Play It Safe or Take a Risk? Computational Modeling & Statistical Inference for the Effect of Emotional Valence on Risk-Taking

Play it Safe or Take a Risk? Computational Modeling & Statistical Inference for the Effect of Emotional Valence on Risk-Taking

A thesis presented
by
Jessica Li Huang
to
The Department of Computer Science &
The Department of Statistics

in partial fulfillment of the requirements
for the degree of
Bachelor of Arts
in the subject of
Computer Science & Statistics

Harvard University
Cambridge, Massachusetts
April 2020

Play it Safe or Take a Risk? Computational Modeling & Statistical Inference for the Effect of Emotional Valence on Risk-Taking

Abstract

Decision making is an important part of life. In the past 50 years, there has been a steady rise in the number and percentage of decision-making papers which also investigate the role of emotions. In this thesis, we are interested in whether the pleasantness of an emotion, or emotional valence, affects risk-taking decisions. This research is important for two main reasons: to better understand a component of the psychological phenomenon of decision making and to inform treatments and interventions for disorders which are symptomatic of altered emotion and decision making (e.g. depression, gambling addiction).

In the literature review, we present comprehensive scientific frameworks for explaining risk-taking and emotional valence during a risk-taking task. We also examine three hypotheses about how emotional valence affects gambling: the mood-maintenance hypothesis, the affect-infusion model, and the reward processing hypothesis. We also explore hurdles in incorporating these studies into mathematical models, which motivate our computational modeling and statistical analysis.

In our risk-taking task and data chapters, we explain why we chose the set of participants and the experimental setup, and we carefully profile participant behavior during the risk-taking task, to incorporate into our behavioral models.

In our theoretical developments and exploratory data analysis, we validate our data, we select covariates, we construct models of risk-taking which are scientifically informed and which demonstrate promise to learn three scientific hypotheses about how emotional valence affects risk-taking (mood-maintenance hypothesis, affect infusion model, reward processing hypothesis), and we test the stability of our models across regularization and resampling.

In our hypothesis test chapter, we build a scientifically informed statistical hypothesis test to begin to answer: does emotional valence affect risk-taking? This hypothesis test relies on conditional randomizations to generate empirical null distributions which boosts power relative to using estimated null distributions. This hypothesis test has advantages of comprehensively accounting for covariates which might affect risk-taking and integrating scientific hypotheses (the mood-maintenance hypothesis, the affect-infusion model, and the reward processing hypothesis) about how emotional valence affects risk-taking. One disadvantage of our hypothesis test is its large computational cost.

Finally, in our conclusion, we propose future computational, statistical, and psychological work to answer our research question, with respect to solving prevailing challenges about computational cost, statistical power, bidirectional relationships, going beyond psychological tasks and towards real-world utility, and modeling the effects of additional components of emotion beyond valence.

Contents

List of Figures

List of Tables

This thesis is dedicated to Susan Li and Zhimin Huang.

Acknowledgments

I am grateful for the opportunity to work with my advisors. Thank you Dr. Charles Zheng for training me and inspiring me to become a better statistician and better person. Thank you Dr. Lucas Janson for being my statistics mentor and for shining a light on the world of statistics within and beyond my project. Thank you Dr. Francisco Pereira for giving me an opportunity to work on this project and with your team. Thank you also for providing perspective on research and life. Thank you Dr. Boaz Barak for teaching me, encouraging me, and breaking down "impossible" barriers in research.

Thank you to the many people who taught me about computer science and statistics. Dr. Joseph Blitzstein, Dr. Jelani Nelson, Dr. Susan Murphy, Dr. Finale Doshi-Velez, Dr. Mark Glickman, Dr. Alexander Young, Dr. Patrick McClure, Dr. Kevin Rader, Dr. David Parkes, Dr. Salil Vadhan, Dr. Stuart Shieber, Dr. David Malan, Dr. Neil Shepard, Dr. Xilin Shen, Dr. Oliver Knill, Matthew Alpert, Christopher Graves, Shane Liebenthal, and Dr. Walter Dempsey.

Thank you to the many people who encouraged me to grow a passion for healthcare: Dr. Ernest Moy, Dr. Anas El Turabi, Dr. Hanna Keren, Dr. Argyris Stringaris, Dr. David Mou, Dr. Carine Simon, Grace Lam, Dr. Serena McCalla, Victoria Smith, Merianne Spencer, Dr. Cassandra Peitzman, Giada Tridello-Mattern, Dr. Nancy Long Sieber, Dr. Leah Somerville, Dr. Amy Bucher, and Dr. Joshua Buchholtz.

Thank you to my mentors in academia: Dr. John Huth, Dr. Melissa Haskell, Dr. Ella Striem-Amit, Dr. Emily Finn, Dr. Robert Constable, Dr. Matthew Nock, Dr. Nada Amin, Pamela Nwakanma, Dr. Michael Parzen, and Dr. Stephen Chong.

Finally, I am grateful for my friends and my family for their crucial contributions to helping me grow as a person and as a researcher.

Chapter 1

Introduction

Decision making is an important part of life. Some researchers study decision making by examining gambling, or the process of deciding among choices with quantifiable rewards that involve risk. Broadly, we can examine whether people play it safe by choosing an option with minimal risk or if people take a gamble by choosing an option with higher risk.

In the past 50 years, there has been a steady rise in the number and percentage of decision-making papers which also investigate the role of emotions. Studies have revealed how emotions have surprising and important effects on decision making [1]. In this paper, we will specifically focus on the valence component of emotion. Valence is a measurement of the pleasantness of the subjective feeling state of an emotion. Valence fails to capture the arousal dimension of the subjective feeling state of emotions. For example, excited and calm are two vastly different emotions because excited is a high arousal emotion and calm is a low arousal emotion; yet, valence classifies them as the same because both are pleasant. Additionally emotional valence fails to capture other complexities of emotion such as physiological response, expressions, or cognitive appraisals. While valence is a crude summary of emotion, it is still popular, useful, and powerful for investigations of decision making and emotion [1].

In this thesis, we are interested in whether emotional valence affects risk-taking. Scientific literature has a lot to say about the possible effects of emotional valence on risk-taking. In this literature review and for the rest of this thesis, we will examine three possible effects of emotional valence on risk-taking. In the following chapter, we will present a summary of these three possible effects and propose a multidisciplinary approach to deciding whether emotional valence affects risk-taking.

We shall start with a comprehensive look at literature about emotional valence and literature about risk-taking, both of which will be crucial for our later experimental design and data analysis.

1.1 Scientific Framework for Emotional Valence

Emotional valence is just one small component of emotions. Emotions have been defined as coordinated responses to eliciting events that manifests on multiple levels [2] [3], including "affect (subjective experiences of valence and arousal), physiology (arousal and stress responses via the peripheral nervous system), expression (facial, verbal, or action tendency), and in appraisals (cognitive evaluations of significance to self)." While it is only one component of the broader set of responses associated with emotion, emotional valence is a crude but powerful estimate of emotions such that most literature about emotions and decision making have implicitly or explicitly taken a valence-based approach [1]. This thesis focuses on emotional valence, but in the conclusions, we discuss more comprehensively measuring emotions.

Before attempting to determine whether emotional valence affects risk-taking, we should first review what is known about the causes of emotional valence. Many variables contribute to variation in emotional valence which may also have indirect effects on risk-taking. Without considering what affects emotional valence, we risk misattributing indirect effects on risk-taking from other covariates to emotional valence. Instead, by exploring scientific literature for variables that contribute to variation in emotional valence, we can condition on these variables to account for their indirect effects on risk-taking and to isolate the variation in emotional valence that contributes to variation in risk-taking.

Average emotional valence levels can depend on many covariates which our thesis aims to beging to control for: age [4], gender [5], depression status. Moreover, one core symptom of depression is anhedonia, which is the loss of interest in pleasurable activities or reduced ability to experience pleasure [6], which includes a lower average emotional valence for participants with depression. Notably, our hope is to understand whether emotional valence affects risk-taking within a framework that is general enough to include both healthy controls and patients with depression, so we consider both in this thesis. Considering both healthy controls and patients with depression is consistent with how there is very little evidence that depression is a true category. Instead, in this thesis, we consider depression diagnosis and depression severity as an indication of risk of symptoms associated with depression instead of a category. In this thesis, we begin to control for some effects of age, gender, and depression severity on emotional valence, helping us isolate whether emotional valence itself affects gambling.

Changes in emotional valence can depend on many things which can confound the relationship between emotional valence and risk-taking. In this thesis, we focus on changes to emotional valence in response to rewards. We first define some notation. In the midst of a gambling task, if someone decides to gamble at time t, they have an expectation of what rewards the gamble will return (E_t), then at time $t+1$, once the reward is given (R_{t+1}), we can define the prediction error as

$$PE_{t+1} = R_{t+1} - E_t$$

the difference between an expected and received value. In this way, a larger-than-expected reward will yield a positive PE, a smaller-than-expected reward will yield a negative PE, and a fully expected reward will yield a $PE = 0$. Prediction error is thought to be biologically meaningful because it has been shown to be encoded by dopamine neurons in animals [7] [8] and humans [9]. [10] showed that momentary emotional valence during a rewards task is informed less so by the task earnings, R and more so by "the cumulative influence of recent reward expectations E and prediction errors PE resulting from those expectations." In this thesis, we will begin to control for the effects on emotional valence of expected rewards E, actual rewards R, reward prediction error PE, and baseline emotional valence (the intercept) in an effort to isolate whether emotional valence affects gambling.

In this thesis, we will integrate this scientific information about emotional valence into a mathematical model of emotional valence to ultimately try to uncover whether emotional valence affects risk-taking.

1.2 Scientific Framework for Risk-Taking

Risk-taking is a broad subject and in this thesis we will focus on lab measurements of risk-taking which are not longitudinal and not ambiguous. That is, we deem studies outside the scope of this literature review if the decisions being examined are embedded in and related to ongoing activities in the participant's life, because in the lab, we are measuring decisions that are not connected to the participants ongoing activities. Ideally speaking, in the lab, information is not ambiguous, as possible rewards are explicitly quantified. Preferences are not ambiguous because they are aligned so that choices that yield more reward are more preferred. While this is limiting, it is still informative to examine this subset of risk-taking, and we will discuss broadening our perspective in the conclusions.

Crucially, the relationship between emotional valence and risk-taking does not exist in a vacuum, such that altering other confounding variables can dissolve, strengthen, or even reverse the relationship. On the other hand, a lack of consideration of many potential covariates of risk-taking can yield a misrepresentation of an effect.

As a result, we rely on existing scientific literature about risk-taking to inform us about covariates that contribute to variation in risk-taking. By understanding how these covariates contribute to variation in risk-taking, we can control for these effects and better isolate the effect of emotional valence.

Some covariates that alter risk-taking concern the risk-taker's characteristics. For example, males take more risks than females [11]. Risk-taking has been shown to decrease with age [12]. Altered risk-taking is a symptom of many disorders (e.g. alcoholism, drug addiction, depression). In this thesis, we focus on not only healthy controls but also subjects that have been diagnosed with depression. While one of the core symptoms of depression is altered emotional valence, another area of concern for patients with depression is altered risk-taking behavior. For example, stress can impair dopamine systems, which in turn might increase probability of developing depression as well as increase risk-taking behavior as a means to compensate [13]. This is just one example of how a third variable (stress) can increase both risk-taking behavior and depression risk. Moreover, dopamine depletion, which is a biological marker of depression, has been shown to significantly decrease risk-taking in a gambling task [14], indicating that depression can also cause a reduction in risk-taking behavior. Also, for children of parents who have depression, low risk-taking during a gambling task was predictive of depressive symptoms [15], indicating that risk-taking behavior is linked to depression. In this thesis, we will build frameworks which integrate information about how age, gender, and depression severity all explain variation in risk-taking, so that our analysis can better isolate the variation in risk-taking that is explained by emotional valence.

