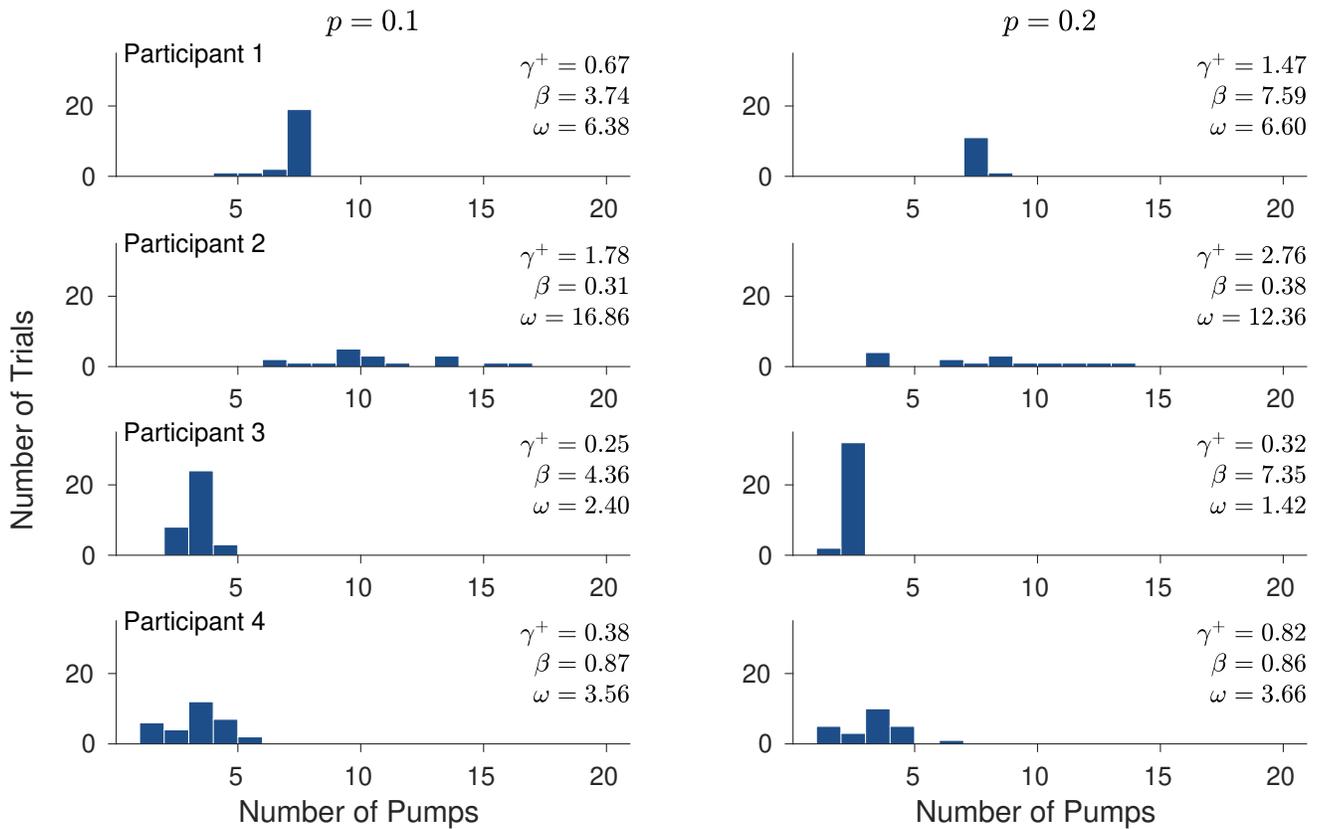
FIGURE 2: Observed behavior and inferred parameter values for four representative participants. The left column corresponds to the condition with bursting probability $p = 0.1$ and the right column corresponds to the condition with bursting probability $p = 0.2$. The distributions show the number of pumps each participant made before banking, excluding problems where the balloon burst. The inferred values of the γ^+, β, and ω parameters are also shown.
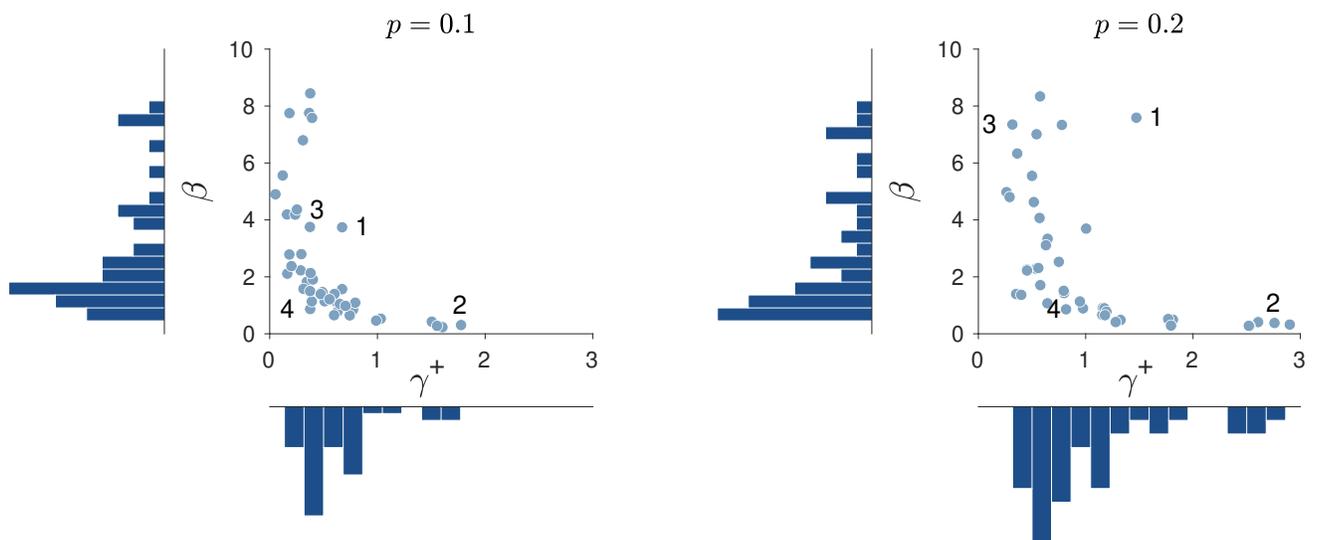
FIGURE 3: Joint and marginal distributions of β and γ^+ posterior expectations across the two conditions for each participant. The four representative participants shown from Figure 2 are labeled.

the origin and closer to the top-left corner, they become consistently risk averse.

4 Gambling Task

Perhaps the most common task for studying decision-making under risk and uncertainty involves people choosing between pairs of gambles (De Martino et al., 2006; Russo & Dosher, 1983; Rieskamp, 2008). Each gamble is defined in terms of the probabilities of different monetary outcomes, and people are asked to choose the gamble they prefer. For example, a person might be asked to choose between Gamble A, which leads to winning$50 with probability 0.6 and losing $50 with probability 0.4, and Gamble B, which leads to winning $100 with probability 0.65 and losing $100 with probability 0.35.

4.1 Method

Participants completed two gambling tasks conditions. One condition was framed in terms of gains and the other was framed in terms of losses. In the gain condition, participants were instructed to maximize their monetary reward over the entire set of problems. In the loss condition participants were instructed to minimize their monetary losses. All of the participants completed the same 40 problems in each condition, but the order of problems within each condition was randomized across participants.

The pairs of gambles were presented as pie charts labeled with their respective payoffs and probabilities. A screenshot of the experimental interface is provided in the supplementary materials. Participants chose between gambles by clicking the corresponding pie chart. The expected values of the outcomes were not provided to the participants and no feedback was given.

4.2 Cumulative Prospect Theory Model

Important cognitive models of how people choose between gambles include regret theory (Loomes & Sugden, 1982), decision-field theory (Busemeyer & Townsend, 1993), the priority heuristic (Brandstätter et al., 2006), anticipated utility theory (Quiggin, 1982), and prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1981). All of these models extend the standard economic account of choice as maximizing expected utility (von Neumann & Morgenstern, 1947) and attempt to provide an account in terms of cognitive processes and parameters.

We use cumulative prospect theory (CPT), which makes a set of assumptions about how people subjectively weigh the value of outcomes and probabilities. CPT 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. In particular, it assumes that the same

absolute value of a loss has a larger impact on the decision than a gain, consistent with the phenomenon of loss aversion (Kahneman & Tversky, 1979). In addition, prospect theory assumes that people subjectively represent probabilities, typically overestimating small probabilities and underestimating large probabilities.

We use a variant of the CPT model developed and implemented by Nilsson et al. (2011). In this model, the expected utility of an alternative O is defined as

$$EU(O) = \sum_i p_i u(x_i),$$

where $u(\cdot)$ defines the subjective utility of x_i.

This subjective utility is weighted by the probability p_i that the ith outcome occurs. According to the CPT model, if alternative O has two possible outcomes, then the subjective value V of O is defined as

$$V(O) = \sum_i \pi(p_i) v(x_i),$$

where $\pi(\cdot)$ is a weighting function of the objective probabilities and $v(\cdot)$ is a function defining the subjective value of the ith outcome. The probability weighting function and the value function differ for gains and losses. The subjective value of payoff x is defined as

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

where $0 < \alpha < 1$ is a parameter that controls the curvature of the value function. Nilsson et al. (2011) used different value functions for gains and losses. We use a simplification of the model in which the shape of the value function, determined by α, is the same for gains and losses. If $\lambda > 1$, losses carry more weight than gains, corresponding to the theoretical assumption of loss aversion. The larger the value of λ, the greater the relative emphasis given to losses. When $0 < \lambda < 1$, in contrast, gains have more impact on the decision than losses. Although prospect theory expects loss aversion, we use a prior $\lambda \sim \text{uniform}(0, 10)$ that tests this assumption.

The CPT model generates subjective probabilities by a weighting function which, for two possible outcomes, is defined as

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

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

Finally, our CPT model allows for probabilistic decision making by assuming a choice rule in which choice probabilities are a monotonic function of the differences of the subjective values of the gambles. Specifically, the exponential
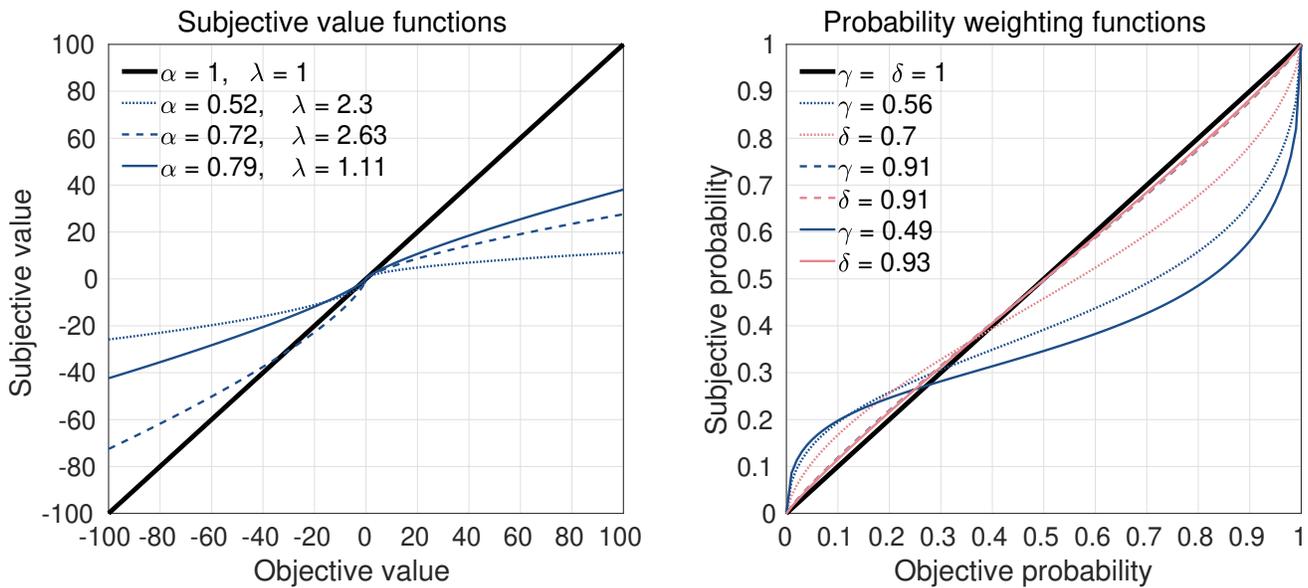
FIGURE 4: Inferred subjective value function curves and probability weighting function curves for representative participants. The three participants in the left panel span the range of inferred individual differences in α and λ. The three participants in the right panel span the range of inferred individual differences in γ and δ.

Luce choice rule, rewritten as a logistic choice rule, assumes that the probability of choosing Gamble A over Gamble B is

$$\theta_{A,B} = \frac{1}{1 + \exp \phi \left(V\left(B \right) - V\left(A \right) \right)}.$$

The parameter ϕ can be interpreted as a measure of the consistency of choice behavior. When $\phi = 0$, the probability of choosing Gamble A over Gamble B becomes $\frac{1}{2}$, and choice behavior is random. As ϕ increases, choice behavior becomes increasingly determined by the difference in subjective value between Gamble A and Gamble B. As $\phi \to \infty$, choices become increasingly consistent in the underlying preference, until in the limit the preferred gamble is always chosen.