Variables that change throughout a risk-taking task can explain variation in risk-taking. For example, the size and likelihood of rewards for the choice to take a risk and for the choice to play it safe will both affect the decision to take a risk. How much people value incremental extra rewards (marginal utility) can affect the decision to take a risk [16]. Also, the outcomes of past trials can affect the decision to take a risk. In an effort to control for confounders of the relationship between emotional valence and risk-taking, we fold all these variables that potentially affect risk-taking into our models of risk-taking.

Now that we have examined many of the confounders for whether emotional valence affects risk-taking, we can focus on emotional valence itself. In order to begin answering our question on whether emotional valence affects risk-taking, we need to start investigating possible mechanisms for the effect. We draw from three hypotheses about how emotional valence affects risk-taking: the mood-maintenance hypothesis, the affect-infusion model, and the reward processing hypothesis.

1.2.1 Mood-Maintenance Hypothesis

The mood-maintenance hypothesis begins to explain how positive emotional valence decreases risk-taking. Under the mood-maintenance hypothesis, relative to people in a neutral affective state, people with more positive emotional valence are less likely to take risks because they are more motivated to maintain their feelings [17]. Conversely, according to the affect regulation hypothesis, people experiencing negative emotional valence are motivated to regulate towards more positive affect, thus taking risks to try to obtain a positive outcome [17]. The mood-maintenance hypothesis can explain how more positive emotional valence decreases risk-taking. In this thesis, we propose a way to test for a version of this hypothesis.

1.2.2 Affect-Infusion Model

The affect-infusion model posits how positive emotional valence is correlated with increased risk-taking. [18] suggests that the affect-infusion model can explain increased risk-taking. The affect-infusion model generally posits that as situations become more complicated and unexpected, emotional valence becomes more influential in how someone processing information. The affect-infusion model assumes that while people are making decisions, they have different perceptions of potential outcomes and they use different information-processing strategies between outcomes. [18] suggests that people experiencing positive affect (high emotional valence) rely more on positive cues via heuristic processing and also perceive outcomes of risky options as desirable and thus make more risky decisions. On the other hand, people experiencing negative affect (low emotional valence) may have more cautious and systematic ways of processing information that involves negatively evaluating options and which yields more conservative decisions. The affect-infusion model can be extended to explain how positive emotional valence is correlated to increased risk-taking. In this thesis, we propose a test for a version of this hypothesis.

1.2.3 Reward Processing Hypothesis

The reward processing hypothesis explains how increasing emotional valence is linked with either less or more risk-taking, depending on what decision also occurred alongside increasing emotional valence. Reward processing refers to anticipation and consumption of reward stimuli. From an evolutionary standpoint, the reward processing system helps humans learn behaviors to maximize contact with beneficial stimuli and minimize contact with harmful stimuli, increasing the likelihood of survival and reproduction [19]. Changes in emotional valence are potentially a metric of beneficial or harmful contact. Thus, changes in emotional valence can moderate the relationship between past choices and present choices. If a participant notices that the choice to gamble has previously been accompanied by an increase in emotional valence, the participant's reward processing system might activate and assign gambling as a beneficial stimuli. One behavioral outcome may be that the participant gambles more in the future. Alternatively, if a participant notices that the choice to gamble is accompanied by a decrease in emotional valence, the participant's reward processing system may work to assign gambling as a harmful stimuli. One behavioral outcome may be that the participant gambles less in the future. In this thesis, we propose a test for a version of this hypothesis.

1.2.4 Hurdles in Synthesizing Literature

There are some key limitations that prevent us from drawing unified conclusions about the relationship between emotional valence and risk-taking from available literature: 1) models lack task specificity; and 2) mathematical models are too rigid.

One barrier to synthesizing these scientific hypotheses is how the relationship found in studies can heavily depend on the specific gambling task. For example, two papers which make claims about the same relationship can use drastically different gambling tasks. [18] used a horse betting task that had 2 rounds, no feedback about rewards until the end, and explicit ratings of risk, while [20] used an acey-ducey game with 120 rounds, continual feedback after every round, and implicit ratings of risk depending on the participant's own mental calculations. Understandably, the two papers came to different conclusions. [18] reported a quadratic relationship between valence and gambling, while [20] reported a positive relationship between valence and gambling. Concretely, in this thesis, we can move towards a more consistent understanding on whether emotional valence affects risk-taking by both examining the relevant details of the task and abstracting away those details in favor of more general trends. In this thesis, we carefully examine task dynamics and integrate information into mathematical models, to begin to control for the effect of task dynamics on the relationship between emotional valence annd risk-taking.

Also, the lack of consideration of the complexity of a task is another potential hurdle to synthesizing these hypotheses on emotional valence and risk-taking. The complexity of the task changes the effect of emotional valence on risk-taking. Indeed, risk-taking can be interpreted as a competition between affective processes and deliberate cognitive-control processes [21], so to be comprehensive scientists, we need to try to control for cognitive processes in our experiments to be able to make claims about whether emotional valence affects risk-taking. In this thesis, we focus on a simple task and do not claim any generalizations at higher levels of task complexity. Future studies should examine tasks with different levels of complexity.

Also, some studies of these scientific hypotheses assume unidirectional relationships between emotional valence and risk-taking [20] [22] which can be a source of inconsistency between studies. Under this assumption, the mood-maintenance hypothesis and affect-infusion model directly contradict one another because the mood-maintenance hypothesis predicts higher emotional valence will cause less risk-taking behavior and the affect-infusion model less risk-taking.

Alternatively, in this thesis, we allow for more flexible effects. [18] already started to consider the possibility that the effect may be bidirectional. Participants were asked to place two bets, one for each horse race. In each race, there were 2 low-risk horses, 2 medium-risk horses, and 2 high-risk horses. The participants get higher payoffs for selecting more risky horses. They observed that participants with negative affect will take fewer risks, which is consistent with the affect infusion model. They also observed that participants with positive affect will take fewer risks, which is consistent with the mood-maintenance hypothesis. More broadly, [18] provides evidence that this investigation of how emotional valence influences gambling could benefit from incorporating different types of effects beyond unidirectional effects. In this thesis, we want to detect flexible effects which capture all three hypotheses about how emotional valence affects risk-taking (the mood-maintenance hypothesis, the affect infusion model, the reward processing hypothesis), so we consider unidirectional, bidirectional, and moderator effects.

1.3 Computer Science and Statistics to Synthesize Literature

In this thesis, we balance real-world utility with statistical power while analyzing computational cost.

One method to test whether emotional valence informs risk-taking is to design a simple, highly-controlled experiment that would give us the most statistical power, but with limited real-world application. We could use data from an experiment where only emotional valence and risk-taking behavior are allowed to change while all other covariates that inform emotional valence or risk-taking are held constant. Using a simple experiment would enable us to make strong claims about whether emotional valence informs risk-taking. However, this approach is not without limitations. Such an analysis is more vulnerable to unknown confounders. Moreover, our knowledge of what covariates affect emotional valence and risk-taking is not without gaps, and, if we miss any confounders which contribute to the variation of emotional valence and risk-taking, then we could misattribute the effect to the relationship between emotional valence and risk-taking. Also, our experiment may not generalize to more real world situations, where much less is held constant. Another limitation is that it may not be practically feasible to implement such an experiment.

An alternative method to test whether emotional valence informs risk-taking involves a first step of using a more complex experiment, supplemented with computer science and statistics, in order to be able to make strong claims while considering applicability to the real-world. For the first step, we can use data from an experiment where many variables are randomized and we rely on existing scientific literature, computer science, and statistics to help us begin to control for these variables. By randomizing some variables, we distribute the effect of some confounders and are less vulnerable to confounders. Additionally, our experiment is more likely to generalize to real world situations where numerous variables are uncontrolled. Beyond the findings of this first step, we can attempt to strengthen our conclusions in follow-up experiments.

In this thesis, we take the second approach to balance statistical power with real-world utility. Specifically, we examine data from an experiment with randomization. We find ways to consider many potential covariates of emotional valence and risk-taking, we consider flexible effects of covariates, and we take into account task-specific dynamics - all of which can otherwise be important barriers to determining whether emotional valence affects risk-taking.

We develop statistical machinery to tackle our research question. That is, we formalize many scientific ideas about emotional valence and risk-taking via mathematical models. We set up models for our target variables which are computationally inexpensive, informed by science, flexible, and have desirable theoretical properties. We begin to answer our research question with respect to an inferential framework of a Markov Blanket and conditional randomization testing which empowers us to detect many scientifically informed relationships between emotional valence and risk-taking in a way that is neat, statistically powerful, and which does not require as many assumptions as in other similar approaches. By taking a multidisciplinary approach that combines domain knowledge with computational and statistical tools, we hope in this thesis to create a solid scientific foundation for investigating the effect of emotional valence on risk-taking behavior.

Chapter 2

Thesis Goal

Our thesis goal is to build a scientifically informed hypothesis test to answer the question on whether emotional valence informs risk-taking. First, we will explain why we selected the set of participants and the experimental setup, then we will carefully examine participants' behavior during the experiment. Then, we will validate our data, we will select covariates, we will construct models of risk-taking which are scientifically informed and which demonstrate promise to learn three scientific hypotheses about how emotional valence affects risk-taking (mood-maintenance hypothesis, affect infusion model, reward processing hypothesis), and we will examine the stability of models. Also, we will rely on scientific literature and models to build a conditional randomization test for deciding whether emotional valence affects risk-taking. We will examine advantages and disadvantages of the conditional randomization test. Finally, we will present future directions to answer our research question with respect to computer science, statistics, and psychology.

Chapter 3

Risk-Taking Task

With access to some datasets from collaborators at the National Institutes of Health (NIH), we endeavored to select a dataset in which the experimental setup and the set of participants would be appropriate for helping us understand whether emotional-valence affects risk-taking. In this thesis, we rely on a simple and data rich task performed by a cohort of healthy controls and patients with depression. In the conclusion, we will examine next steps for going beyond the experiment chosen for this thesis.

3.1 Source of Data

This thesis is in collaboration with the Mood Brain and Development Unit and the Machine Learning Team at the National Institute of Mental Health. Among the many data sources of the two groups, we selected two data sources for this thesis, one sample of adolescents (age: 13 - 18) and one sample of college students (age: 19 - 25). All data were from experiments conducted after approval by the NIH Institutional Review Board. Data was shared while following HIPAA guidelines.

All subjects performed the same task, except the adolescents completed the task in an MRI scanner while the college students completed the task while sitting at a computer. All participants were compensated for their time and received additional compensation proportional to the number of points they earned during the gambling task. After this thesis is due, we will conduct the hypothesis test for whether emotional valence affects risk-taking by using data from a third set of participants who completed the gambling task on Amazon's Mechanical Turk.

3.2 Participants

We had a total of 38 participants across the two data sets, as summarized in Table 3.1.

Table 3.1: Participants had the following subject characteristics.

Category			Total	Train	Test
College Student	Female	Healthy	7	4	3
College Student	Female	Depressed	2	1	1
College Student	Male	Healthy	6	3	3
College Student	Male	Depressed	3	1	2
Adolescent	Female	Healthy	4	2	2
Adolescent	Female	Depressed	13	7	6
Adolescent	Male	Healthy	2	1	1
Adolescent	Male	Depressed	1	1	0
		Totals:	38	20	18

For the adolescents, we had direct access to a depression diagnosis. We also had access to results from a Snaith-Hamilton pleasure scale modified for clinician administration (SHAPS-C [23]) test, with a support of 14-56. This score is a measurement of anhedonia, which is a dampened ability to experience pleasure. A higher SHAPS-C score corresponds to more severe anhedonia and less ability to experience pleasure. Anhedonia is a core symptom of depression and is relevant to emotional valence, or the pleasantness of an emotion. The SHAPS-C score only captures the anhedonia symptom of depression so it is not a perfect indicator of depression, such that SHAPS-C score and depression diagnosis have a correlation of r=0.1991. Despite their small correlation, for this study, SHAPS-C is still informative and we use the SHAPS-C score as an approximate measure of depression severity, a covariate we later use for emotional valence and for risk-taking. We use a threshold at 24 so that any score at or above the threshold approximately corresponds to a depression diagnosis.

For the college students, we had no access to a depression diagnosis; however, we did have access to a depression severity score, measured with Beck Depression Inventory (BDI [24]) with a support of 0-21, with a threshold at 8 so that any score at or above the threshold approximately qualifies for a depression diagnosis. We used the depression severity score to impute depression diagnosis.

For both datasets, we normalized depression severity on a scale of 0 to 1, with the threshold value at 0.5 so that measurement was comparable between groups. One limitation of this data and this thesis is that we merged two different psychological test scores into one covariate, which is a useful approximation of depression severity but certainly can be improved. For example, future studies can have all participants take the same psychological test.

For the rest of our thesis, we aimed for a 50-50 split of hierarchically sampled participants (Table 3.1), so that the train set and the test set are similar in the number of participants who satisfy any combination of (college vs. adolescent) $\times$ (female vs. male) $\times$ (depressed vs. healthy). Then, the training set would be used for exploratory data analysis and the test set would be used for our eventually preregistered analysis. Also, to conduct the hypothesis test for whether emotional valence affects risk-taking, the test dataset would be supplemented by another set of participants, who are crucially not used in the exploratory data analysis.

3.3 Gambling Task

All the participants performed the same probabilistic reward task (Figure 3.1), with 90 trials of decision making over the course of approximately 12 minutes. In each decision, participants can choose to play it safe and win a certain number of points or to take a risk and earn a reward that is one of two options. If the participant chose to gamble, the participant was not explicitly told the probability of earning the higher reward, but the experiment was designed to grant the higher reward with 50% chance. In this way, the two choices (not-gamble or gamble) were designed to have similar expected reward. The reward amounts of each trial were generated independently.

To also measure emotional valence, participants were asked every 2-3 trials to rate "How happy do you feel right now?" on a sliding scale from "unhappy" to "happy" (totalling to 37 self-reports throughout 90 trials).

Figure 3.1: The gambling task presents a participant with 90 gambling decisions and 37 asks of emotional valence, interspersed every 2-3 gambling decisions. For each gambling decision, the participant is given the choice between receiving a fixed point value (left side) or gambling and receiving either of two higher and lower values (right side) with an unknown probability of receiving either reward. The participant is given 3 seconds to make a decision and if the participant does not decide, the gamble is chosen by default. Then, the participant is shown the actual reward of the decision, which concludes the gambling trial. To measure emotional valence, participants were asked to rate "How happy do you feel at this moment?" on a sliding scale centered in between unhappy and happy. From start to end, the game lasted approximately 12 minutes.

Characteristics of this gambling task make it well-suited for beginning to answer our research question with high statistical power. The task is meant to "elicit rapid changes in affective state [10]" which will be helpful in inducing variance in emotional valence, which will increase our signal-to-noise ratio in the emotional valence measurement in our data. Then, with a higher signal-to-noise ratio, we would have more statistical power to fit accurate predictive models, discover true effects, and test statistical hypotheses related to our research question.

Unlike other gambling games, where there exist provably superior strategies that would yield larger rewards on average than random decision-making, in this task, there are no objectively better or worse strategies because all choices (gamble or not-gamble) in all trials are equally rewarding on average. Moreover, on average, the strategy of choosing to gamble every trial yields equal rewards as the strategy of choosing not-to-gamble every trial, or, indeed, any mixture of gambling and not gambling.

This is important because in games where there exist objectively better or worse strategies, the participants' level of strategic thinking becomes a confounding factor to their choice to gamble or not. Moreover, risk-taking literature suggests that the complexity and cognitive demand of a task affects if and how emotional valence affects risk-taking [21], so participants who are more or less adept at playing may have their risk-taking behavior differentially affected by changes in mood due to the effect of playing skill on cognitive demand and task complexity.

On the other hand, in our task where success is purely based on luck rather than skill, we eliminate the effect of skill on risk-taking behavior. Therefore we avoid the difficulty of modeling this confounder or its downstream effects on the relationship between emotional valence and risk-taking.

One limitation of this gambling task is the possibility that participants will demonstrate behavior that doesn't reflect risk preferences. Participants might make choices based on preferences unrelated to risk. For example, a participant might make any choices in an effort to earn compensation while spending as little effort as possible. One feature of our task that might boost the incentive to make risk-driven decisions is that we compensate participants for their time and proportional to how many points a participant earns during the task, so if a participant cares about larger compensation, then a participant will try to earn more points and be less likely to make random decisions. This is despite the fact that there are objectively no such compensation-maximizing strategies–participants may be unaware of this fact and still attempt to make decisions in order to maximize reward.

Finally, the simplicity of each trial also means we can ask participants to repeat many trials without getting fatigued. The additional trials will be helpful to build statistical models with less uncertainty.

Chapter 4

Participants' Behavior During Gambling Task

We profiled the behavior of participants during the gambling task. The goal of this chapter is to examine emotional valence, risk-taking, and all our covariates of interest (all variables summarized in Figure 4.1), to set up the next parts of this thesis, the exploratory data analysis and the hypothesis test, in which we construct models to comprehensively connect these variables and begin to answer our question on whether emotional valence affects risk-taking.

This section is crucial because neglecting the complexity of a task and failing to integrate task-specific information into modeling can be a barrier for combining results from prior risk-taking literature. Thus, in this section, we play close attention to task-specific dynamics such as the different parameters of every trial, the outcomes of previous trials, the participant's individual characteristics' contribution to how the participant plays the task, the changes in emotional valence throughout the task, and the various memory effects that could be at play throughout the task. Overall, this attention to task dynamics will be fruitful for creating a model that combines existing literature about emotional valence and risk-taking.

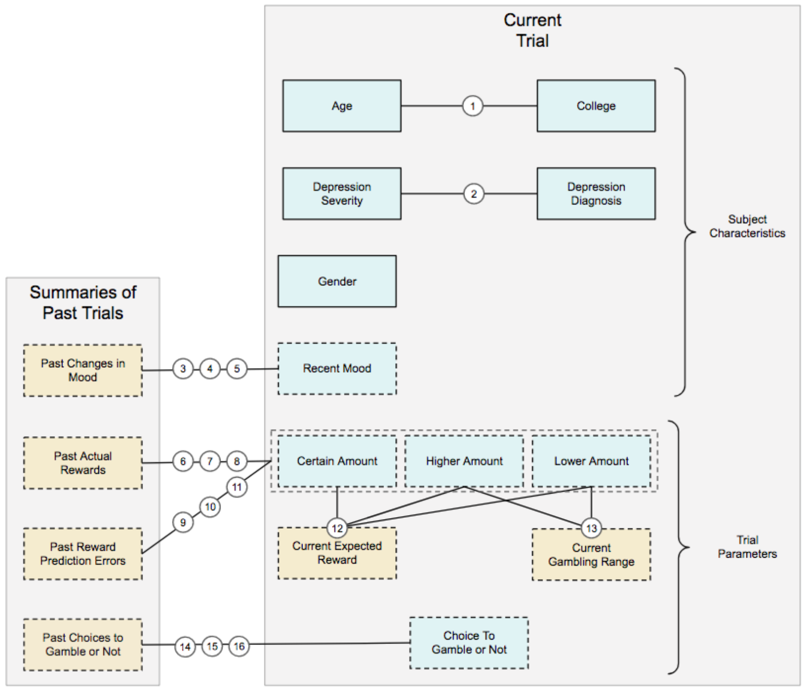

Figure 4.1: The full set of variables we considered in this thesis. Blue boxes indicate explicitly measured variables. Yellow boxed indicate implicit variables that are relevant to summarizing the task or summarizing the results of past trials, which probably influence the present trial. Variables with a solid line do not change from trial to trial. Variables with a dashed line change throughout the task. Our response variable is the choice to gamble or not ('gambling'). The black lines between variables indicate mathematical relationships between variables. The numbered mathematical equations are explicitly written on the next page.

Equations for Variables. E-Valence abbreviates Emotional Valence

$$\textcircled{1}\ \text{College} = \begin{cases} 0 & \text{if Age} < 18 \\ 1 & \text{if Age} \geq 18 \end{cases}$$

$$\textcircled{2}\ \text{Depression Diagnosis} = \begin{cases} 0 & \text{if College} = 1\text{, Depression Severity} < 0.5 \\ & \text{if College} = 0\text{, Doctor Diagnosed without Depression} \\ 1 & \text{if College} = 1\text{, Depression Severity} \geq 0.5 \\ & \text{if College} = 0\text{, Doctor Diagnosed with Depression} \end{cases}$$

$$\textcircled{3}\ \text{Past Changes in E-Valence (Primacy)}_t = \frac{\sum_{i=0}^{t-1} \gamma^i \text{Recent E-Valence}_i}{\sum_{i=0}^{t-1} \gamma^i}, \gamma = 0.5$$

$$\textcircled{4}\ \text{Past Changes in E-Valence (Recency)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Recent E-Valence}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 0.5$$

$$\textcircled{5}\ \text{Past Changes in E-Valence (No Memory Change)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Recent E-Valence}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 1.0$$

$$\textcircled{6}\ \text{Past Reward (Primacy)}_t = \frac{\sum_{i=0}^{t-1} \gamma^i \text{ Reward}_i}{\sum_{i=0}^{t-1} \gamma^i}, \gamma = 0.5$$

$$\textcircled{7}\ \text{Past Reward (Recency)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Reward}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 0.5$$

$$\textcircled{8}\ \text{Past Reward (No Memory Change)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Reward}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 1.0$$

$$\textcircled{9}, \textcircled{10}, \textcircled{11}\ \text{Reward Prediction Error}_t = \begin{cases} \text{Reward}_t - \frac{\text{Higher Amount}_t + \text{Lower Amount}_t}{2} & \text{if chose to gamble} \\ 0 & \text{otherwise} \end{cases}$$

$$\textcircled{9}\ \text{Past RPE (Primacy)}_t = \frac{\sum_{i=0}^{t-1} \gamma^i \text{RPE}_i}{\sum_{i=0}^{t-1} \gamma^i}, \gamma = 0.5$$

$$\textcircled{10}\ \text{Past RPE (Recency)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{RPE}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 0.5$$

$$\textcircled{11}\ \text{Past RPE (No Memory Change)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{RPE}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 1.0$$

$$\textcircled{12}\ \text{Current Expected Reward}_t = \frac{\text{Certain Amount}_t + \text{Higher Amount}_t + \text{Lower Amount}_t}{3}$$

$$\textcircled{13}\ \text{Gambling Range}_t = \text{Higher Amount}_t - \text{Lower Amount}_t$$

$$\textcircled{14}, \textcircled{15}, \textcircled{16}\ \text{Gamble}_t = \begin{cases} 1 & \text{if chose to gamble} \\ 0 & \text{if chose not-to-gamble} \end{cases}$$

$$\textcircled{14}\ \text{Past Gamble (Primacy)}_t = \sum_{i=0}^{t-1} \gamma^i \text{Sign(Gamble)}_i, \gamma = 0.5$$

$$\textcircled{15}\ \text{Past Gamble (Recency)}_t = \sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Sign(Gamble)}_i, \gamma = 0.5$$

$$\textcircled{16}\ \text{Past Gamble (No Memory Change)}_t = \sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Sign(Gamble)}_i, \gamma = 1.0$$

4.1 Target Variables

We are interested in whether emotional valence affects risk-taking. For both variables, we examine how they change within participants (throughout the task) and between participants. Participants demonstrate high variation in emotional valence and in risk-taking, which is a promising sign that our data has a nonzero signal-to-noise ratio and thus that we have nonzero statistical power. We also plot the two variables against each other to examine the unconditional relationship.