We use independent priors for all five parameters for each participant. Besides the prior $\lambda \sim \text{uniform}(0, 10)$ already mentioned, the remaining parameters have priors $\alpha \sim \text{uniform}(0, 1)$, $\gamma \sim \text{uniform}(0, 1)$, $\delta \sim \text{uniform}(0, 1)$, and $\phi \sim \text{gamma}(2, 1)$. Note that the final prior on the response consistency gives the highest density to $\phi = 1$, which corresponds to probability matching, while also allowing for more random or more deterministic behavior.

To measure risk propensity using the CPT model we focus on the loss aversion parameter λ. The motivation is that an individual who exhibits strong loss aversion can be interpreted as being risk averse, since their preference will be for gambles that avoid the possibility of a large loss. For the inference about loss aversion to be meaningful, there must be some level of behavioral consistency, and so we place a secondary focus on the ϕ parameter. We acknowledge that there are other ways in which the CPT model could

be interpreted in terms of risk propensity. For example, if the probability weighting function infers that an individual perceives probabilities in extreme ways, significantly underestimating small probabilities and overestimating large probabilities, this could be seen as supporting a risky perception of the gambles. Alternatively, a lack of consistency in decision-making corresponds to a form of risk-seeking, but is more in line with erratic behavior than the underlying risk propensity trait we aim to measure.

4.3 Modeling Results

4.3.1 Removing Contaminants

We used a simple guessing model of contamination that assumes the probability any participant will choose Gamble A over Gamble B is $\theta_{A,B} = \frac{1}{2}$. This guessing model was applied using a latent-mixture procedure based on model-indicator variables (Zeigenfuse & Lee, 2010). A total of 22 of the participants were inferred to be using the guessing model, and were removed from the remainder of the analysis.

4.3.2 Descriptive Adequacy

We checked the descriptive adequacy of the CPT model using the mode of the posterior predictive distribution for each participant on each problem. This measure of the choice described by the model agreed with 77% of the decisions that participants made. Given that the chance level of agreement for choosing between two gambles is 50%, we interpret these results as suggesting that the CPT model provides a reasonable account of people's behavior in the gambling task.
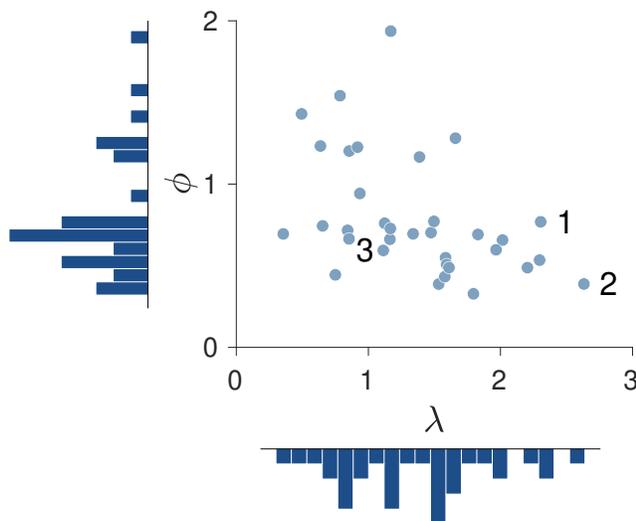
FIGURE 5: The joint and marginal distributions of the posterior expectations of the λ risk aversion and ϕ consistency parameters over all participants. The representative participants from the left panel of Figure 4 are labeled.

4.3.3   Inferred Subjective Value Functions

We found large individual differences in the subjective value functions and probability weighting functions that participants use. Figure 4 shows the inferred functions for a set of representative participants. In the left panel 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 the sense that the magnitude of losses are weighed more heavily 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. The third participant, shown by the solid line, has a relatively high value of α but λ is close to one. Consequently, the effect of undervaluing the magnitude of both gains and losses is weaker.

The first participant in the right panel of Figure 4, shown by the dotted line, 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 line, has relatively high values of both γ and δ. Their probability weighting functions are extremely close to the diagonal, which corresponds to good calibration. The third participant, shown by the solid line, has a relatively low value of γ but high value of δ. This participant significantly underestimates small probabilities and overestimates large probabilities.

Figure 5 shows the joint and marginal distributions of the posterior expectations of the loss aversion parameter λ as a measure of risk propensity, and the consistency parameter ϕ, over all participants. The representative participants from the left panel of Figure 4 are labeled. It is clear from that there is a range of inferred individual differences in both loss aversion and consistency. About one-third of the participants exhibit the opposite of loss aversion, with λ values below 1. About one-third of the participants exhibit relatively strong loss aversion with values of λ values over 1.5. All of the ϕ consistency parameters are inferred to be well above 0, as expected given the removal of guessing contaminants, but many are less consistent than probability matching.

5   Optimal Stopping Problems

5.1   Theoretical Background

Optimal stopping problems are sequential decision-making tasks in which people must choose the best option from a sequence, under the constraint that an option can only be chosen when it is presented (Ferguson, 1989; Gilbert & Mosteller, 1966). These problems are sometimes called secretary problems, based on the analogy of interviewing a sequence of candidates for a job with the requirement that offers must be made immediately after an interview has finished, and before the next candidate is evaluated.

People's behavior on optimal stopping problems has been widely studied in a variety of contexts, using a number of different versions of the task (Bearden et al., 2006; Christian & Griffiths, 2016; Kogut, 1990; Lee, 2006; Seale & Rapoport, 1997, 2000). Some studies have used 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 & Rapoport, 1997, 2000; Bearden et al., 2006). Other studies have used the full-information version of the task, in which the values of the alternatives are presented (Goldstein et al., 2020; Lee, 2006; Guan et al., 2014, 2015; Shu, 2008). For both of these versions there are known optimal solution processes to which people's performance can be compared (Ferguson, 1989; Gilbert & Mosteller, 1966).

We use the full-information version of the problem, for which the optimal solution is to choose the first number that is both currently maximal and above a threshold that depends upon the position in the sequence. The values of the optimal thresholds also depend on two properties of the problem. One is the number of options in the sequence, known as the length of the problem. Intuitively, the more options a problem has, the higher thresholds should be, especially early in the sequence. The second property is the distribution from which values of the options are chosen, known as the environment distribution. Intuitively, distributions that generate many large values require setting higher thresholds, while

distributions that generate many small values require setting lower thresholds.

5.2 Method

Participants completed four types of optimal stopping problems, made up of combining problem lengths of four and eight with environment distributions we call neutral and plentiful. In the neutral environment, values were generated from the uniform$(0, 100)$ distribution. In the plentiful environment, values were generated by scaling values drawn from the beta$(4, 2)$ distribution to the range from 0 to 100. All participants completed the same 40 problems within each condition, and the order of problems within each condition was randomized across participants.

To complete each problem, participants were instructed to pick the heaviest cartoon cat out of a sequence, with each cat's weight ranging from 0 to 100 pounds. A screenshot of the interface is provided in the supplementary material. Participants were told the length of the sequence, that a value could only be chosen when it is presented, that any value that was not the maximum was incorrect, and that the last value must be chosen if no values were chosen beforehand. Participants indicated whether or not they chose each presented value by pressing either a "select" or a "pass" button. The values that participants rejected in a sequence were not shown once the next value in the sequence was displayed. The values in the sequence after the one the participant chose were never presented. After each problem, participants were provided with feedback indicating whether or not they chose the option with the maximum value.

5.3 Bias-From-Optimal Model

Previous work modeling decision making in optimal stopping problems has found evidence that people use a series of thresholds to make decisions, and that there are large individual differences in thresholds (Goldstein et al., 2020; Guan et al., 2014; Guan & Lee, 2018; Lee, 2006). A surprising but reliable finding is that, beyond the initial few problems in an environment (Goldstein et al., 2020), there is relatively little learning or adjustment of thresholds (Baumann et al., 2018; Campbell & Lee, 2006; Guan et al., 2014; Lee, 2006). This justifies modeling an individual's decisions in terms of the same set of thresholds being applied to all of the problems.

We use the previously-developed Bias-From-Optimal (BFO) model to characterize the thresholds people use. (Guan et al., 2015). The BFO model represents the thresholds an individual uses in terms of how strongly they deviate from the optimal thresholds for the problem length and environmental distribution. We denote the optimal thresholds as $\tilde{\tau}_1, \ldots, \tilde{\tau}_m$ for a problem of length m (Gilbert & Mosteller, 1966, Table 2). Naturally, the last threshold in
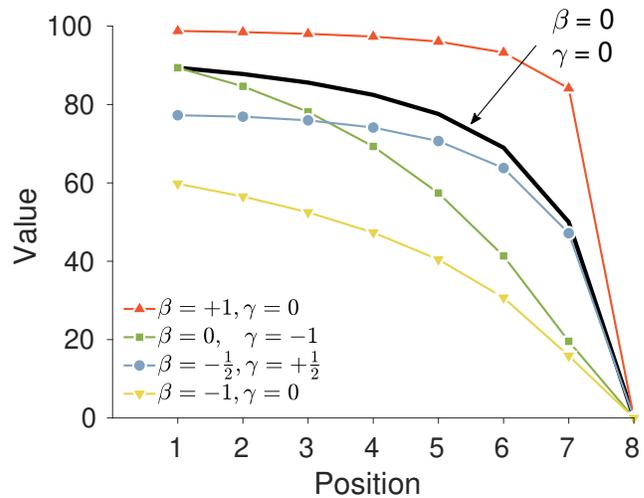
FIGURE 6: The thresholds produced by the bias-from-optimal threshold model under different parameterizations of β and γ. The optimal threshold, corresponding to $\beta = \gamma = 0$, is also shown.

the sequence must be 0 since the last value must be chosen. The ith participant's thresholds depend on a parameter $\beta_i^m \sim \text{Gaussian}(0, 1)$ that determines how far above or below their threshold is from optimal, and a parameter $\gamma_i^m \sim \text{Gaussian}(0, 1)$ that determines how much their bias increases or decreases as the sequence progresses. Formally, under the BFO model, the ith participant's kth threshold in a problem of length m is

$$\tau_{ik}^m = 100 \times \Phi\left(\Phi^{-1}\left(\frac{\tilde{\tau}_k^m}{100}\right) + \beta_i^m + \frac{k}{m}\gamma_i^m\right)$$

for the first $m - 1$ positions, and $\tau_{im}^m = 0$ for the last. The link functions $\Phi(\cdot)$ and $\Phi^{-1}(\cdot)$ are the Gaussian cumulative distribution and inverse cumulative distribution functions, respectively.

According to the BFO model, the probability that the ith participant will choose the value they are presented in the kth position on their jth problem is

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

for the first m positions and

$$\theta_{ijm}^m = 1 - \sum_{k=1}^{m-1} \theta_{ijk}^m$$

for the last position. The parameter $\alpha_i^m \sim \text{uniform}(0, 1)$ is the individual-level accuracy of execution that corresponds to how often the deterministic threshold model is followed (Guan et al., 2014; Rieskamp & Otto, 2006).

Figure 6 shows how the shape of threshold functions changes with different values of β and γ, as compared to

the optimal decision threshold for a problem of length eight in the 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 so that $\gamma, \beta \sim \text{Gaussian}(0, 1)$ and $\alpha \sim \text{uniform}(0, 1)$.

Our decision to use the BFO model was based on the direct interpretability of its parameters in terms of risk propensity. It is an unrealistic model of the cognitive processes involved in optimal stopping problem decisions, because it assumes perfect knowledge of the optimal thresholds, which are difficult to derive and compute. Alternative models based on fixed and linearly decreasing thresholds provide more realistic cognitive processing accounts (Baumann et al., in press; Goldstein et al., 2020; Lee, 2006; Lee & Courey, in press). The BFO model is better interpreted as a measurement model, with the β and γ parameters quantifying how a set of thresholds are more or less risky than optimal.

One interpretation is that higher thresholds that require higher values represent risk seeking and lower thresholds represent risk aversion. Larger values β increase thresholds, and larger values of γ maintain higher thresholds for longer. Under this interpretation larger values of β and γ correspond to greater risk propensity. In contrast, smaller values of β and γ both lead to lower thresholds over the course of the sequence and correspond to lower risk propensity.

5.4 Modeling Results

Before applying the BFO model, we checked that there was no clear evidence of learning or adaptation. As discussed above, this is a basic empirical pre-condition for the application of threshold models. Figure 7 shows the performance of participants, measured by the proportion of problems for which they correctly chose the maximum. The problems were split into four blocks of 10 problems each. In the two length-four conditions mean performance is between about 0.5 and 0.6. In the two length-eight conditions mean performance is between about 0.3 and 0.5. Participant performance is better in the shorter problems, but there do not appear to be large differences in performance between the neutral and plentiful environments. These results do not suggest there is any significant learning or adaptation.
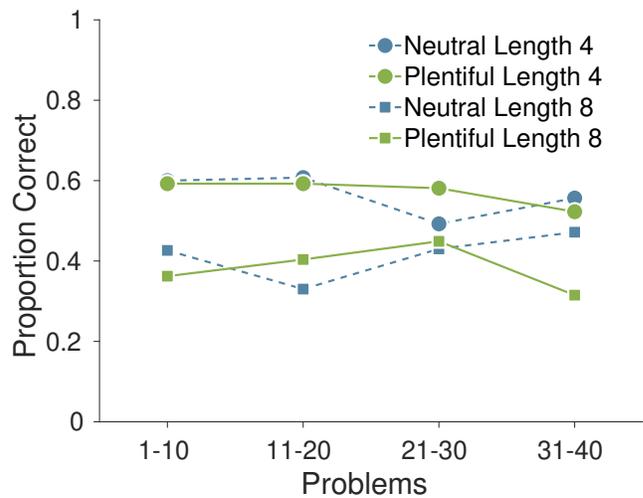
FIGURE 7: Mean proportion correct over all participants on successive blocks of 10 problems for the four different optimal stopping conditions.

5.4.1 Removing Contaminants

We developed two contaminant models for the optimal stopping task. The first assumes that people simply picked the first option in the sequence repeatedly across all problems, regardless of its value. The second assumes that people choose randomly, so that each option in the sequence is equally likely to be chosen. A latent-mixture analysis identified three participants as using the first contaminant model, and these were removed from subsequent analysis.

5.4.2 Descriptive Adequacy

As a posterior predictive check, we took the mode of the posterior predictive distribution for each participant on each problem as the decision the model expects. By this measure, the BFO model successfully described about 77% of the decisions that participants made. Given that the base rate or chance level of agreement is 25% for length-four problems and 12.5% for length-eight problems, we interpret these results as evidence that the model provides a reasonable account of people's behavior.

5.4.3 Inferred Thresholds

Figure 8 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 generally sensitive to both length of the problem and the environmental distribution from which values are drawn. The thresholds in the plentiful environment conditions are relatively higher than the thresholds in the neutral environment conditions. The thresholds in the length-eight
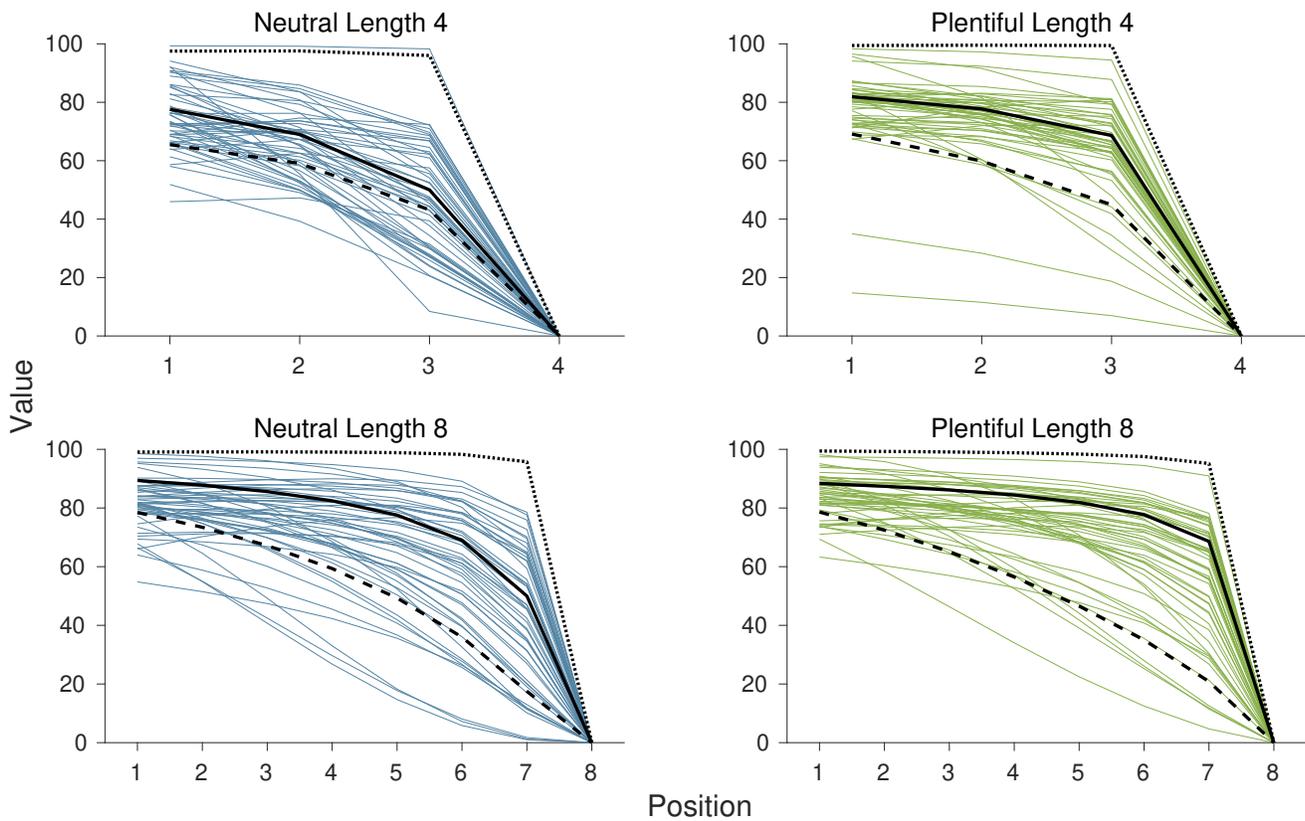
FIGURE 8: The inferred thresholds for all participants in the optimal stopping conditions corresponding to the length-four neutral environment (top-left), the length-four plentiful environment (top-right), the length-eight neutral environment (bottom-left), and the length-eight plentiful environment (bottom-right). Two representative participants are shown by the dashed and dotted lines.

conditions remain higher longer into the sequence than the thresholds in the length-four conditions. Interestingly, it appears that participants in the length-eight conditions tend to use thresholds that are lower than optimal in both environments, and especially so in the plentiful environment.

It is also clear that there are individual differences in thresholds in all four conditions. Two participants are highlighted by dotted and dashed lines in Figure 8, showing their inferred thresholds in all four conditions. These participants were chosen because they show very different patterns of risk propensity in terms of their thresholds. The participant represented by the dotted lines can be seen to be risk seeking, because their thresholds for all four conditions are much higher than optimal. The participant starts their threshold high and maintains it at a high level as the sequence progresses. This risk-seeking behavior is quantified by their β and γ parameter values, which are both positive and relatively large. Conversely, the participant represented by the dashed lines can be seen to be risk averse, because their thresholds are much lower than optimal in all four conditions. The participant starts their threshold low and lowers it quickly as the sequence progresses. This risk-averse behav-

ior is also quantified by their large negative β and γ parameter values.

Figure 9 summarizes the individual differences across all participants for all four conditions. The posterior expectations of the β and γ risk parameters are shown jointly in the scatter-plot in the center panel, and their marginal distributions are shown as histograms on the bottom and left margins. The two participants highlighted in Figure 8 are labeled in the joint distribution. The dotted lines represent where β and γ are equal to 0. Where the dotted lines meet in the center represents the optimal threshold. It is clear that there is a wide range of both quantitative and qualitative individual differences in risk propensity, because all four quadrants around optimality are populated.

6 Bandit Problems

6.1 Theoretical Background

Bandit problems are widely used to study human decision making under risk and uncertainty (Banks et al., 1997; Daw et al., 2006; Meyer & Shi, 1995; Lee et al., 2011). In bandit
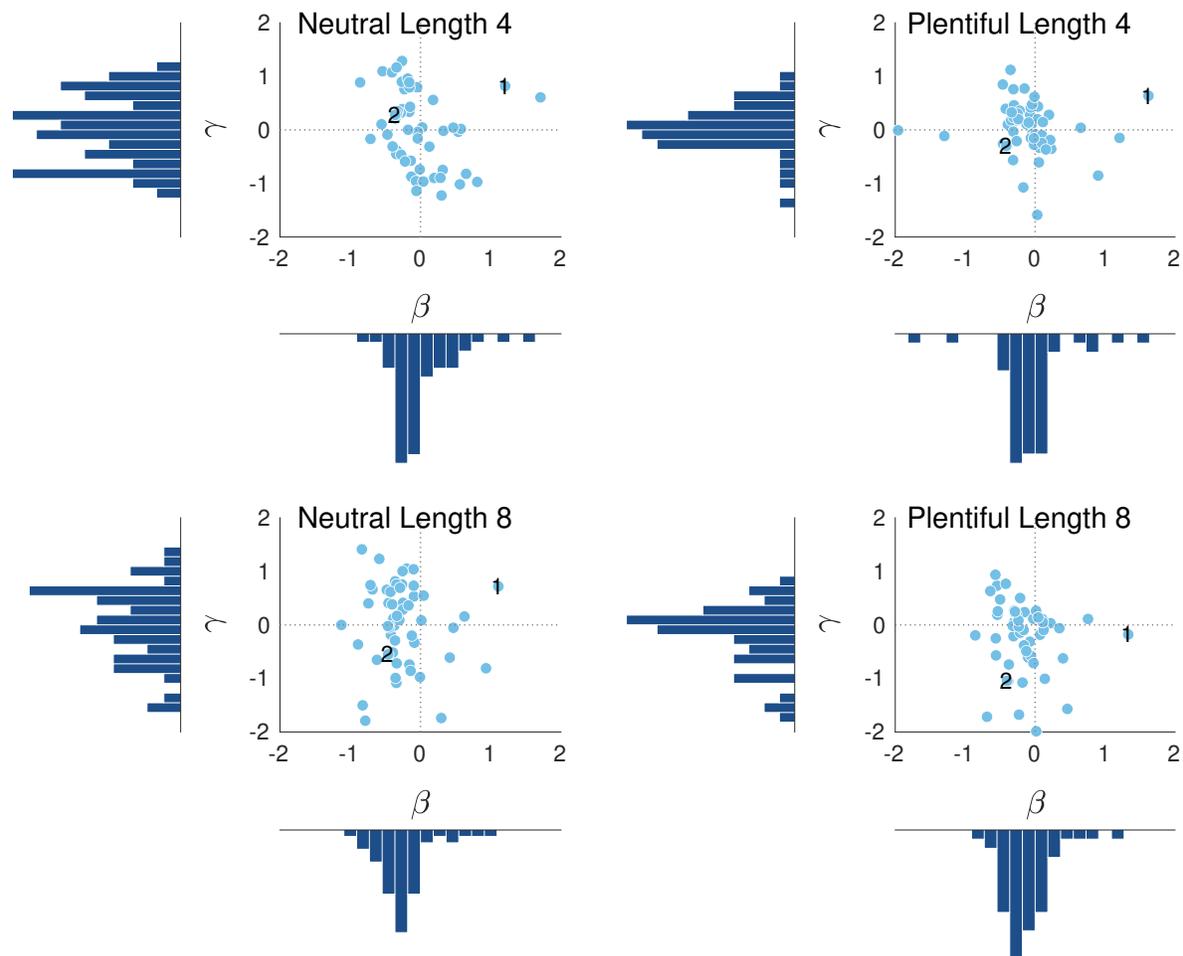
FIGURE 9: The joint and marginal distributions of the posterior expectations of the β and γ parameters, across the four conditions for all of the participants. The risk-seeking and risk-averse participants from Figure 8 are labeled.

problems, people must choose repeatedly between a set of alternatives. Each alternative has a fixed reward rate that is unknown to the decision maker, and each time it is chosen this probability is used to generate either a reward or a failure. The goal is to maximize the total number of rewards over the sequence of decisions. Bandit problems are psychologically interesting because they require that the exploration of new good alternatives be balanced with the exploitation of good existing alternatives (Mehlhorn et al., 2015). People generally start by exploring the different available alternatives before shifting to exploit the alternative with the highest reward rate.

Bandit problems can differ in terms of how many alternatives are available and in terms of how many decisions are made within a problem. In infinite-horizon bandit problems the total number of decisions to be made is not known in advance, but there is some probability that the problem stops after any decision. In finite-horizon bandit problems the total number of decisions to be made within a problem is fixed and known in advance. This corresponds to the length of a

problem. Bandit problems can also differ in terms of the distributions of reward rates that underlie each alternative. This distribution corresponds to the environment for the problem.

6.2 Method

Participants completed four types of finite-horizon bandit problems, all involving two alternatives. The four conditions combined problem lengths of eight and 16 with neutral and plentiful environmental distributions. In the neutral environment, reward probabilities were generated from the $\mathrm{uniform}(0, 1)$ distribution. In the plentiful environment, reward probabilities were generated from the $\mathrm{beta}(4, 2)$ distribution. Consequently, the plentiful environments contained alternatives that had relatively higher reward rates. All participants completed 40 problems within each condition and the order of problems within each condition was randomized across participants.

Participants were instructed to maximize the number of rewards by pulling the arms of two cartoon slot machines. A

screenshot of the interface is provided in the supplementary material. Before beginning each condition, participants were informed that the reward probabilities for each machine were different for each problem in the block, but the same for all choices within a problem. They were also told how many choices were required for each problem. They were not, however, told the underlying distribution of the reward probabilities.

Participants made their choice selection by clicking a "pull" button under one or the other of the two slot machines. The reward or failure outcome was then provided, in the form of a green or red bar. If a choice resulted in a reward, a green bar was added to the left side of the chosen slot machine. If a choice resulted in a failure, a red bar was added to the left side of the chosen slot machine. Thus, the bars showed the cumulative pattern of reward and failure over the course of the problem, and the total reward points earned on the current problem was also shown at the top of the screen. A problem was completed once the participant completed all of the choices.