4.1.1 Emotional Valence

Emotional valence was measured 37 times for each person throughout the 90 trial gambling task, based on self-reports to the question "How happy are you at this moment" on a sliding scale from "unhappy" to "happy," scaled onto 0 to 1 respectively. To leverage the data throughout the task, we wanted to impute emotional valence even when it wasn't explicitly asked for. Since every participant had a self-report by the 4th trial, we imputed emotional valence with the most recent self-report then we threw out the first 3 trials for the rest of our data analysis.

Participants' large shifts in emotional valence increases the chance that this experiment has a nonzero signal-to-noise ratio and the gambling task indeed induced changes in emotional valence. Within each subject, emotional valence changed dramatically throughout the task (Figure 4.2), so that participants spanned an average of 34% of the emotional valence scale. Still, one participant had a steady emotional valence that didn't span more than 8% of the emotional valence scale while the largest span was 88% of the emotional valence scale.

Across participants, the average emotional valence was neutral valence, 0.54. One participant was close to happy on average, with the highest average of 0.89. Another participant was close to unhappy on average, with the lowest average of 0.26 (Figure 4.3). The variation is an indication that our task and sample are valid because in a random sample of healthy and depressed participants playing a probabilistic game, we would expect some participants to be close to unhappy on average and some participants to be close to happy on average.

4.1.2 Risk-Taking

Risk-taking was measured 90 times and coded with a 1 to indicate someone took a risk to gamble and a 0 to indicate someone played it safe and choose not to gamble. To accommodate imputed emotional valence, we focused on the last 87 trials of the task. Participants' choice to gamble or not changed often over time (Figure 4.4).

The global mean of gambling across all 87 trials of all 20 participants in the training set was 0.55, which indicates that participants gambled slightly more than not. This could be an artefact of the task automatically choosing to gamble if the participant does not make a decision within 3 seconds.

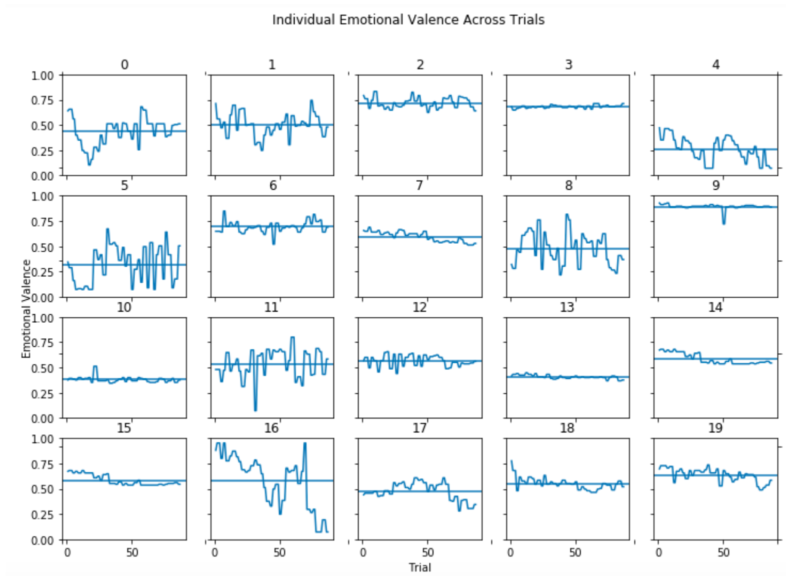

Figure 4.2: For all 20 participants in the train set, the plots show the 87 emotional valence ratings (line plot), and the average emotional valence (horizontal line) throughout the gambling task. For most subjects, emotional valence had large changes throughout the task.

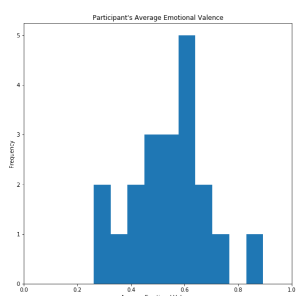

Figure 4.3: For each subject, we averaged their self-reported and imputed emotional valence scores across the last 87 trials.

Having a large variation in gambling makes it plausible that our experiment has a large signal-to-noise ratio. There were no subjects who gambled exclusively, nor were there subjects who never gambled. (Figure 4.5). Across subjects, the smallest gambling rate was 26% and the highest 70%.

4.1.3 Risk-Taking Based on Emotional Valence

To get an initial picture of the association between risk-taking rate based on self-reported emotional valence (without controlling for covariates), we examined

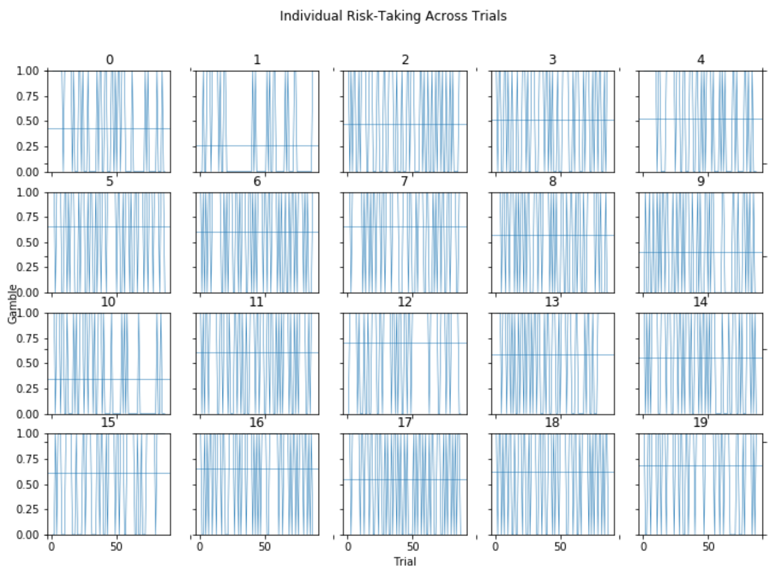

Figure 4.4: People often changed between gambling and not gambling (blue dots). The average rate of gambling (horizontal line) also varied between participants.

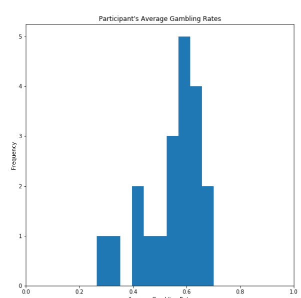

Figure 4.5: Participant's gambling rate varied from 26 % to 70 %, with an average gambling rate of 55 %.

small ranges of emotional valence and all the risk-taking decisions made at any trial by any participant with an emotional valence in that range. We noticed that for each range, there were roughly as many decisions to gamble as decisions to not gamble (Figure 4.6). This indicates no obvious trends about the relationship between emotional valence and risk-taking.

Regardless of the lack of obvious unconditional association between our target variables, we are still motivated to examine other covariates of risk-taking which could potentially reveal a conditional relationship between emotional va-

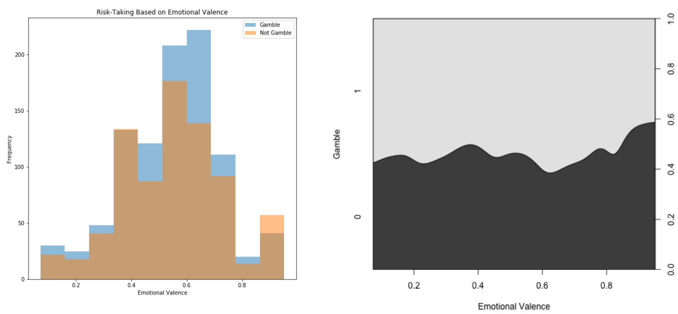

(a) Histogram **(b)** Conditional Density Plot

Figure 4.6: Risk-Taking based on self-reported emotional valence for all 87 trials for all subjects showed no obvious trends.

lence and risk-taking. We divide these covariates into two groups: explicit and implicit covariates of risk-taking as elucidated by scientific literature. We analyze the dynamics of these two groups of variables separately, and delve into constructing a bigger picture of connecting the covariates and the target variables in the exploratory data analysis and the hypothesis test.

4.2 Other Explicit Covariates of Risk-Taking

We know many variables affect risk-taking. In this section, we examine some of the scientifically relevant covariates that we explicitly measured (shown in blue in Figure 4.1). Understanding these covariates will be crucial to isolating variation in risk-taking due to emotional valence instead of any of these other influential covariates.

Some covariates of risk-taking stayed constant throughout the task (age, gender, and depression diagnosis, Figure 4.7). Age ranged from 14 to 28 with an average age of 18.9. Our sample had 7 males and 13 females. Six of the participants had depression severity ≥ 0.5 which is indicative of depression. This is a subset of the 10 participants we considered to have depression, that is 4 participants had depression severity scores below the threshold but were separately diagnosed with depression.

Other covariates changed throughout the task (trial parameters: the certain amount R_1 for the not-gamble option, higher amount R_2 for the gamble option, and lower amount R_3 for the gamble option, Figure 4.8). The rewards for each trial ranged from $-3, 3$. Importantly, from the design of the gambling task and as seen in Figure 4.8, the decision to gamble has approximately the same expected value as the decision not-to-gamble, helping scientists observe risk-taking behavior instead of strategic behavior. That is, the trial parameters were designed according to the following relationship, $R_1 \simeq \frac{R_2+R_3}{2}$. This means that

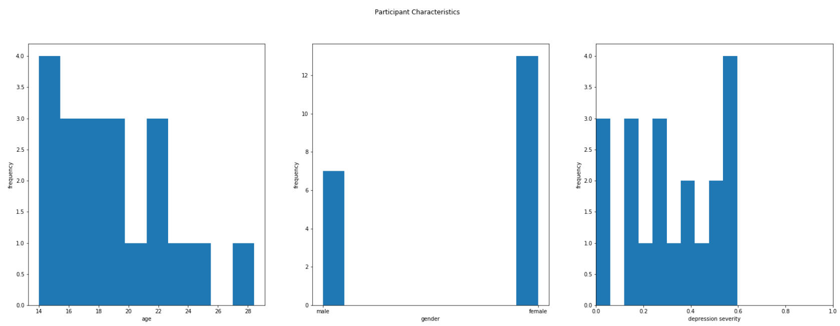

Figure 4.7: Some explicitly measured variables stayed constant throughout the task according to the following distributions across the $n = 20$ participants in the training group. Depression severity has a threshold so severity ≥ 0.5 (blue vertical line) is considered a depression diagnosis.

we can replace our three variables that summarize the trial, the certain amount R_1, higher amount R_2, and lower amount R_3 with two other implicit variables the current expected reward X_1 and the gambling range X_2 which capture the same information and thus equally inform risk-taking. This partially motivates our next section, about implicit variables that contribute to risk-taking.

Figure 4.8: Some explicitly measured variables changed throughout the task. The trial parameters for the 87 trials of all 20 participants are shown. The certain amount corresponds to the reward for the not-gamble option, higher amount corresponds to the higher reward for the gamble option, and lower amount corresponds to the lower reward for the gamble option.