6.3 Extended Win-Stay Lose-Shift Model

There are many different models of human decision making on bandit problems, including the ϵ-greedy, ϵ-decreasing and the τ-first model (Sutton & Barto, 1998). We use a variant of perhaps the simplest and most widely used model, known as win-stay lose-shift (WSLS: Robbins, 1952; Sutton & Barto, 1998). In its deterministic form, this model assumes that people stay with the most recently-chosen alternative if it provides a reward, but shift to another alternative if it does not. In the standard stochastic version of the WSLS strategy, there is a probability γ of following this rule for every decision.

In our extended WSLS model there is a probability γ_w of staying after a reward and a potentially different probability γ_l of shifting after a failure. This WSLS model allows there to be a psychological difference between reacting to reward and failure in the decision-making process. This model has been found to account well for people's behavior (Lee et al., 2011; Zhang & Lee, 2010a).

The extended WSLS model does not require memory of previous actions and outcomes beyond the immediately preceding trial. It is also insensitive to whether the horizon is an infinite or finite. Despite this simplicity, it provides a measure of risk propensity. A person who is risk seeking is likely to shift to another alternative with a high probability following a failure, in order to explore the other available options. In contrast, a person who is risk averse is likely to shift to another alternative with a relatively lower probability following a failure.

We represent the behavioral data as $y_{ijk} = 1$ if the ith participant chose the left alternative on the kth trial of their jth problem, and $y_{ijk} = 0$ if they chose the right option. The

extended WSLS model assumes the probability of choosing the left alternative is

$$\theta_{ijk} = \begin{cases} 1/2 & \text{if } k = 1 \\ \gamma_w & \text{if chose left and } r_{ij(k-1)} = 1 \\ 1 - \gamma_l & \text{if chose left and } r_{ij(k-1)} = 0 \\ 1 - \gamma_w & \text{if chose right and } r_{ij(k-1)} = 1 \\ \gamma_l & \text{if chose right and } r_{ij(k-1)} = 0, \end{cases}$$

where $r_{ij(k-1)} = 1$ if the previously selected alternative resulted in a reward, and $r_{ij(k-1)} = 0$ if the previously selected alternative resulted in a failure. The observed rewards and failures on each trial r_{ijk} are generated by $r_{ijk} \sim \text{Bernoulli}(p^x)$, where p^{left} and p^{right} are the reward rates for the two alternatives. These reward rates are generated from either the neutral or plentiful environment. The behavioral data are modeled as $y_{ijk} \sim \text{Bernoulli}(\theta_{ijk})$. Finally, our model uses the priors $\gamma_w, \gamma_l \sim \text{uniform}(0, 1)$.

6.4 Modeling Results

6.4.1 Removing Contaminants

We used a guessing contaminant model in which, for every trial of a problem, the participant chooses at random. Using the latent-mixture approach, there was overwhelming evidence in favor of the extended WSLS model over the guessing model for all of the participants. Consequently, no contaminant participants were removed and the modeling analysis used all 56 participants.

6.4.2 Descriptive Adequacy

As a posterior predictive check, the mode of the posterior predictive distribution for each participant on each problem was used as the decision that the model expected to have been made. The extended WSLS model was able to describe 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, we interpret this result as showing that the extended WSLS model provides a good account of people's behavior.

6.4.3 Inferred Win-Stay Lose-Shift Probabilities

Figure 10 shows the numbers of shifts following rewards and failures across positions for four representative participants. These participants were chosen because they span the range of inferred individual differences. The left panels show the length-eight conditions while the right panels show the length-16 conditions. The numbers of shifts following failure are shown in blue for the neutral condition, and in green for the plentiful condition, while the numbers of shifts following reward are shown in gray. In all four conditions, Participant 1 shifts relatively often after a failure
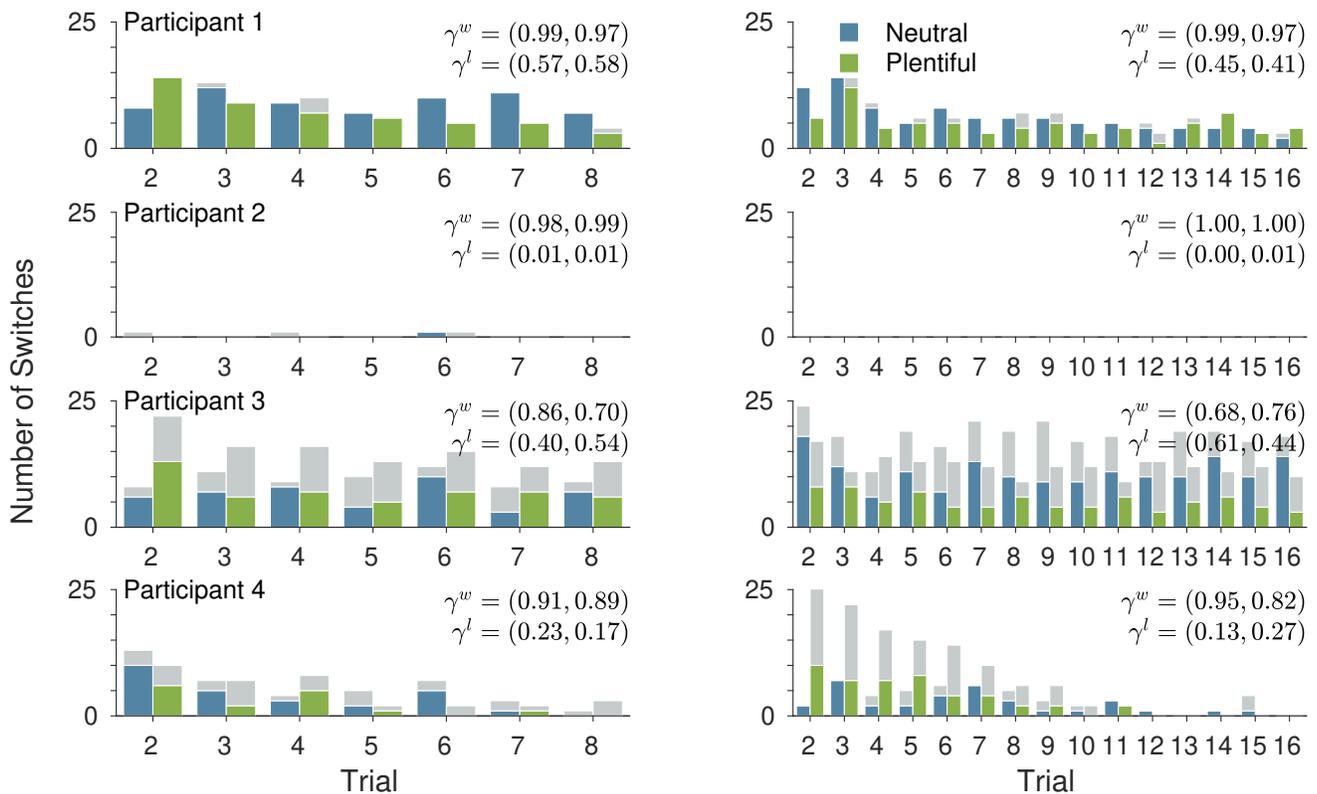
FIGURE 10: The number of shifts following reward versus failure for four representative participants in each of the four conditions. The left panels show the length-eight conditions while the right panels show the length-16 conditions. The numbers of shifts following failure are shown in blue for the neutral condition, and in green for the plentiful condition. The numbers of shifts following reward are shown in gray. The inferred γ_w and γ_l parameters for each participant in the the neutral and plentiful conditions are also shown.

but rarely after a reward. Participant 2 almost never shifts, either following a reward or a failure. Participant 3 shifts relatively more often following failure than Participant 2, but also shifts sometimes following a reward. Participant 4 shifts moderately often following both reward and failure for early decisions in the sequence, but shifts less often as the sequence progresses. The inferred γ_w and γ_l parameters for each participant in the the neutral and plentiful conditions are also shown, and correspond to the observed staying and shifting behavior.

Figure 11 shows the joint and marginal distributions of the posterior means of the γ^w and γ^l for each participant, for all four conditions. The four representative participants from Figure 10 are labeled. It is clear that the γ^w and γ^l parameters capture the consistent differences in their behavior observed in Figure 10. For example, Participant 1, who almost always stays after a reward and shifts after a failure, is consistently in the top right of the scatter plot, corresponding to high values of both the γ^w and γ^l parameters. In contrast, Participant 2, who rarely switches, is consistently located in the bottom right of the scatter plot, corresponding to a high value of the γ^w parameter and a small value of the γ^l

parameter.

Overall, it is clear that there is a range of individual differences in both win-stay and lose-shift probabilities, and that 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.

7   Questionnaires

Participants completed three questionnaires: the Risk Propensity Scale (Meertens & Lion, 2008), the Risk Taking Index (Nicholson et al., 2005), and the Domain Specific Risk Taking scale Blais & Weber (2006). The questions involved in these instruments are provided in the supplementary materials.
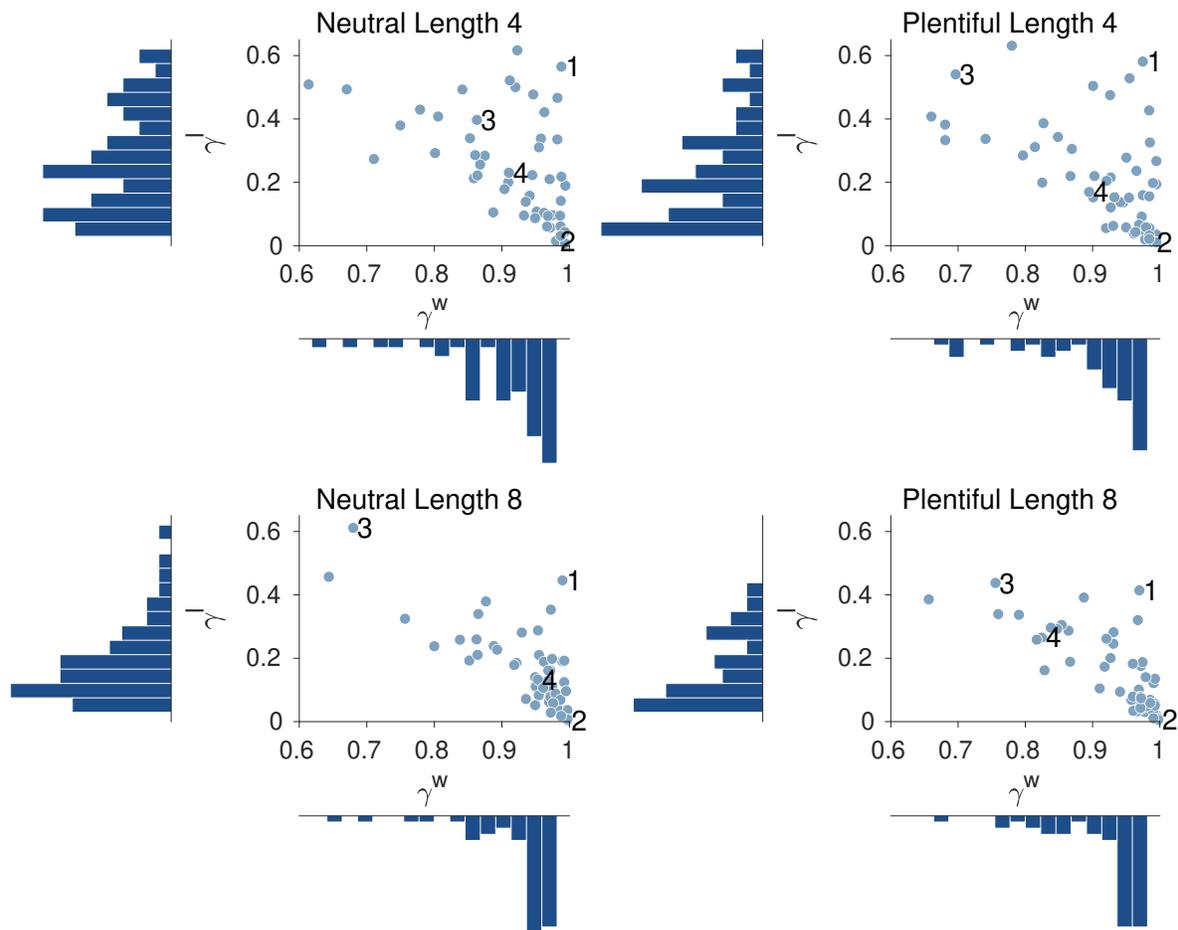
FIGURE 11: Joint and marginal distributions of the means of the γ^w and γ^l posterior expectations across the four conditions for each participant. The four representative participants shown from Figure 10 are labeled.

7.1 Risk Propensity Scale

The Risk Propensity Scale (RPS) was designed to be a short and easily administered test for measuring general risk-taking tendencies. The RPS originally consisted of only nine items, from which two items were later removed. The version of the RPS we use consists of the seven remaining items. All of the items involve statements that are rated on a nine-point scale ranging from "totally disagree" to "totally agree," except for the last item, which involves a nine-point rating from "risk avoider" to "risk seeker." Items 1, 2, 3, and 5 were reverse-scored so that high scores represented high risk propensity. Meertens & Lion (2008) reported an internal reliability coefficient measured by Cronbach's α of 0.77.

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 was taken. The left panel of Figure 12 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 & Lion (2008), who reported a mean score of 4.63 and standard deviation of 1.23. The Cronbach's α observed in this sample of 56 participants was 0.90.

7.1.1 Risk Taking Index

The Risk Taking Index (RTI) assesses overall risk propensity in six domains: recreation, health, career, finance, safety, and social. There is only one item for each of the six domains, but each item is answered twice: once for current attitudes, and once for past attitudes. All of the answers are given using a five-point Likert scale ranging from "strongly disagree" to "strongly agree."

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 the current and past contexts. 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. Nicholson et al. (2005) reported high internal consistency for the

general risk propensity scale with a Cronbach's α of 0.80. The left panel of Figure 12 shows the distribution of RTI scores across all 56 participants. The RTI scores are possibly bi-modal and range from 12 to 42. There is a large group of participants with a peak around 20 and a 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. The Cronbach's α observed in this sample of 56 participants was 0.84.

7.2 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 & Weber (2006) to be shorter and more broadly applicable. The original version was revised from 40 items down to 30 items, evaluating risky behavioral intentions originating from five domains: ethical, financial, health/safety, social, and recreational risks. Each domain involves six items.

The DOSPERT differs from the RPS and RTI in that it attempts to distinguish people's tendency to be risk seeking from people's perception of risk. Blais & 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, separating risk taking from risk perception. Participants rated each of the 30 statements in terms of self-reported likelihood of engaging in risky behaviors to measure risk taking, and in terms of their gut-level assessment of the riskiness of these behaviors to measure risk perception. In the risk-taking assessment, a seven-point rating scale was used, ranging from "extremely unlikely" to "extremely likely." In the risk-perception assessment, a seven-point rating scale was used ranging from "not at all risky" to "extremely risky".

Participants indicated their selection by checking the appropriate box under each number. Ratings were summed 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. Blais & Weber (2006) reported Cronbach's α's ranging from 0.71 to 0.86 for the risk-taking scores, and Cronbach's α values ranging from 0.74 to 0.83 for the risk-perception scores.

The right panel of Figure 12 shows the relationship between the risking taking and risk perception scores from the DOSPERT across all 56 participants, along with the marginal distributions of each. The risk taking scores also appear to

be slightly bi-modal, with a large group of participants centered around about 16–18 and then a smaller group near 30. Risk perception scores are unimodal and centered around 27. These results are consistent with the findings from Blais & Weber (2006), in the sense that there is a negative relationship between risk taking and risk perception scores ($r = -0.22$). The Cronbach's α observed in this sample of 56 participants was 0.92 for the overall risk-taking score, and 0.92 for the overall risk-perception score.

8 Correlation Analysis

Our main goal is to examine the relationship between the risk propensity and consistency parameters within and across tasks, and their relationship to the questionnaire measures. Before doing this, however, we compared the behavioral performance of participants within and across each cognitive task. Performance in the BART was computed as the average dollar amount collected on each problem. Performance in the gambling task was computed as the proportion of problems for which the participant chose the gamble with the maximum expected utility. Performance on the optimal stopping problem was computed as the proportion of problems where the participant correctly chose the maximum. Performance in the bandit task was computed as the average proportion of trials that resulted in reward.

Figure 13 shows the correlations of participant performance across each condition for all of the decision-making tasks. The area of the circles represent the magnitude of Pearson's correlation r, with blue circles representing positive correlations and red circles representing negative correlations. These empirical results suggest that participant performance is highly correlated within tasks, but that it is less strongly correlated across tasks.

8.1 Cognitive Task Overview

Table 1 provides an overview of the four decision-making tasks, models, and relevant parameters. The BART has two risk parameters, $\gamma_{1:2}^{+}$, and two consistency parameters, $\beta_{1:2}$. The gambling task has one risk parameter, λ, and one consistency parameter, ϕ. The optimal stopping task has eight risk parameters, $\beta_{1:4}$ and $\gamma_{1:4}$, and four consistency parameters, $\alpha_{1:4}$. The bandit task has four risk parameters, $\gamma_{1:4}^{l}$, and four consistency parameters, $\gamma_{1:4}^{w}$. In total, there are 26 relevant parameters from the decision-making tasks to be compared within and across tasks for each individual.

8.2 Estimating Correlations with Uncertainty

The correlations between each risk and consistency parameter from all four decision-making tasks were estimated, using a Bayesian approach, based on Lee & Wagenmakers (2013,
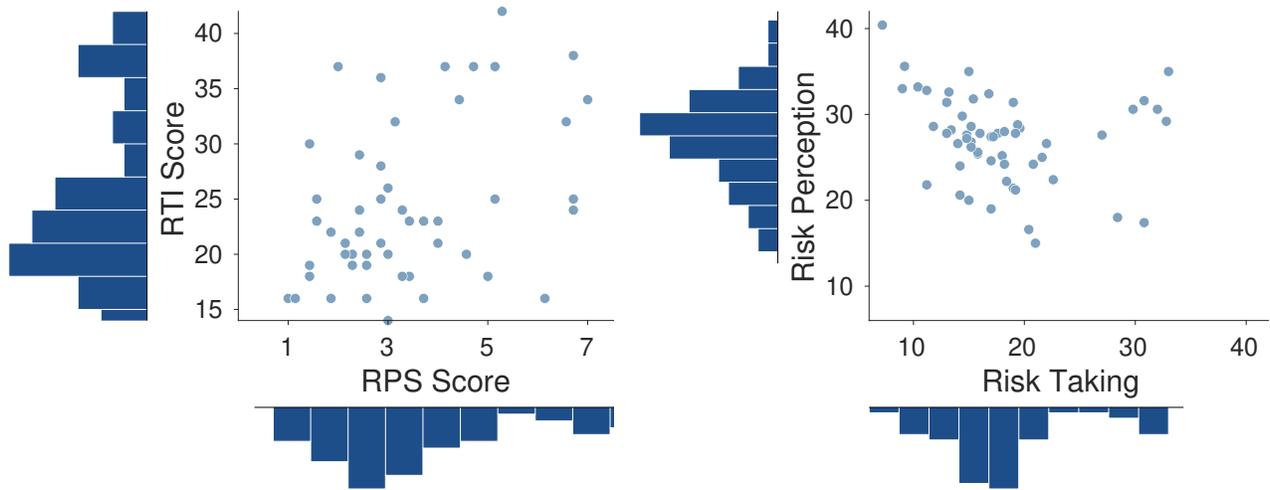
FIGURE 12: The distributions of questionnaire-based measures of risk. The left panel shows the joint and marginal distributions of RPS and RTI scores. The right panel shows the joint and marginal distributions of the DOSPERT risk taking and risk perception scores.

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	$\gamma_1^l, \ldots, \gamma_4^l$	$\gamma_1^w, \ldots, \gamma_4^w$
Gambling	2	CPT	λ	ϕ

TABLE 1: Overview of tasks and parameters.

Chap. 5). A key feature of this approach is that it incorporates uncertainty in the inferences of the parameters themselves (Matzke et al., 2017). That is, we do not use point estimates of the various risk and consistency parameters, but instead acknowledge that participant's behavior is consistent with a range of possible values, given the limited behavioral data. Our inferences about the correlations between parameters are thus sensitive to the precision with which their values are determined from the cognitive models and decision-making tasks we used.

Formally, for each pair of parameters, we correlate a set of samples for the ith participant, rather than just a single best estimate for each participant. These samples are generated by assuming Gaussian marginal posterior distributions

$$x_{ij} \sim \text{Gaussian}\left(y_{ij}, \lambda_j^e\right),$$

where $\mathbf{y}_i = (y_{i1}, y_{i2})$ represents the latent true value of the parameters, and $\boldsymbol{\lambda}^e = \left(\lambda_1^e, \lambda_2^e\right)$ denotes the precision of the inference about them. The precisions are estimated as the standard deviations of the marginal posterior distributions from the inferences of the decision-making models. The correlation focuses on the latent true values of the cognitive measures, by modeling them as a draw from a multivariate

Gaussian distribution

$$\mathbf{y}_i \sim \text{Gaussian}\left((\mu_1, \mu_2), \begin{bmatrix} \sigma_1^2 & r\sigma_1\sigma_2 \\ r\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}^{-1}\right).$$

Our hierarchical correlation model uses the following priors on r, σ_1, σ_2, μ_1, and μ_2:

$$
\begin{aligned}
r &\sim \text{uniform}(-1, 1) \\
\sigma_1^2, \sigma_2^2 &\sim \text{invGamma}(0.001, 0.001) \\
\mu_1, \mu_2 &\sim
\begin{cases}
\text{uniform}(0, 1) & \text{for OS } \alpha, \text{Bandit } \gamma^w, \gamma^l \\
\text{uniform}(0, 10) & \text{for BART } \beta, \gamma^+ \\
\text{Gaussian}(0, 0.001) & \text{for OS } \beta, \gamma \\
\text{uniform}(1, 9) & \text{for RPS} \\
\text{uniform}(10, 50) & \text{for RTI, DOSPERT.}
\end{cases}
\end{aligned}
$$

The correlation analysis was implemented as a Bayesian graphical model in JAGS. It was applied independently to all possible parameter combinations, inferring the posterior distribution of the correlation coefficient in each case. We generally use the posterior mean as a summary of the inference, but also use the Savage-Dickey method (Wetzels et al., 2010) to estimate Bayes factors to compare the hypotheses of correlation and no correlation.
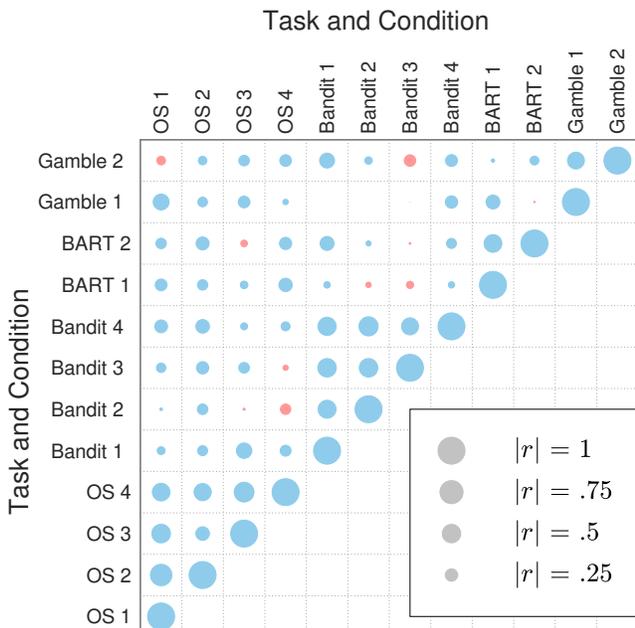
FIGURE 13: Pearson's correlations of performance across each condition in all of the decision-making tasks. Blue circles represent positive correlation, while red circles represent negative correlations. The areas of the circles correspond to the magnitudes of the correlations.

An advantage of Bayesian analysis is that it can find evidence in favor of a null hypothesis such as no correlation. Whereas null hypothesis significance testing can either find evidence for a correlation, or fail to find evidence for a correlation, the Bayesian analysis can produce three outcomes. These possible outcomes are evidence for a correlation, evidence for the absence or a correlation, or no strong evidence for either possibility. This is important in evaluating whether the data contain enough information to make meaningful claims about the correlations. To the extent that the Bayes factors provide evidence in favor of either the presence or absence or correlations, the data can be considered sufficiently powerful to have answered the research question. Evidence for the data being insufficient would be provided by Bayes factors that provide no strong evidence in either direction.

8.3 Correlation Results

Combining the scores from the three questionnaires to the parameters from the four decision-making tasks gives a total of 30 risk and consistency measures to be compared, which leads to 435 pairwise correlations. Figure 14 shows the results for all of these correlations. The dashed lines divide the grid into the three questionnaires and four decision-making tasks. The circles indicate parameter pairs for which the Bayes factor provides evidence of a correlation. We used a cutoff of 3 for the Bayes factor, because it is a standard

boundary corresponding what is variously labeled "substantial" (Jeffreys, 1961), "positive" (Kass & Raftery, 1995), and "moderate" (Lee & Wagenmakers, 2013) evidence.[1] The areas of the circles correspond to the magnitudes of the correlations, given by the posterior expectation of r. Blue circles indicate that a correlation is positive while red circles indicate that a correlation is negative. Meanwhile, the cross markers correspond to those comparisons where Bayes factor was at least 3 in favor of the null hypothesis of no correlation.

It is clear 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 are the risk parameters within the BART. This is clear from the patterns of blue circles along the diagonal. The positive correlations across conditions within the same task are expected, given the stability we observed in representative participants across conditions in the decision-making task analyses. Furthermore, the RTI, RPS, and DOSPERT RT are also positively correlated with each other, replicating previous findings.

There also appear to be some negative correlations between different parameters within tasks. For example, the γ^w and γ^l parameters in the bandit task are negatively correlated with each other, and the risk and consistency parameters in the BART are also negatively correlated. As we noted in the task-specific analyses, there is some trade-off between parameters for some of these tasks.