4.3 Other Implied Covariates of Risk-Taking

While our explicitly measured variables offer a start to accounting for variation in risk-taking, there is still more to examine. In this section, we explore the dynamics of implicit variables that inform risk-taking. We examine two implicit task parameters that completely summarize our three explicit task parameters. We also examine four implicit variables that summarize the past and are derived from explicit parameters that describe the outcomes of trials. Understanding the dynamics of these implicit covariates (shown in yellow in Figure 4.1) is important because these variables will be crucial in isolating the relationship between emotional valence and risk-taking.

4.3.1 Current Expected Reward, Gambling Range

The explicit task parameters of the current trial, that is the certain amount R_1 for the not-gamble option, higher amount R_2 for the gamble option, and lower amount R_3 for the gamble option can be summarized by two implicit task parameters, the current expected reward X_1 and the gambling range X_2 (defined below).

The current expected reward X_1 is a measure of the baseline approximate expected reward of the current trial whether the participant chooses to gamble or not $X_1 = \frac{R_1+R_2+R_3}{3}$. The gambling range X_2 is a measure of the risk involved in choosing to gamble, with a wider range $X_2 = R_2 - R_3$ representing a higher risk of only earning the lower amount R_3 instead of the higher amount R_1 or certain amountR_2. Overall, then,

$$R_1 \simeq X_1;$$
$$R_2 \simeq X_1 + \frac{X_2}{2};$$
$$R_3 \simeq X_1 - \frac{X_2}{2}.$$

These two features X_1, X_2 capture the information from the initial features R_1,R_2,R_3 while reducing the redundancy, and are pairwise independent (by design), so that instead of using the explicit task parameters R_1,R_2,R_3, we can use the implicit task parameters X_1, X_2 as covariates for risk-taking.

4.3.2 Summaries of the Past

We know that participants likely remember outcomes of past trials and use that information to make a decision in the present trial. For example, the reward processing hypothesis suggests that participants might pay attention to past changes in emotional valence and past choices to gamble or not to determine which choice (to gamble or not) is beneficial or harmful. Additionally, the past rewards and past reward prediction errors may give some indication of which choice to make. It is important for us to consider all four of these variables in our model for variation in risk-taking.

However, we didn't explicitly measure these variables, so we need to come up with some method to mathematically summarize them. We take inspiration from scientific literature about memory effects and from past literature that has tried to represent these implicit variables with exponential sums [10].

When people summarize a sequence of events, there are three common possibilities of serial position effects [25]: primacy, recency, and equal/ none. The primacy serial position effect explains that people remember the beginning of a sequence more than the rest. The recency serial position effect explains that people remember the end of the sequence, the most recently observed more than the rest. The equal serial position effect explains that people remember the contents of a sequence equally well, or that the order of the sequence does not change the probability of remembering. We will consider all of these serial position effects in this thesis.

We know there are many ways to summarize the past in this gambling task. As a slight improvement on previous literature, we will rely on exponential averages instead of exponential sums as used in previous literature that relied on the same gambling task [10]. We use exponential averages instead of sums to reduce the variance of these terms which might otherwise compromise its usefulness as a predictor of risk-taking or emotional valence. The primacy effect on a variable corresponds to a decaying weight of more recent terms at $\gamma = 0.5$. The recency effect on a variable corresponds to an increasing weight of more recent terms at $\gamma = 0.5$. The equal effect corresponds to no change in the weight of terms at $\gamma = 1$. For example, we can derive three versions of a summary for past changes in emotional valence, corresponding to the three scientific representations of memory:

- Primacy Effect:

$$\text{Past Changes in Emotional Valence (Primacy)}_t = \frac{\sum_{i=0}^{t-1} \gamma^i \text{Recent Emotional Valence}_i}{\sum_{i=0}^{t-1} \gamma^i}, \gamma = 0.5$$

- Recency Effect:

$$\text{Past Changes in Emotional Valence (Recency)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Recent Emotional Valence}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 0.5$$

- Equal Effect, No Memory Change:

$$\text{Past Changes in Emotional Valence (No Memory Change)}_t = \frac{\sum_{i=0}^{t-1} \gamma^{t-i-1} \text{Recent Emotional Valence}_i}{\sum_{i=0}^{t-1} \gamma^{t-i-1}}, \gamma = 1.0$$

We explicitly define all our 4 variables of interest with respect to the 3 serial position effects in the equations of the Figure 4.1.

In order to test all possible combinations of the serial position effects of these 4 summaries, it would require us to build 3^4 models. However, to save computation time while still maintaining comprehensiveness, in the exploratory data analysis, we narrow down each of the 4 variables into a specific serial position effect, then we use only this permutation of variables in our hypothesis test for whether emotional valence affects risk-taking.

Chapter 5

Theory and Exploratory Data Analysis

In this chapter, we validate our data, we select covariates, we construct models of risk-taking which are scientifically informed and which demonstrate promise to learn three scientific hypotheses about how emotional valence affects risk-taking (mood-maintenance hypothesis, affect infusion model, reward processing hypothesis), and we test the stability of our models across regularization and resampling.

5.1 Is the data invalid?

In this section, we examine two ways our data could be invalid: low signal-to-noise ratio and inconsistency with scientific literature about emotional valence and risk-taking. We discover that the signal-to-noise ratio in this data collected throughout the gambling task is high enough to predict risk-taking better than by chance and we show that we do not have evidence that our data is inconsistent with scientifically established null hypotheses about risk-taking and emotional valence.

First, we want to know if the signal-to-noise ratio is high enough in this data to predict risk-taking behavior better than chance. If we cannot predict better than chance, we would not proceed with the analysis and we would look for a different data set. The baseline model we are trying to beat is the global mean model which predicts the probability of gambling of every trial of every participant to be the global mean gambling rate, 0.55. We built 7 increasingly complex logistic regression models to predict the probability participant i would gamble at trial t. Models used covariates related to participant characteristics, the current trial, and the outcomes of past trials (See Appendix B for full methods).

Five of the models (models 4,5,6,7,8) of risk-taking were more predictive than chance. Relative to the global mean model, these five models had higher accuracy, higher AUC, and lower cross-entropy loss in leave-one-subject-out cross validation (Table 5.1). These models had significantly lower rate of incorrect classification in leave-one-subject-out cross validation (McNemar Test, $p < 0.01$, Table 5.2). These models had significantly higher likelihood of the data (likelihood ratio test, $p < 0.01$, Table 5.3). Overall, using this data, we can predict risk-taking better than chance.

Table 5.1: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss. Bolds indicated most desirable value among the models.

Model #	Model	ACC	AUC	Loss
1	Global Mean	0.500	0.374	4.0E-04
2	Participant	0.479	0.501	4.0E-04
3	Current Trial	0.668	0.702	3.6E-04
4	Past Trials	0.591	0.622	3.8E-04
5	Current Trial + Past Trials	**0.684**	**0.739**	**3.4E-04**
6	Participant + Current Trial	0.664	0.698	3.6E-04
7	Participant + Past Trials	0.579	0.622	3.8E-04
8	Participant + Current Trial + Past Trials	0.669	0.734	3.5E-04

Table 5.2: McNemar Test To Beat The Baseline. We performed McNemar tests to see if models significantly out predicted risk-taking relative to the global mean model.

Model	Model Name	statistic	p-value
2	Participant	0.0	1.0
3	Current Trial	267.0	5.8e-15
4	Past Trials	178.0	4.9e-07
5	Current Trial + Past Trials	231.0	2.5e-20
6	Participant + Current Trial	266.0	2.4e-14
7	Participant + Past Trials	192.0	5.9e-05
8	Participant + Current Trial + Past Trials	242.0	1.1e-16

Second, our data did not indicate we should reject scientifically informed null hypotheses about how participant characteristics inform average emotional valence (depression) and risk-taking rates (age, gender, depression diagnosis). Our full methodology is explained in Appendix C.

Table 5.3: Likelihood Ratio Tests To Beat The Baseline. We performed likelihood ratio tests for nested models. Model numbers correspond to Table B.2. Dev. stands for deviance and p stands for p-value.

Model	Model Name	DF	Dev.	p
2	Participant	4	1588.5	2.2e-16
3	Current Trial	2	0.624	0.73
4	Past Trials	4	4285.9	2.2e-16
5	Current Trial + Past Trials	6	4288.2	2.2e-16
6	Participant + Current Trial	6	4288.2	2.2e-16
7	Participant + Past Trials	8	4813.8	2.2e-16
8	Participant + Current Trial + Past Trials	10	4816.5	2.2e-16

According to [12], people with depression have lower average emotional valence than healthy controls. Using a t-test in our data set, we successfully fail to reject this null hypothesis. The null hypothesis is that the difference in average emotional valence of healthy participants and depressed participants > 0. In our data, participants with depression were closer to unhappy on average (average $= 0.46$; Figure 5.1) than healthy controls (average $= 0.63$; Figure 5.1) and we successfully failed to reject the null hypothesis that people with depression have lower average emotional valence than healthy controls (95% CI of difference $[-\infty, 0.267]$, estimated difference $= 0.175$, t $= 3.28$, p-value $= 0.99$).

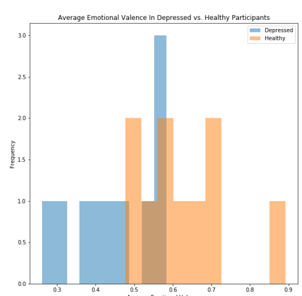

Figure 5.1: The average emotional valence was centered at 0.54 with range $0.26, 0.89$. Participants with depression had a lower average emotional valence than healthy participants $(0.46 < 0.63)$

According to [12], older people gamble at lower rates. Using a Fisher's exact test in our data, we successfully failed to reject this null hypothesis. The null hypothesis corresponds to the odds ratio of gambling in college students to gambling in adolescents ≤ 1. In our data, older participants gambled less than younger patients and we successfully failed to reject the null hypothesis that older participants gambled less than younger participants (college students' gambling rate $= 0.51$; adolescents' gambling rate $= 0.59$; 95% CI of odds ratio $[0.61, \infty]$, odds ratio estimate $= 0.72$, p-value $=0.9997$; Figure 5.2a).

According to [11], men gamble more than women. Using a Fisher's exact test, we successfully failed to reject this null hypothesis. The null hypothesis corresponded to the odds ratio of gambling in men to gambling in women≥ 1. In our data, men gambled at a similar rate as women, and we successfully could not reject the null hypothesis that men gamble more than women (male gambling rate = 0.547; female gambling rate = 0.553; 95% CI of odds ratio [0.82, ∞], odds ratio estimate = 0.97, p-value = 0.625; Figure 5.2b).

According to [13] [14] [15], we are unsure of the relationship between depression and risk-taking. Using a Fisher's exact test, we successfully fail to reject this null hypothesis. The null hypothesis corresponds to the odds ratio of gambling in participants with depression to gambling in healthy participants = 1. In our data, patients with depression gambled at a similar rate as healthy participants and we successfully could not reject the null hypothesis that participants with depression and healthy participants gambled at the same rate (depression gambling rate = 0.552; healthy gambling rate = 0.551; 95% CI of odds ratio [0.83, 1.22], odds ratio estimate = 1.01, p-value $\simeq 1$; Figure 5.2c).

Low signal-to-noise ratio and inconsistency with scientific literature are not the only ways our data could be invalid. Future studies could examine other ways our data could be invalid.