There appears, however, to be less evidence for systematic correlations across tasks. Indeed, there is generally evidence for a *lack* of correlation between parameters from different tasks, and between cognitive parameters and the questionnaire measures. The one exception relates to the gambling tasks parameters, for which there is no evidence for or against correlations with other cognitive parameters and questionnaire measures. This result likely reflects a failure of the experimental design to measure the risk aversion and consistency parameters with enough precision. In contrast, the results in Figure 14 show that there is enough information to make inferences, either in favor or against the presence of a correlation, for all of the other cognitive parameters and questionnaire measures. This finding speaks directly to the adequacy of the data to address the main research question about correlations between model parameters and questionnaire measures.

Figure 15 provides a different presentation of the correlation analysis that focuses on the comparisons for which there is evidence for correlations. Only pairs of parameters or measures with Bayes factors greater than 10 in favor of the alternative model are considered in this analysis, to focus on those pairs for which the evidence of correlation is strongest. The left panel shows the 95% Bayesian credible intervals of r for each comparison. The right panel shows

---

[1]See https://www.nicebread.de/grades-of-evidence-a-cheat-sheet/.
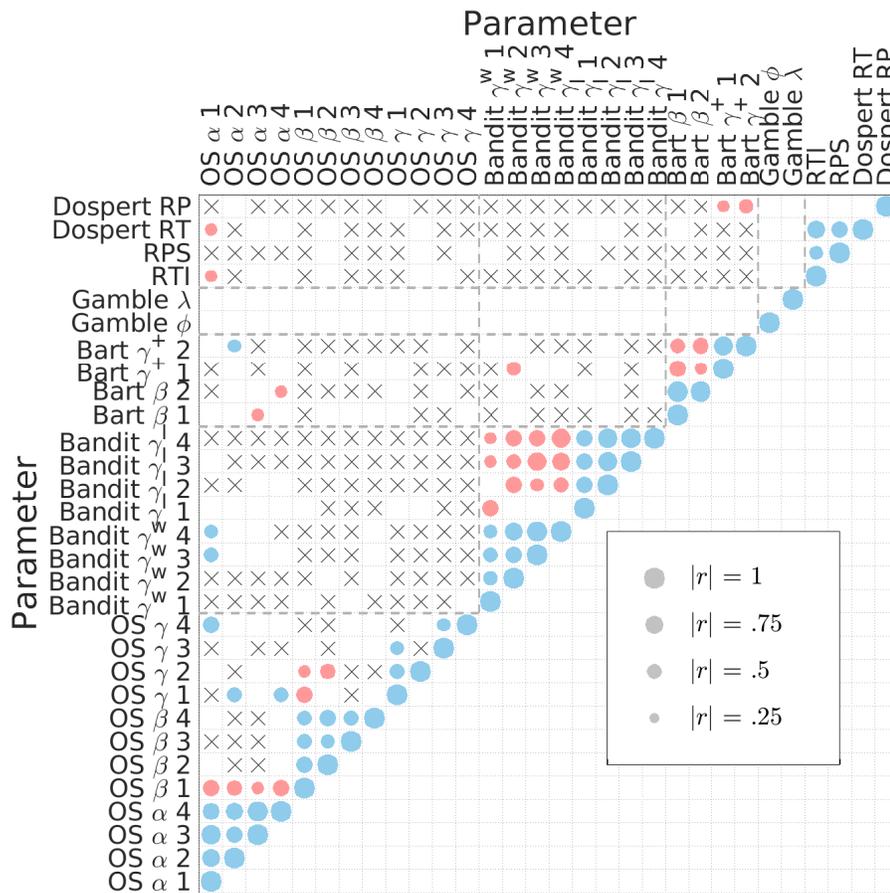
FIGURE 14: Correlation matrix of the risk and consistency parameters. Blue circles represent positive correlations for which the Bayes factor provided at least moderate evidence, while red circles represent negative correlations for which the Bayes factor provided at least moderate evidence. The areas of the circles correspond to the absolute values of the posterior expectation of the correlation r. Cross markers indicate that the Bayes factor provided at least moderate evidence for the absence of a correlation. The parameters within tasks are identified 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.

the log Bayes factors for the corresponding comparisons. The strong positive positive correlations between the same cognitive parameters across different conditions within tasks are clear, as are the trade-offs between different parameters within tasks, shown by the strong negative correlations.

9   Cognitive Latent Variable Analysis

The correlation analysis is one way to test the idea that there is a general risk factor underlying the cognitive parameters that control people's risk propensity on the cognitive tasks, and is also measured by the questionnaires. As a second complementary approach to testing the same idea, we explored the factorial structure of the tasks using a cognitive latent

variable model analysis (CLVM: Vandekerckhove, 2014; Pe et al., 2013). CLVMs are a broad category of models that involves a latent variable structure built on top of cognitive process models and other measures of behavior, to allow inference of latent variables that have higher-order cognitive interpretations.

A CLVM is defined by a factor matrix Φ, which contains a score ϕ_{fi} for each participant $i = 1, \ldots, I$ on each of F latent factors $f = 1, \ldots, F$, and a loadings matrix Ψ, which has F columns corresponding to latent dimensions or factors, and E rows $e = 1, \ldots, E$ corresponding to cognitive parameters or other behavioral measures. The values ψ_{ef} in the loadings matrices, corresponding to factor-parameter pairings, may be set to assume there is no association ($\psi_{ef} = 0$),
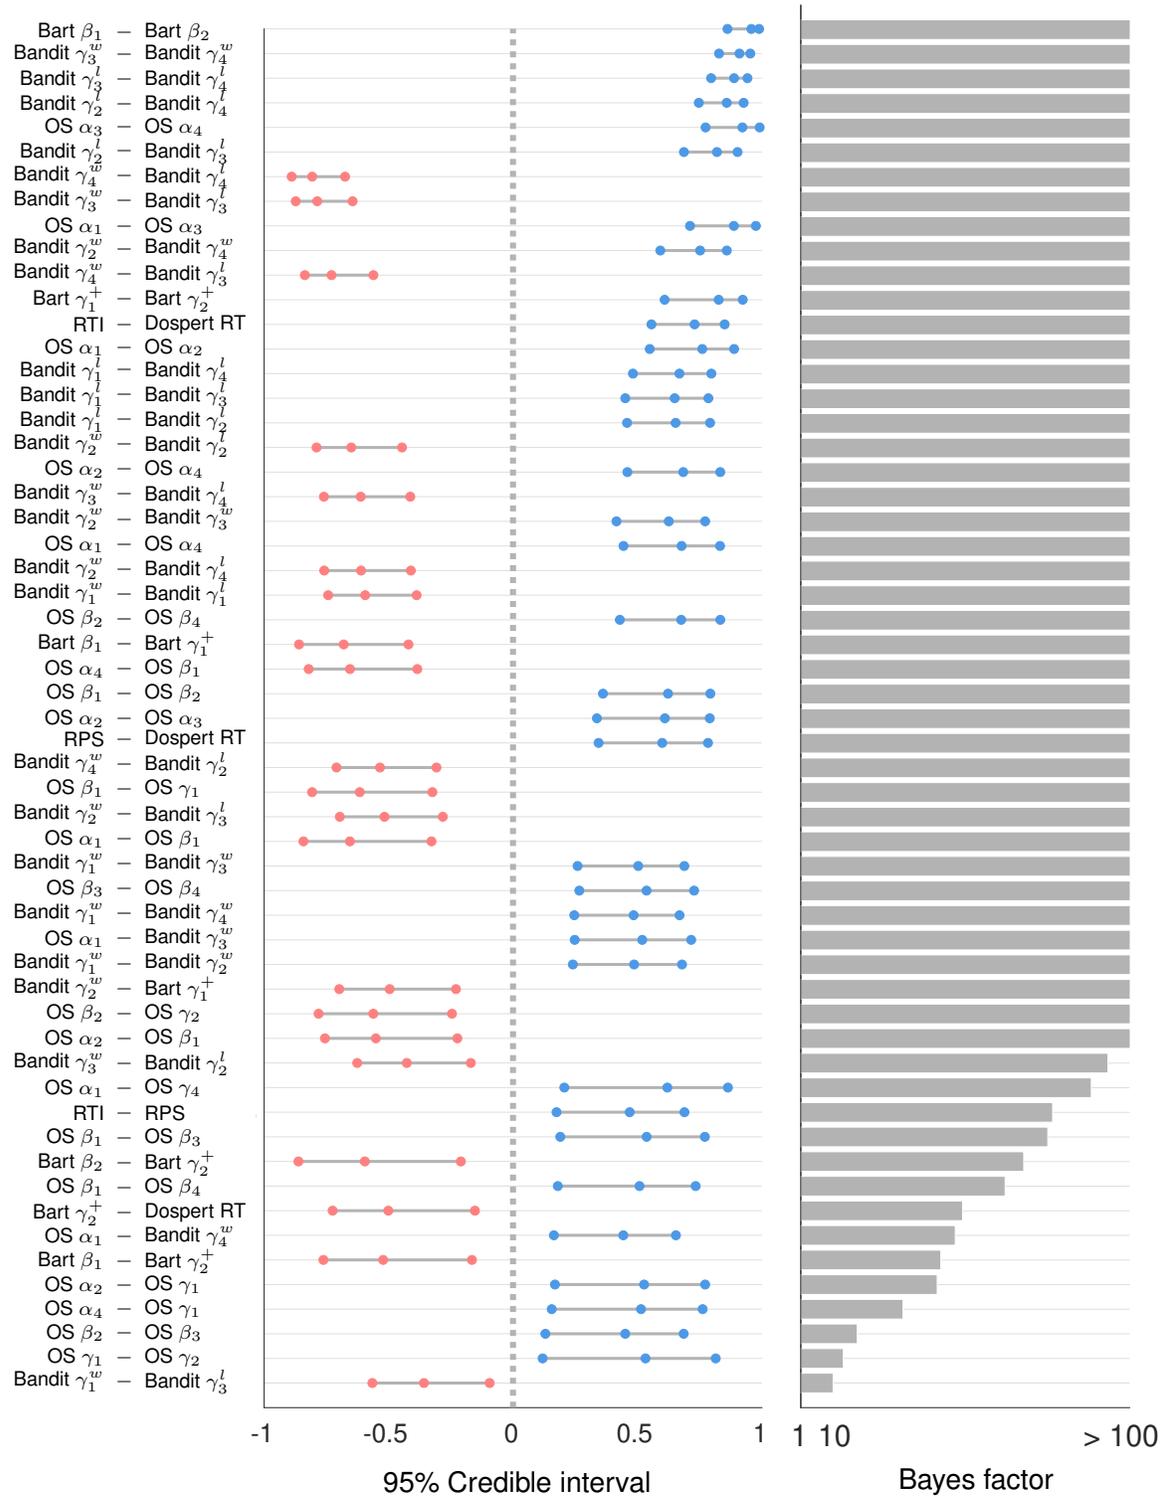
FIGURE 15: 95% Bayesian credible intervals (left panel) and Bayes factors (right panel) for Pearson's correlation coefficient r for parameter pairs with strong evidence of a correlation. Positive correlations are shown in blue and negative correlations are shown in red.

assume there is an association ($\psi_{ef} = 1$), or allow for the possibility there is some level of association to be inferred. These assumptions formalize different models of the factor structure underlying the relationships between the cognitive model parameters and questionnaire measures. Each cognitive model parameter and questionnaire measure e has an expected value given by the weighted average of all factors: $E(e_{ij}) = \sum_f^F \psi_{ef} \phi_{fi}$. The likelihood of the model is

$$e_{ij} \sim \text{Gaussian}\left(\sum_f^F \psi_{ef} \phi_{fi}, \lambda^e\right),$$

where the uncertainty λ^e is estimated as the standard deviation of the marginal posterior distribution of parameter e as obtained from the preceding analyses. In all cases, the latent factor scores have multivariate Gaussians priors with mean zero and precision matrix the identity matrix: $\phi_{\cdot,i} \sim \text{multivariateGaussian}(\mathbf{0}_{F\times 1}, \mathbf{1}_{F\times F})$. Similarly, the free loadings (i.e., those K loadings not constrained to be 0 or 1) were given the same multivariate Gaussian prior $\psi_{\cdot,\cdot} \sim \text{multivariateGaussian}(\mathbf{0}_{K\times 1}, \mathbf{1}_{K\times K})$.

We consider eight CLVMs. Three of these models capture what we believe are sensible theoretical positions, and three are based on the data and are exploratory in nature. The remaining two models are "bookend" models, which serve as reference points for assessing the merit of the substantive models based on theory and data (Lee et al., 2019).

9.1 Theory-based models

The first theoretical model is the "general risk" model. It has one latent factor for each cognitive model parameter, and combines their independent replication across experimental conditions. For example, with respect to the optimal stopping model, there is one factor for all four of the α error-of-execution parameters applied to the four experimental conditions, one factor for all of the β bias parameters, and one factor for all of the γ decrease parameter. The same separation and grouping of parameters applies to the other cognitive models. In addition, the general risk model has a general factor that all parameters share and is assumed to correlate with the risk surveys. The theoretical motivation for this model is based on the possibility that there is a general factor, which can be conceived as a risk propensity equivalent to the general intelligence factor "g" from cognitive abilities and psychometric testing. The general risk model emphasizes this general factor, while also allowing for the uniqueness of the cognitive tasks.

The left panel of Figure 16 details the structure of the general risk model. Rows represent the cognitive model parameters and questionnaire measures and columns represent the assumed factors. Dark blue squares indicate that a parameter or measure is assumed to load on a factor. Light yellow squares indicate that some level of association is possible. Empty squares assume a lack of association. Thus, the first factor has dark blue squares for the questionnaire measures, since these are assumed to index general risk, and light yellow squares for the cognitive parameters, allowing for the possibility they may also index risk. The remainder of the model structure loads each cognitive parameter in each task on a separate factor.

The second theoretical model is the "two-factor" model. It is a simpler model, with only two latent factors. The middle panel of Figure 16 details the structure of this model. One factor corresponds to risk propensity and the other corresponds to behavioral consistency. The risk propensity factor loads on the specific cognitive model parameters we interpret as controlling risk propensity in the tasks. These are the β bias and γ decrease parameters in the optimal stopping model, the γ^l lose-shift parameters in the extended-WSLS model, the γ risk propensity parameter in the BART model, and the λ loss aversion parameter in the cumulative prospect theory model. It also loads on the risk measures produced by the four questionnaires. The behavioral consistency factor loads on the other cognitive model parameters, which control the error of execution and response determinism within the models.

The third theoretical model is the "three-factor" model. It is detailed in the right panel of Figure 16. The three-factor model is an extension of two-factor model that loads the four questionnaire measures on a separate third factor, rather than on the risk propensity factor. This model was included to test the possibility of a difference between behavioral risk taking, as potentially expressed in the cognitive tasks, and self-reported risk taking, as measured by the questionnaires.

9.2 Exploratory models

The exploratory models were constructed based on inspection of the correlation analyses presented in Figure 14. We measure the performance of these models relative to two bookend models. The first bookend is the "unitary" model, which has a single latent factor for all cognitive model parameters and questionnaire measures. It is a very simple CLVM account of the data that provides a lower bound on the goodness-of-fit that can be achieved. The other bookend is the "saturated model", which has one latent factor for each of the 30 cognitive model parameters and questionnaire measures. It is the most complicated CVLM account of the data. It provides an upper bound on the goodness-of-fit. The role of bookend models is to provide comparison points for substantively interesting models. A useful substantive model should outperform both bookends in terms of a model evaluation measure that balances goodness-of-fit and complexity. In addition, requiring substantive mod-
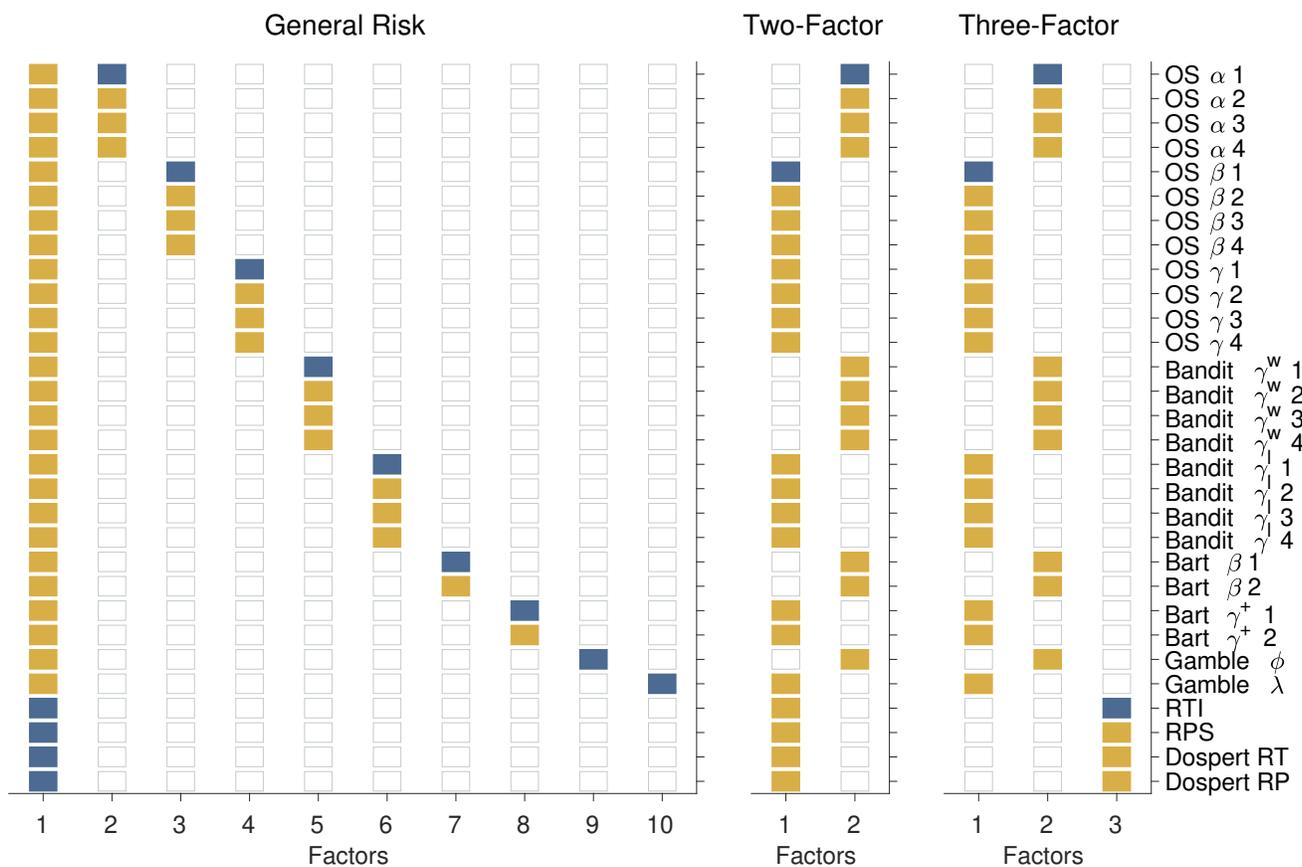
FIGURE 16: The factor loadings structure for the general risk (left panel), two-factor (middle panel), and three-factor (right panel) theory-based CLVMs. In each panel, rows correspond to cognitive model parameters and questionnaire measures, and columns correspond to model factors. Dark blue squares indicate an assumed association between a factor and a parameter or measure. Light yellow squares indicate a possible association, to be estimated. Empty squares indicate an assumed lack of association.

els to outperform the saturated model provides confidence that they are descriptively adequate, because their balance between goodness-of-fit and complexity is better than an account that has high descriptive adequacy. We use the Deviance Information Criterion (DIC: Spiegelhalter et al., 2002, 2014), which has theoretical limitations, but provides a useful practical measurement for a coarse-grained assessment of competing models.

The first exploratory model we found is the "questionnaires only" model. It simplifies the saturated model by assuming that a single latent factor underlies all four questionnaire measures, but that the cognitive model parameters continue to have their own factors. The second exploratory model is the "BART β" model. It simplifies the saturated model by assuming that a single latent factor underlies the two Bart β parameters. Finally, the "questionnaires and BART β" model combines the constraints of the first two exploratory models, so that a single factor underlies all of the four questionnaires and the BART β parameters.

9.3 CVLM results

Table 2 summarizes the results of the CLVM analysis. According to the DIC measure, none of the theory-based models performed well. The exploratory models lead to slight improvements. We could not find any other CLVM that improved on the unitary and saturated bookend models. These results are largely consistent with the results of the correlational analysis above: There is not much evidence for a jointly explanatory underlying structure between the cognitive tasks. Even within tasks, the CLVM analysis provides evidence for models with multiple underlying dimensions per task. Perhaps the most interesting exploratory finding from the CLVM analysis is that it is the BART task, and its associated cognitive model parameter measuring behavioral consistency, that most closely aligns with the measure of risk produced by the questionnaires.

Model	Factors	DIC	Δ DIC
General Risk	10	$62,221$	$62,859$
Two-factor	2	$170,725$	$171,363$
Three-factor	3	$170,737$	$171,375$
One-factor	1	$220,896$	$221,534$
Saturated	30	-588	50
Surveys only	27	-600	38
BART β	29	-611	27
Surveys, BART β	26	-618	0

TABLE 2: Results of the CLVM analysis. DIC = deviance information criterion. ΔDIC measures difference in DIC to the best-performed "Surveys, BART β" model.

10 Discussion

The goal of this article was to explore the psychological construct of risk propensity in the context of cognitive tasks and the inferred latent parameters of cognitive models that can be interpreted as the psychological variables that control risk seeking and risk avoiding behavior. We compared these measures of risk across four sequential decision-making tasks and measures obtained from more traditional questionnaires based on self-report. In each of the independent analyses of the four decision-making tasks we used a cognitive model that provided an adequate account of people's behavior. The inferred parameters of the cognitive models have natural interpretations as measures of risk propensity and decision-making consistency, and appear to capture stable individual differences across conditions within each task. The measures found using the questionnaires were generally consistent with previous studies, with similar means and standard deviations.

If risk propensity is a stable psychological construct that can be measured by these decision-making tasks, then the risk parameters and questionnaires are expected to correlate across tasks. We found strong within-task correlations and interpretable consistency in the key parameters for representative participants across task conditions. We did not, however, find evidence for any systematic between-task relationships consistent with stable underlying risk propensity or consistency traits in individuals. A complementary analysis based on cognitive latent variable modeling reached the same conclusion. The data provided no evidence for any model that incorporated an interpretable general risk factor that spanned the four cognitive tasks. There was some evidence for a relationship between cognitive models of risk propensity in the BART and the RPS, RTI, and DOSPERT scale measures. Of the four cognitive tasks we considered,

the BART has been the most widely used as a psychometric instrument for measuring risk propensity (e.g. Taşkın & Gökçay, 2015; White et al., 2008), including examining its correlation with questionnaire measures (e.g. Asher & Meyer, 2019; Courtney et al., 2012), and as a predictor of real-world risk-taking behavior Lejuez et al. (2003b, 2007).

Overall, however, our results do not find evidence for a common underlying risk trait. This lack of evidence arose despite the use of cognitive models to make inferences about latent parameters, rather than relying on simple behavioral measures. Similar findings of weak relationships between measures from behavioral tasks and questionnaires has been found in psychological research on individual differences in other domains such as the description-experience gap (Radulescu et al., 2020), self-regulation (Eisenberg et al., 2018), intelligence (Friedman et al., 2006), and theory of mind (Warnell & Redcay, 2019),

10.1 Limitations and Future Directions

An obvious potential limitation of this study is the relatively small sample size. Generally, studies focusing on individual differences use larger sample sizes, typically over 100 participants, with some studies recruiting many more than that (Eisenberg et al., 2018; Frey et al., 2017). A common reaction to our use of 56 participants is to question whether our experimental design was sufficiently "powerful" to address the research questions it aimed to answer. We think this question reflects a (widely-held) conceptual misunderstanding, sometimes called the power fallacy (Wagenmakers et al., 2015). Power is a pre-experimental concept and is not relevant once data have been collected. Power analyses consider, before data have been collected, the results an experimental design could produce, and whether those results would be informative. Once the data have been collected, the uncertainty is resolved, and it is not logical to continue considering what are now counterfactual possibilities. From a Bayesian perspective, scientific inferences should be conditioned on only the observed data.

This means that whether our data are sufficiently informative can be answered by the direct examination of the inferences they produce. The key results are presented in Figure 14, where it is shown that for the large majority of parameter pairs, the Bayes factor provides clear evidence in favor of either the presence or the absence of a correlation. The one exception, as we noted, is for the gambling task. Here, we believe the lack of evidence is caused by our use of relatively few conditions and trials compared to previous literature (Nilsson et al., 2011). All of the other tasks and measures, however, have sufficient information about the cognitive parameters and behavioral measures to answer our research questions. Thus, overall, we believe our results demonstrate that the experiment was well enough designed, had enough participants, and was completed by sufficiently

motivated participants, to address the research question of whether behavior on the task is controlled by a common underlying risk trait.

A different limitation of our study involves the specific cognitive models we used, and the details of how they were applied to the behavioral data. There are many other possible accounts of the BART, gambling behavior, optimal stopping, and bandit problem decision making. We referenced a number of alternative models for each task before we presented the model we used. While our models provide reasonable starting points, there are clearly many alternative models that could be explored. Similarly, we made practical choices about contaminant behavior that could be extended or improved. Different modeling possibilities are not limited to just different assumptions about cognitive processes. Alternative cognitive models could also be explored by considering more informative priors, which corresponds to making different assumptions about the psychological variables controlling the processes (Lee & Vanpaemel, 2018). As one concrete example, it could be reasonable to in the extended WSLS model of bandit problem behavior to assume that the probability of winning and staying is greater than the probability of losing and shifting. This order constraint would lead to more informative priors. As another example, it is probably possible to develop better priors for the BART task than the uniform priors we used, by seeking choices that lead to empirically reasonable prior predictive distributions (Lee, 2018).

We did not attempt to use common-cause models that capture the consistency of individuals across conditions for the same decision-making task (Lee, 2018). This has previously been done successfully for the specific BFO model of optimal stopping (Guan et al., 2015), and could likely be done for the other models we used. Indeed, the consistency of within-participant parameters across conditions for the same tasks makes this an obvious extension. Common-cause modeling could easily be implemented hierarchically in the graphical modeling framework we used, and would have the advantage of reducing the number of risk and consistency parameters to one per task, rather than one per condition. The parameters should also be more precisely measured, because they would be based on the entirety of each participant's behavior in a task. On the other hand, we would expect this commonality to emerge from the cognitive latent variable modeling we conducted, and so we think it is likely that there simply is no evidence for the common construct in our data and modeling analysis.