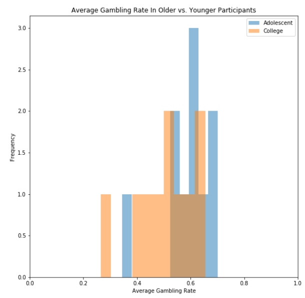

(a) By Age

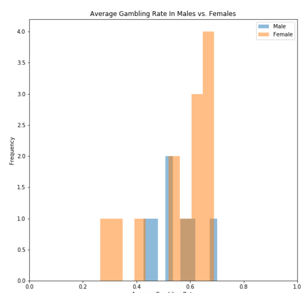

(b) By Gender

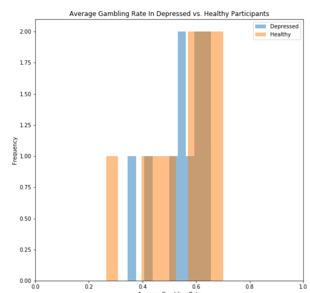

(c) By Depression

Figure 5.2: Participant's gambling rate varied from 26 % to 70 %, with an average gambling rate of 55 %. Adolescent participants gambled more than college students (college students' gambling rate = 0.51, adolescents' gambling rate = 0.59). Females and males gambled at a similar rate (female gambling rate = 0.547, male gambling rate = 0.553). Participants with depression and healthy participants gambled at a similar rate (depression gambling rate = 0.552, healthy gambling rate = 0.551).

5.2 What kinds of models should we use to model risk-taking?

In this section we propose using logistic regression and neural networks to model risk-taking, to mimic human decision making behavior, to accurately capture possible scientific relationships, and to perform statistical inference.

Generative models mimic human decision making. Given a choice whether or not to gamble, one might weigh the advantages and disadvantages of gambling, and if the advantages outweigh the disadvantages, one might proceed to gamble. Similarly, we can emulate this latent variable of the advantages minus the disadvantages of gambling as Z, explained by observed covariates η (e.g. the participant's characteristics, the trial parameters, the outcomes of past trials) and unobserved covariates ϵ. Ultimately, we can say if the advantages out weigh the disadvantages, i.e. $Z > 0$, the person decides to gamble $Y = 1$.

Moreover, we assume that a person i is making a decision Y_{it} at trial t such that

$$Y_{it} = \begin{cases} 1 & Z_{it} \geq 0 \\ 0 & Z_{it} < 0 \end{cases}$$

and the latent variable Z_{it} is a combination of observed variables η_{it} and unobserved variables ϵ_{it}. We define η_{it} to be the function of observed variables that affect Z_{it} and ϵ_{it} to be a function of unobserved variables that affect Z_{it} so that $Z_{it} = \eta_{it} + \epsilon_{it}$

We assume the unobserved variables follow a logistic distribution

$$\epsilon_{it} \sim \text{Logistic}(\mu = 0, s = 1)$$

$$\text{E}(Z_{it}) = \text{E}(\eta_{it}) + \text{E}(\epsilon_{it}) = \eta_{it}$$

$$\begin{aligned} \text{P}(Y_{it} = 1) &= \text{P}(Z_{it} \geq 0) \\ &= \text{P}(\eta_{it} + \epsilon_{it} \geq 0) \\ &= \text{P}(\epsilon_{it} \geq -\eta it) \\ &= \text{logit}(\eta_{it}) \end{aligned}$$

In this generative scheme, we have created η_{it} as a mirror for the probability of risk-taking , so that when η_{it} increases, the probability of gambling $\text{P}(Y_{it} = 1)$ increases and when η_{it} decreases, the probability of gambling $\text{P}(Y_{it} = 1)$ decreases too. In this way, if we model how the covariates affect η_{it}, we know the same direction of the relationship is mirrored for how the covariates affect the probability of gambling $\text{P}(Y_{it} = 1)$. For the rest of this thesis, we proceed to model the log odds of gambling, η_{it}.

For all the generative models, we have two goals,

- To capture accurate relationships: we want to account for all possible relationships between all covariates $X_1, ..., X_p$ and η_{it}
- To perform inference: we want to examine the models for our scientific hypotheses that mood X_j has a first or second order relationship with η_{it}

We settle on using a logistic regression and a neural network as our models, both of which have different strengths and weaknesses towards achieving these two goals.

The first goal is important because we want to isolate the effect of mood X_j on Y_{it} while controlling for all other covariates $X_1, ..., X_p$. From a scientific perspective, a lack of control or consideration of other covariates could be an explanation for the lack of consensus over how emotional valence affects decisions making Y_{it}. Notably, the neural network is best suited for getting closer to this goal because the neural network can theoretically learn any relationships among covariates X_j and η_{it}. The logistic regression is less well-suited for this goal because it can only learn linear relationships between $X_1, ..., X_p$ and η_{it}. Still, if we want the logistic regression to learn any nonlinear relationships between $X_1, ..., X_p$ and η_{it}, we can hard-code those nonlinearities and treat them as additional covariates, so that the logistic regression can also learn nonlinearities.

We also want to perform inference for the second goal, to detect whether valence X_j affects η_{it} as proposed by three hypotheses (the mood-maintenance hypothesis, the affect infusion model, the reward processing hypothesis). Notably, the logistic regression has an upper hand over the neural network when we estimate variance of the effects of covariates. That is, we have less uncertainty with the logistic regression than the neural network. This is for two reasons. The neural network is sensitive to randomness in the training set, so adding or dropping a couple points can drastically change the neural network [26] more than it would change the logistic regression. Also the neural network has an internal randomness such that the backpropagation solving algorithm critically depends on the random initialization of starting weights [26]; the logistic regression has none of this randomness.

We can rely on generative models of risk-taking, specifically logistic regression models and neural network models. For both these models, we can perform inference on covariates' effects on the response and capture flexible and accurate relationships between the covariate and response. In the next three chapters, we examine if our models are stable and if our models can learn flexible relationships.

5.3 Are risk-taking models stable?

We examined if the directions of estimated effects of covariates on risk-taking stayed the same across regularization levels and across data resamplings. The logistic regression prediction function demonstrates stability across regularization levels and some stability across data resamplings. The neural network prediction function was not stable across regularization levels and was stable across data resamplings at some levels of regularization. This exploratory experiment (full methods in Appendix D, E) is not statistically rigorous.

5.3.1 Logistic Regression

The logistic regression was stable across regularization levels and had some stability across data resamplings (for full methods, see Appendix D).

The estimated directions of the effects of covariates on risk-taking are the same even over a large grid of L1 regularization values. For each covariate, the estimates of the effect of the covariate on risk-taking $\hat{\beta}_j^\lambda$ in L1 regularized logistic regression prediction functions were either all ≥ 0 or all ≤ 0 (Figure 5.3). This indicates that our models' estimated directions of effects were not sensitive to this range of regularization strengths.

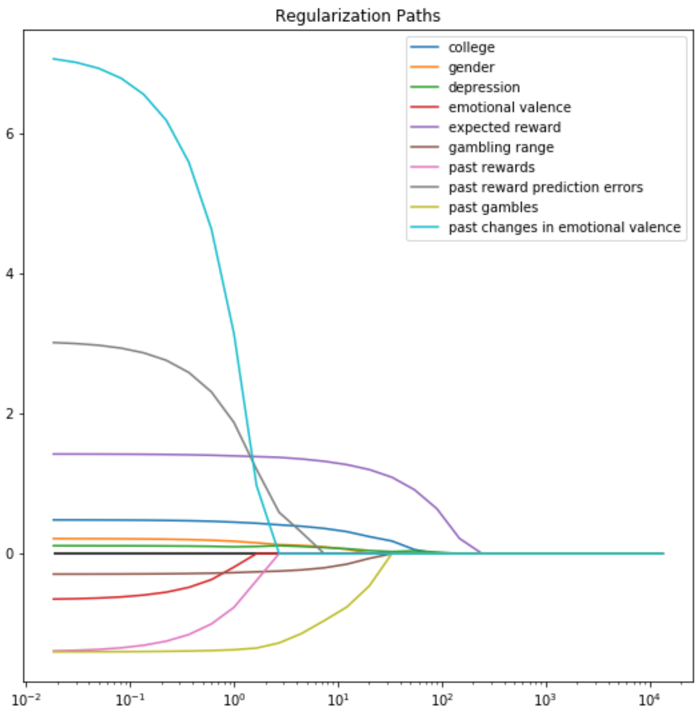

Figure 5.3: The regularization path of all covariates. The x-axis is a log scale of regularization strengths, λ. The y-axis is the estimated beta weight $\hat{\beta}_j^\lambda$ for each covariate X_j at regularization strength λ. The black line indicates $\hat{\beta}_j^\lambda = 0$.

For some of the covariates, the estimated direction of the covariate's effect on risk-taking stayed the same across 100 resamplings of $\frac{n}{2} = 10$ participants. For each covariate and regularization level, let stability+ be the percentage of resamplings on which the estimated effect of a covariate on the response was positive and let stability- be the percentage of resamplings on which the estimated effect of a covariate on the response was negative. For some of the covariates, such as expected reward, depression, past gambles, and college, either stability+ or stability- was at 100% until both were at 0% for large enough regularization λ (Figure 5.4), which indicates that for these covariates, the estimated direction of the effect of the covariates on risk-taking was the same across resamplings of the data and stayed the same (or the beta weight was estimated to be 0) across the grid of regularization strengths.

On the other hand, for some covariates, the estimated direction of the covariate's effect on risk-taking did change across resamplings. For some covariates, the stability+ and stability- were lower than 100% and greater than 0% at many levels of regularization, which indicates that the estimated direction of the effect of the covariates on risk-taking varied across resamplings of the data (Figure 5.4).

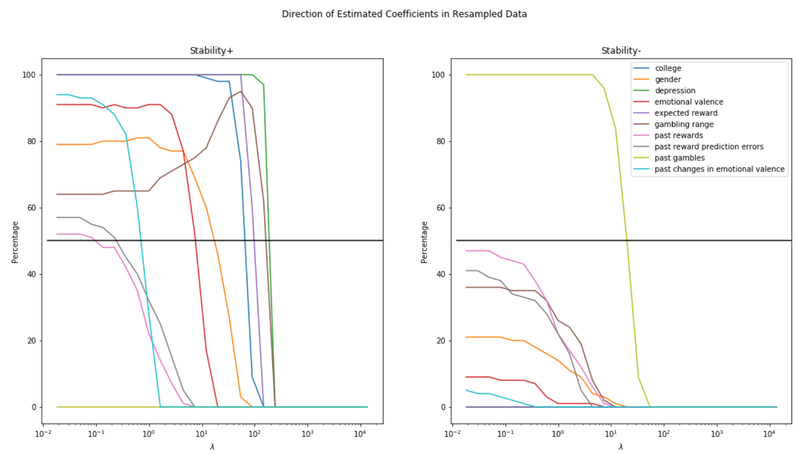

Figure 5.4: The stability+ and stability- path of all covariates. The x-axis is a log scale of regularization strengths, λ. The y-axis is stability+ and stability- for each covariate X_j at regularization strength λ. The black line indicates 50%.

5.3.2 Neural Network

The neural network prediction function was not stable across regularization levels and was stable across data resamplings at some levels of regularization. We can examine how the estimated log odds of risk-taking $\hat{\eta}_{it}$ changes with respect to each covariate X_j, $\frac{\partial \hat{\eta}_{it}}{\partial X_j}$, evaluated at each observation and averaged (i.e. "the empirical average gradient"), as an estimate of an effect a covariate on risk-taking (subject to smoothing, see Appendix E for more details).

The neural network prediction function is unstable across regularization levels (for the full explanation of how we fit neural networks, see Appendix E). Looking across regularization levels λ, ϵ corresponding to l1-regularization and adversarial noise, the sign of the empirical average gradient changes at different levels of regularization.

The neural network prediction function was stable across data resamplings at some levels of regularization. At some regularization levels, the sign of the empirical average gradient remained the same in 4 out of the 5 bootstrap resamplings of the $n = 20$ participants. This is exploratory evidence that the neural network was stable across data resamplings at some levels of regularization.