While all of the decision-making tasks we used were sequential decision-making tasks involving risk and uncertainty, there are fundamental differences between them. There is debate about exactly whether and how the tasks and questionnaires measures risk propensity (e.g. De Groot & Thurik, 2018), and even more scope for debate about whether and how the cognitive model parameters relate to

the relevant psychological concepts. As such, there is no clear consensus that either the tasks or the cognitive models we used capture risk propensity and consistency in the same way, or capture it at all. What we did is choose tasks that depend on risk seeking and avoidance in some way, and provide a rationale for the interpretations of the cognitive modeling parameters in terms of risk propensity.

A finer-grained version of this general issue is that the different cognitive tasks provide information about risk and uncertainty in different ways, and these differences could affect the way any latent risk construct is able to be inferred. 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. The preferential choice gambling task requires people 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. Meanwhile, the BART only provides feedback when a balloon bursts, and by keeping track of the total banked amount over problems. These nuances suggest that each of the decision-making tasks require related but different cognitive processes. It is thus entirely plausible that risk seeking or avoidance 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 psychologically equivalent to balancing exploration and exploitation in a bandit task.

Collectively, these sorts of considerations raise the issue of whether risk propensity can usefully be salvaged as a multi-dimensional construct. While we sought a single latent trait to explain individual differences across the tasks, it is possible that how people manage risk is better conceived in terms of a few inter-related but distinct traits. Theoretically, of course, this is a slippery slope. As the number of traits expands to match the number of tasks, the usefulness of the notion of an underlying risk propensity controlling behavior is lost. It becomes better understood as a temporary psychological state than a permanent psychological trait.

10.2   Conclusion

We used cognitive models to analyze four sequential decision-making tasks that are sensitive to people's propensity for risk. We found stable individual differences within tasks for model parameters corresponding to the psychological variables of risk and consistency. However, we found little evidence for commonality or stability when we compared conceptually similar parameters across the tasks. In addition, we found little evidence for any meaningful relationships between the model-based measures of risk and standard widely-used questionnaires for measuring risk propensity based on

self-report. Our 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 cognitive behavioral tasks.

Acknowledgements

References

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

Asher, N. B. & Meyer, J. (2019). Three Dimensions of User Risk-Taking: Individual Differences in the TriRB. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 63 (pp. 267–271).: SAGE Publications Sage CA: Los Angeles, CA.

Banks, J., Olson, M., & Porter, D. (1997). An experimental analysis of the bandit problem. *Economic Theory*, *10*, 55–77.

Baron, J., Badgio, P., & Gaskins, I. W. (1986). Cognitive style and its improvement: A normative approach. In R. J. Sternberg (Ed.), *Advances in the Psychology of Human Intelligence*, volume 3 (pp. 173–220). Hillsdale, NJ: Erlbaum.

Baumann, C., Gershman, S. J., Singmann, H., & von Helversen, B. (in press). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Science*.

Baumann, C., Singmann, H., Kaxiras, V. E., Gershman, S., & von Helversen, B. (2018). Explaining Human Decision Making in Optimal Stopping Tasks. In C. Kalish, M. Rau, J. Zhu, & T. T. Rogers (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 1341–1346). Austin, TX: Cognitive Science Society.

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

Berg, J., Dickhaut, J., & McCabe, K. (2005). Risk preference instability across institutions: A dilemma. *Proceedings of the National Academy of Sciences*, *102*, 4209–4214.

Blais, A.-R. & Weber, E. U. (2006). A domain-specific risk-taking (DOSPERT) scale for adult populations. *Judgment and Decision Making*, *1*, 33–47.

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review*, *113*, 409–432.

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

Busemeyer, J. R. & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432–459.

Campbell, J. & Lee, M. D. (2006). The Effect of Feedback and Financial Reward on Human Performance Solving 'Secretary' Problems. In R. Sun (Ed.), *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (pp. 1068–1073). Mahwah, NJ: Erlbaum.

Cavanagh, J. F., et al. (2012). Individual differences in risky decision-making among seniors reflect increased reward sensitivity. *Frontiers in Neuroscience*, *6*, 111.

Christian, B. & Griffiths, T. (2016). *Algorithms to live by: The computer science of human decisions*. New York, NY: Henry Holt and Co.

Courtney, K. E., et al. (2012). The relationship between measures of impulsivity and alcohol misuse: an integrative structural equation modeling approach. *Alcoholism: Clinical and Experimental Research*, *36*, 923–931.

Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*, 876–879.

De Groot, K. & Thurik, R. (2018). Disentangling risk and uncertainty: When risk-taking measures are not about risk. *Frontiers in Psychology*, *9*, 2194.

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

Deary, I. J. (2020). *Intelligence: A very short introduction*. Oxford: Oxford University Press.

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

Eisenberg, I. W., Bissett, P. G., Enkavi, A. Z., Li, J., MacKinnon, D., Marsch, L., & Poldrack, R. (2018). Uncovering mental structure through data-driven ontology discovery. *PsyArXiv*, (pp. 1–10).

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

Figner, B., Mackinlay, R. J., Wilkening, F., & Weber, E. U. (2009). Affective and deliberative processes in risky

choice: age differences in risk taking in the Columbia Card Task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 709–730.

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

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

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

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

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science*, *66*, 1375–1394.

Guan, M. & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*, 339–361.

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

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* (pp. 824–829). Austin, TX: Cognitive Science Society.

Harrison, J. D., Young, J. M., Butow, P., Salkeld, G., & 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*, 1385–1396.

Holmes, M. K., et al. (2009). Conceptualizing impulsivity and risk taking in bipolar disorder: importance of history of alcohol abuse. *Bipolar Disorders*, *11*, 33–40.

Hopko, D. R., Lejuez, C., Daughters, S. B., Aklin, W. M., Osborne, A., Simmons, B. L., & 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*, 95–101.

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

Jeffreys, H. (1961). *Theory of Probability*. Oxford, UK:

Oxford University Press.

Josef, A. K., Richter, D., Samanez-Larkin, G. R., Wagner, G. G., Hertwig, R., & Mata, R. (2016). Stability and change in risk-taking propensity across the adult life span. *Journal of Personality and Social Psychology*, *111*, 430.

Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *XLVII*, 263–291.

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

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., & 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*, 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. (2018). Bayesian methods in cognitive modeling. In J. Wixted & E.-J. Wagenmakers (Eds.), *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience. Volume 5: Methodology* chapter 2, (pp. 37–84). John Wiley & Sons, fourth edition.

Lee, M. D. & Courey, K. A. (in press). Modeling optimal stopping in changing environments: A case study in mate selection. *Computational Brain & Behavior*.

Lee, M. D., et al. (2019). Robust modeling in cognitive science. *Computational Brain & Behavior*, *2*, 141–153.

Lee, M. D. & Vanpaemel, W. (2018). Determining informative priors for cognitive models. *Psychonomic Bulletin & Review*, *25*, 114–127.

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

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

Lejuez, C., Aklin, W., Daughters, S., Zvolensky, M., Kahler, C., & Gwadz, M. (2007). Reliability and validity of the youth version of the Balloon Analogue Risk Task (BART–Y) in the assessment of risk-taking behavior among inner-city adolescents. *Journal of Clinical Child and Adolescent Psychology*, *36*, 106–111.

Lejuez, C., Aklin, W. M., Jones, H. A., Richards, J. B., Strong, D. R., Kahler, C. W., & Read, J. P. (2003a). The balloon analogue risk task (BART) differentiates smokers and nonsmokers. *Experimental and Clinical Psychopharmacology*, *11*, 26–33.

Lejuez, C., Simmons, B. L., Aklin, W. M., Daughters, S. B., & Dvir, S. (2004). Risk-taking propensity and risky sex-

ual behavior of individuals in residential substance use treatment. *Addictive Behaviors*, *29*, 1643–1647.

Lejuez, C. W., Aklin, W. M., Zvolensky, M. J., & Pedulla, C. M. (2003b). Evaluation of the Balloon Analogue Risk Task (BART) as a predictor of adolescent real-world risk-taking behaviours. *Journal of Adolescence*, *26*, 475–479.

Lejuez, C. W., et al. (2002). Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, *8*, 75–84.

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

Loomes, G. & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, *92*, 805–824.

Mata, R., Frey, R., Richter, D., Schupp, J., & Hertwig, R. (2018). Risk preference: A view from psychology. *Journal of Economic Perspectives*, *32*, 155–72.

Matzke, D., Ly, A., Selker, R., Weeda, W. D., Scheibehenne, B., Lee, M. D., & Wagenmakers, E.-J. (2017). Bayesian inference for correlations in the presence of measurement error and estimation uncertainty. *Collabra: Psychology*, *3*, 25.

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

Mehlhorn, K., et al. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*, 191–215.

Meyer, R. J. & Shi, Y. (1995). Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. *Management Science*, *41*, 817–834.

Mishra, S., Lalumière, M. L., & 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*, 616–621.

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

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

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

Pedroni, A., Frey, R., Bruhin, A., Dutilh, G., Hertwig, R., & Rieskamp, J. (2017). The risk elicitation puzzle. *Nature Human Behaviour*, *1*, 803–809.

Pleskac, T. J. (2008). Decision making and learning while taking sequential risks. *Journal of Experimental Psychol-*

ogy: Learning, Memory, and Cognition, *34*, 167–185.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *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*, 323–343.

Radulescu, A., Holmes, K., & Niv, Y. (2020). On the convergent validity of risk sensitivity measures. psyArxiv. https://psyarxiv.com/qdhx4.

Rao, H., Korczykowski, M., Pluta, J., Hoang, A., & 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*, 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., & Mellers, B. A. (2006). Extending the bounds of rationality: Evidence and theories of preferential choice. *Journal of Economic Literature*, *44*, 631–661.

Rieskamp, J. & Otto, P. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.

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

Russo, J. E. & Dosher, B. A. (1983). Strategies for multiattribute binary choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 676–696.

Schmitz, F., Manske, K., Preckel, F., & Wilhelm, O. (2016). The multiple faces of risk-taking: Scoring alternatives for the balloon-analogue risk task. *European Journal of Psychological Assessment*, *32*, 17–38.

Schonberg, T., Fox, C. R., & Poldrack, R. A. (2011). Mind the gap: bridging economic and naturalistic risk-taking with cognitive neuroscience. *Trends in Cognitive Sciences*, *15*, 11–19.

Seale, D. A. & 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. & 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. & 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*, 1573–1592.

Spiegelhalter, D. J., Best, N. G., Carlin, B., & van der Linde, A. (2002). Bayesian Measures of Model Complexity and Fit (with Discussion). *Journal of the Royal Statistical Society*, *64*, 583–640.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Linde, A. (2014). The deviance information criterion: 12 years on. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *76*, 485–493.

Stewart Jr, W. H. & Roth, P. L. (2001). Risk propensity differences between entrepreneurs and managers: A meta-analytic review. *Journal of Applied Psychology*, *86*, 145–153.

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

Sutton, R. S. & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: The MIT Press.

Szrek, H., Chao, L.-W., Ramlagan, S., & Peltzer, K. (2012). Predicting (un) healthy behavior: A comparison of risk-taking propensity measures. *Judgment and decision making*, *7*, 716–727.

Taşkın, K. & Gökçay, D. (2015). Investigation of risk taking behavior and outcomes in decision making with modified BART (m-BART). In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)* (pp. 302–307).: IEEE.

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

Van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2011). Cognitive model decomposition of the BART: Assessment and application. *Journal of Mathematical Psychology*, *55*, 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. & Morgenstern, O. (1947). *Theory of games and economic behavior, 2nd rev*. Princeton, NJ: Princeton University Press.

Wagenmakers, E.-J., et al. (2015). A power fallacy. *Behavior Research Methods*, *47*, 913–917.

Wallsten, T. S., Pleskac, T. J., & Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychological Review*, *112*, 862–880.

Warnell, K. R. & Redcay, E. (2019). Minimal coherence among varied theory of mind measures in childhood and adulthood. *Cognition*, *191*, 103997.

Weber, E. U., Blais, A.-R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, *15*, 263–290.

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

White, T. L., Lejuez, C. W., & de Wit, H. (2008). Test-retest characteristics of the Balloon Analogue Risk Task (BART). *Experimental and Clinical Psychopharmacology*, *16*, 565–570.

Zeigenfuse, M. D. & Lee, M. D. (2010). A general latent assignment approach for modeling psychological contaminants. *Journal of Mathematical Psychology*, *54*, 352–362.

Zhang, S. & Lee, M. D. (2010a). Cognitive models and the wisdom of crowds: A case study using the bandit problem. In R. Catrambone & S. Ohlsson (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 1118–1123). Austin, TX: Cognitive Science Society.

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

Zhou, R., Myung, J., & Pitt, M. (2019). Hierarchical Bayesian models of choice decisions in sequential risk-taking tasks. In *Proceedings of the 52nd Annual Meeting of the Society for Mathematical Psychology* (pp.1̃6).