Our exploratory data analysis indicate that our logistic regression model is sometimes stable and our neural network is unstable. These results suggest we need to be careful about generalizing our models to other data sets.

5.4 How will models of risk-taking learn scientifically hypothesized relationships?

In our generative models, we can implicitly or explicitly encode unidirectional, bidirectional, or moderator effects between emotional valence and risk-taking with $X_j, X_j^2, X_j \cdot X_i$ respectively, which capture the mood-maintenance hypothesis [17], the affect-infusion model [18], and the reward processing hypothesis [19].

One hurdle in incorporating scientific literature about emotional valence and risk-taking into a mathematical model is how some studies mathematically test for only single direction relationships between emotional valence and risk-taking [20] [22], which can yield contradictions among hypotheses.

Instead, in this thesis, we consider three different scientifically meaningful mathematical representations of covariates $(X_j, X_j^2, X_j X_i)$ which allow us to test for these three different hypotheses about the relationship between emotional valence and risk-taking.

5.4.1 Unidirectional Effect X_j

Given our assumptions, a finding that $\text{sign}(\text{E}(\frac{\partial \eta}{\partial X_j})) = \{+, -\}$ for a given covariate X_j is evidence in favor of a unidirectional effect between a covariate X_j and probability risk-taking. Let $\text{sign}(\frac{\partial \eta}{\partial X_j})$ represent how, on average, a change in the covariate X_j affects the probability of risk-taking $\text{P}(Y_{it} = 1) = \text{logit}(\eta_{it})$. Moreover, if $\text{sign}(\text{E}(\frac{\partial \eta}{\partial X_j}))$ is positive, all other covariates held constant, then a participant or trial with larger covariate X_j has a larger probability of risk-taking Y. If $\text{sign}(\text{E}(\frac{\partial \eta}{\partial X_j}))$ is negative, all other covariates held constant, then a participant or trial with larger covariate X_j has a smaller probability of risk-taking Y. Then, a nonzero $\frac{\partial \eta}{\partial X_j}$ indicates a unidirectional effect between the covariate X_j and log odds of risk-taking $\text{P}(Y_{it} = 1)$.

For X_j as the covariate of emotional valence, if our mathematical models of risk-taking which have nonzero $\text{E}(\frac{\partial \eta}{\partial X_j})$, then our mathematical models have detected if larger emotional valence either increases risk-taking or decreases risk-taking . This corresponds to the affect-infusion model [18] or the mood-maintenance hypothesis [17] respectively.

5.4.2 Bidirectional Effect X_j^2

A finding that $\text{sign}(\text{E}(\frac{\partial^2 \eta}{\partial X_j^2})) = \{+, -\}$ for a given covariate X_j is evidence in favor of a bidirectional effect between a covariate X_j and probability risk-taking. Let $\text{sign}(\text{E}(\frac{\partial^2 \eta}{\partial X_j^2}))$ represent how, on average, a change in the covariate

X_j affects the relationship between the covariate X_j and the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$. Moreover, if $\mathrm{sign}(\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j^2}))$ is positive, all other covariates held constant, then, if X_j is larger, then the effect of X_j on the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$ is more positive, so while at smaller X_j, larger X_j will yield a decrease in the probability of risk-taking $\mathrm{P}(Y_{it} = 1)$, one the other hand at a larger X_j, larger X_j will yield a increase in the probability of risk-taking $\mathrm{P}(Y_{it} = 1)$. If $\mathrm{sign}(\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j^2}))$ is negative, all other covariates held constant, then, for larger X_j, then the effect of X_j on the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$ is more more negative, so while at smaller X_j, an increase in X_j will yield a increase in the probability of risk-taking $\mathrm{P}(Y_{it} = 1)$, one the other hand at a larger X_j, an increase in X_j will yield a decrease in the probability of risk-taking $\mathrm{P}(Y_{it} = 1)$. Thus, a nonzero $\frac{\partial^2 \eta}{\partial X_j^2}$ indicates a bidirectional effect between the covariate X_j and probability of risk-taking $\mathrm{P}(Y_{it} = 1)$.

For the covariate X_j of emotional valence, if our mathematical models of risk-taking which have nonzero $\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j^2})$, then this is an indication that the relationship between emotional valence and risk-taking has two directions. This corresponds to a synthesis of the affect-infusion model [18] and the mood-maintenance hypothesis [17], such that higher emotional valence can both increase and decrease risk-taking.

5.4.3 Moderator Effect $X_j X_i$

A finding that $\mathrm{sign}(\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i})) = \{+, -\}$ for a given covariates X_j, X_i is evidence in favor of a moderator effect between a covariate X_j and probability risk-taking. Let $\mathrm{sign}(\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ represent how, on average, a change in the covariate X_j affects the relationship between X_i and the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$. Moreover, if $\mathrm{sign}(\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ is positive, all other covariates held constant, then for larger X_j, the relationship between X_i and the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$ is more positive. If $\mathrm{sign}(\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ is negative, all other covariates held constant, then for larger X_j, the relationship between X_i and the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$ is more negative. Then, a nonzero $\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i})$ indicates a moderator effect between the covariate X_j and the relationship between the covariate X_i and probability of risk-taking $\mathrm{P}(Y_{it} = 1)$.

For X_j as past changes in emotional valence and X_i as past decisions to gamble, if $\mathrm{sign}(\mathrm{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ is positive then our model has detected the reward processing hypothesis [19]. Moreover,

- $X_j > 0, X_i > 0$: if past changes in emotional valence X_j is positive and past decisions to gamble X_i is positive, then this corresponds to the scenario where the participant has gambled and seen increases in emotional valence, so the choice to gamble is a beneficial decision and under the reward processing hypothesis, the participant's gambling probability for the next trial should be higher and indeed our model mirrors that because $X_j \cdot X_i$ is positive, $\text{sign}(\text{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ is positive, and this term adds a positive contribution to the probability of risk-taking Y

- $X_j > 0, X_i < 0$: if past changes in emotional valence X_j is positive and past decisions to gamble X_i is negative, then this corresponds to the scenario where the participant has chosen not-to-gamble and seen increases in emotional valence, so the choice not-to-gamble is a beneficial decision and under the reward processing hypothesis, the participant's gambling probability for the next trial should be lower and indeed our model mirrors that because $X_j \cdot X_i$ is negative, $\text{sign}(\text{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ is positive, and this term adds a negative contribution to the probability of risk-taking Y

- $X_j < 0, X_i > 0$: if past changes in emotional valence X_j is negative and past decisions to gamble X_i is positive, then this corresponds to the scenario where the participant has gambled and seen decreases in emotional valence, so the choice to gamble is a harmful decision and under the reward processing hypothesis, the participant's gambling probability for the next trial should be lower and indeed our model mirrors that because $X_j \cdot X_i$ is negative, $\text{sign}(\text{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ is positive, and this term adds a negative contribution to the probability of risk-taking Y

- $X_j < 0, X_i < 0$: if past changes in emotional valence X_j is negative and past decisions to gamble X_i is negative, then this corresponds to the scenario where the participant has chosen not-to-gamble and seen decrease in emotional valence, so the choice not-to-gamble is a harmful decision and under the reward processing hypothesis, the participant's gambling probability for the next trial should be higher and indeed our model mirrors that because $X_j \cdot X_i$ is positive, $\text{sign}(\text{E}(\frac{\partial^2 \eta}{\partial X_j \partial X_i}))$ is positive, and this term adds a positive contribution to the probability of risk-taking Y

5.4.4 Explicit and Implicit Effects in Generative Models

The generative models of the logistic regression and neural network have prediction functions that can incorporate all scientifically informed effects. The logistic regression is less well-suited for this goal because first glance, a prediction function can only contain linear relationships between $X_1, ..., X_p$ and η_{it}. Still, if we want the logistic regression prediction function to contain any nonlinear relationships between $X_1, ..., X_p$ and η_{it}, we can hard-code those nonlinearities and treat them as additional covariates, so that the logistic regression prediction function can also include nonlinearities. That is, we can hard-code both first order covariates X_j and second order covariates X_j^2, $X_i \cdot X_j$ so that the logistic regression prediction function includes the unidirectional effects, bidirectional effects, and moderator effects. Notably, the logistic regression would assign a single beta weight to $X_i \cdot X_j$, and for two continuous covariates, this is a very rigid construction of a moderator effect. Future studies should examine less rigid constructions of moderator effects.

An infinitely large neural network can theoretically include any relationship among covariates and the probability of gambling. Our finite network models have more limited function space, but can potentially learn the above one direction, bidirectional, and moderator effects in addition to other less scientifically interpretable effects.

While in this thesis, we limit ourselves to examining first and second order representations of covariates because they already begin to test for all three the scientific hypotheses of emotional valence and risk-taking (the mood-maintenance hypothesis [17], the affect-infusion model [18], and the reward processing hypothesis [19]); however, future studies might examine more complex relationships.

Overall, to capture scientifically informed unidirectional, bidirectional, or moderator effects between emotional valence and risk-taking, we will use generative models that implicitly or explicitly find relationships between $X_j, X_j^2, X_j \cdot X_i$ for all covariates X_j, X_i and risk-taking Y. This will allow us to detect three hypotheses about how emotional valence affects risk-taking: the mood-maintenance hypothesis [17], the affect-infusion model [18], and the reward processing hypothesis [19].

5.5 Do risk-taking models learn flexible relationships?

We have some exploratory but not statistically rigorous evidence that risk-taking models learn flexible relationships, and thus can learn relationships aligned with the mood-maintenance hypothesis [17], the affect-infusion model [18], and the reward processing hypothesis [19]. Additionally, considering other more flexible relationships between emotional valence and risk-taking does not compromise predictive ability (the full methods of these experiments are in Appendix F).

5.5.1 Flexible Logistic Regression

We have some exploratory, not statistically rigorous evidence that the flexible logistic regression found flexible unidirectional, bidirectional, and moderator effects of covariates on risk-taking. We chose l1-ratio, $\alpha = 0.65, 0.001$ which maximized predictive performance.

Unidirectional Effects X_j

We have exploratory evidence that the flexible logistic regression detected unidirectional effects of covariates on risk-taking. For a logistic regression model of probability of gambling, one indication that a covariate X_j had a unidirectional effect on the probability of risk-taking $P(Y_{it} = 1) = \text{logit}(\eta_{it})$ is if the gradient entry of X_j: $\frac{\partial \eta}{\partial X_j}$ is nonzero. We can estimate this quantity with the estimated beta-weight for X_j: $\hat{\beta}_j$. For the flexible logistic regression model, many values had nonzero $\hat{\beta}_j$ with the largest $\hat{\beta}_j = 1.8$ (Figure F.1), indicating that the flexible logistic regression model might have picked up on unidirectional effects. This is completely exploratory, because we didn't set up a hypothesis test for $\hat{\beta}_j \neq 0$.

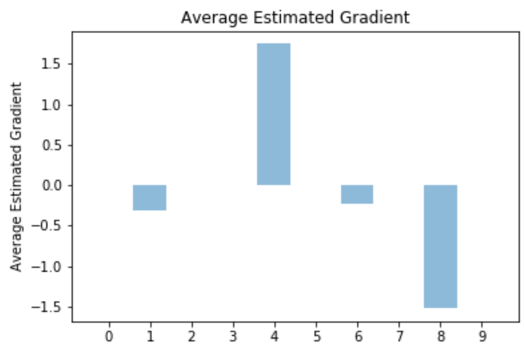

Figure 5.5: The estimated gradient of η with respect to covariates. The numbered covariates corresponded to 0 = 'college', 1 = 'gender', 2 = 'depression', 3= 'emotional valence', 4 = 'expected reward', 5= 'gambling range', 6= 'past rewards', 7= 'past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'

Bidirectional Effects X_j^2

We have exploratory evidence that the flexible logistic regression detected bidirectional effects of covariates on risk-taking. One indication that a covariate X_j had a bidirectional effect on the probability of risk-taking $P(Y_{it} = 1) = \text{logit}(\eta_{it})$ is if the Hessian matrix diagonal entry of X_j: $\frac{\partial^2 \eta}{\partial X_j^2}$ is nonzero. For the logistic regression, we can estimate this quantity with the estimated beta-weight for X_j^2: $\hat{\beta}_{jj}$. For the flexible logistic regression model, many values had nonzero $\hat{\beta}_{jj}$ with the largest $\hat{\beta}_{jj} = 0.40$ (Figure F.2), indicating that the flexible logistic regression model might have picked up on bidirectional effects. This is completely exploratory, because we didn't set up a hypothesis test for $\hat{\beta}_{jj} \neq 0$.

Moderator Effects $X_j \cdot X_i$

We have exploratory evidence that the flexible logistic regression detected moderator effects of covariates on risk-taking. One indication that a covariate X_j had a moderator effect on the probability of risk-taking $P(Y_{it} = 1) = \text{logit}(\eta_{it})$ is if the Hessian matrix off-diagonal entry of X_{ji}: $\frac{\partial^2 \eta}{\partial X_j \partial X_i}$ is nonzero. For the logistic regression, we can estimate this quantity with the estimated beta-weight for $X_j \cdot X_i$: $\hat{\beta}_{ji}$. For the flexible logistic regression model, many values had nonzero $\hat{\beta}_{ji}$ with the largest $\hat{\beta}_{ji} = 2.12$ (Figure F.2), indicating that

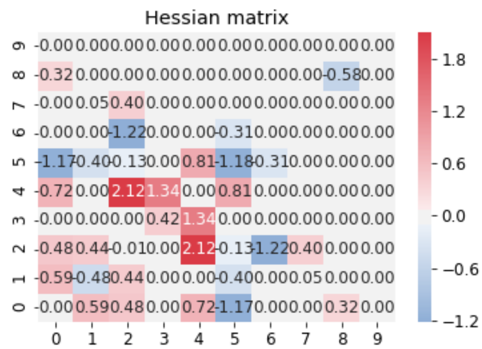

Figure 5.6: The estimated Hessian matrix of η with respect to covariates. The numbered covariates corresponded to 0 = 'college', 1 = 'gender', 2 = 'depression', 3= 'emotional valence', 4 = 'expected reward', 5= 'gambling range', 6= 'past rewards', 7= 'past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'

the flexible logistic regression model might have picked up on moderator effects. This is completely exploratory, because we didn't set up a hypothesis test for $\hat{\beta}_{ji} \neq 0$.

Overall, the estimated gradient vector and estimated Hessian matrix corresponded to estimated weights $\hat{\beta}$ assigned to X_i and X_i^2 or $X_i \cdot X_j$ respectively and since many of the entries of the estimated gradient and estimated Hessian matrix values were nonzero, we have some evidence that the logistic regression did learn unidirectional, bidirectional and moderator effects between covariates and risk-taking.

5.5.2 Neural Network

We settled on the hyperparameters $\lambda = 0.01$ for l1 regularization and $\epsilon = 0.0001$ for adversarial noise (for full methods on how we fit neural networks see Appendix E). We have some exploratory evidence that the neural network detected some nonzero flexible on-direction, bidirectional, and moderator effects.

Unidirectional Effects X_j

We found exploratory evidence that the neural networks detected unidirectional effects. While it is hard to say what it would mean for a neural network to detect a unidirectional effect, we can examine the expected gradient of η with respect to a covariate as an indication that the neural network detected a unidirectional effect on risk-taking. For any covariate, if the empirical average gradient of $\hat{\eta}$ (subject to a smoothing procedure [27]) is large, this is some exploratory evidence that the neural network has learned a unidirectional effect. For the neural network model we trained, some average estimated gradients were nonzero (Figure 5.7a), with the largest average gradient at 1.25 indicating that our neural network may have learned some unidirectional effects of the covariates on risk-taking.

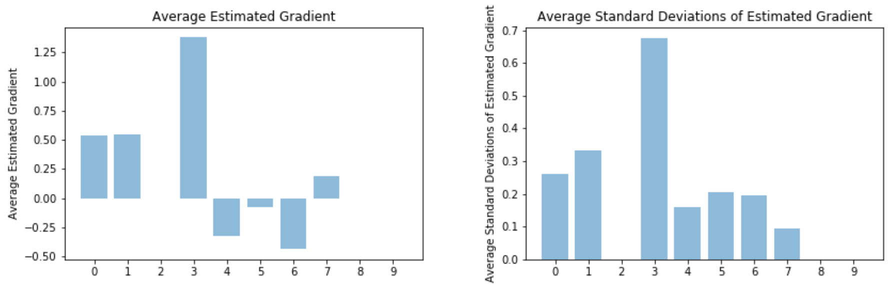

(a) Average Gradient **(b)** Standard Deviation of Gradient

Figure 5.7: The average and standard deviation of the estimated gradient of η with respect to a covariate evaluated at all data points corresponding to all participants and all trials. The numbered indices correspond to 0 = 'gender', 1 = 'depression severity', 2= 'emotional valence', 3 = 'expected reward', 4 = 'gambling range', 5 = 'age', 6 ='past rewards', 7 ='past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'.

Bidirectional Effects X_j^2

We have exploratory evidence that the flexible logistic regression detected bidirectional effects of covariates on risk-taking.

It is hard to say what a bidirectional effect looks like with respect to the Hessian matrices for all data points. In theory, each entry of the expected Hessian matrix, $E(\frac{\partial^2 \eta}{\partial X_j^2})$ is an indication for how a function of risk-taking η changes with respect to one quadratic covariate X_j^2. In our data, we can evaluate a Hessian matrix for each data point, subject to a smoothing procedure [27]. We posit that if the empirical average of the estimated Hessian matrices evaluated at all data points is nonzero on diagonal entries, then we have some evidence that the neural network is learning some bidirectional effects of X_j on Y. Indeed, we found some nonzero diagonal entries on the averaged Hessian, with magnitude up to 0.02 (Figure 5.8a). Also, we posit that if the majority of entries are all positive or all negative, then we have some evidence that the neural network is learning some bidirectional effects of X_j on Y. To quantify this concept, we define prevalence of j as the difference in the ratio of jj^{th} estimated Hessian matrix entries which are positive and the ratio of jj^{th} estimated Hessian matrix entries which are negative, so that larger magnitude of prevalence indicates that many of the Hessian matrix entries were of the same sign. In our neural network, some of the prevalence values of some of the covariates were nonzero (Figure 5.8b), with magnitude up to 0.13 which provides some evidence that the neural networks are learning bidirectional relationships between the covariates and the probability of gambling.

Moderator Effects $X_j \cdot X_i$

We have exploratory evidence that the flexible logistic regression detected moderator effects of covariates on risk-taking. It is hard to say what a moderator effect looks like with respect to the ij^{th} entries of the expected Hessian matrices for all data points. In theory, the expected Hessian entries $E(\frac{\partial^2 \eta}{\partial X_i \partial X_j})$ are a proxy for how a function of risk-taking η changes with respect to two covariates X_j, X_i. In our data, we can evaluate a Hessian matrix for each data point. We posit that if the empirical average of estimated the Hessian matrices over all data points is nonzero on off-diagonal entries, then we have some evidence that the neural network is learning some moderator effects of X_j on the relationship between X_i and Y. Indeed, we found some nonzero diagonal entries on the averaged Hessian, with magnitude up to 0.04 (Figure 5.8a). Also, we posit that if the majority of entries are all positive or all negative, then this is some evidence that the neural network is learning some moderator effects of $X_j \cdot X_i$ on Y. To quantify this concept, we examine prevalence of ij as the difference in the ratio of ij^{th} estimated Hessian matrix entries which are positive and the ratio of ij^{th} estimated Hessian matrix entries which are negative, so that larger magnitude of prevalence indicates that many of the Hessian matrix entries were of the same sign. In our neural network, some of the prevalence values were nonzero (Figure 5.8b), with magnitude up to 0.26, which provides some evidence that the neural networks are learning moderator relationships between the covariates and the probability of gambling.

5.5.3 Prediction

The models which can learn nonlinearities did no worse at prediction than the model that cannot (for full methods see Appendix F).

Table 5.4: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss. The neural network was trained with $\lambda = 1e-3, \epsilon = 1e-2$

Model #	Model	ACC	AUC	Loss
1	Logistic Regression	0.66	0.712	3.57E-04
2	Flexible Logistic Regression (All Effects)	0.670	0.708	3.62E-04
3	Neural Network (All Effects)	0.683	0.683	3.36E-04

It is promising that our generative models which are designed to detect flexible and scientifically informed relationships might actually detect nonzero flexible relationships between emotional valence and risk-taking without compromising predictive ability.

(a) Average Hessian

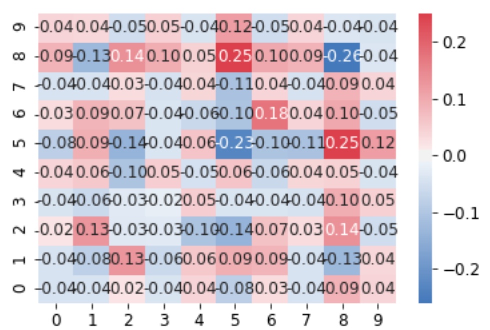

(b) Prevalence

Figure 5.8: The prevalence values, positive ratio, and negative ratio of each entry of the Hessian matrix of η with respect to two covariates. The numbered indices correspond to 0 = 'gender', 1 = 'depression severity', 2= 'emotional valence', 3 = 'expected reward', 4 = 'gambling range', 5 = 'age', 6 ='past rewards', 7 ='past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'.

5.6 Which Covariates Should We Include In Models?

In this section, we show how models which rely on combinations of covariates predict better than smaller models which rely on only subsets. We also reduce the computational cost of our analysis by choosing functional forms of covariates related to the past or rewards.

5.6.1 Models with All Covariates

As shown in Appendix B, a logistic regression model of risk-taking using combinations of covariates related to the participant, the current trial, and the outcomes of past trials, is significantly more predictive than a trivial model with no covariates. Additionally, this full model was significantly more predictive (Table 5.5) and significantly more informative (Table 5.6) than smaller models which only used a subset of the covariates, such as a model trained on only the covariates related to the participant and the current trial. In next chapter, we will use models of risk-taking which use all the covariates related to the participant, the current trial, and the outcomes of past trials.

Table 5.5: McNemar Test Against the Full Model. We performed McNemar tests to see if the full model using covariates from participants, current trial, and past trials significantly out predicted other models.

Model	Model Name	statistic	p-value
1	Global Mean	242.0	1.1e-16
2	Participant	223.0	9.6e-25
3	Current Trial	134.0	0.76
4	Past Trials	213.0	2.3e-06
5	Current Trial + Past Trials	38.0	0.017
6	Participant + Current Trial	126.0	0.50
7	Participant + Past Trials	200.0	1.4e-08

Table 5.6: Likelihood Ratio Tests Against the Full Model. We performed likelihood ratio tests for nested models. Model numbers correspond to Table B.2. Dev. stands for deviance and p stands for p-value.

Model	Model Name	DF	Dev.	p
1	Global Mean	10	4816.5	2.2e-16
2	Participant	6	3227.9	2.2e-16
3	Current Trial	8	4815.8	0.73
4	Past Trials	6	530.53	2.2e-16
5	Current Trial + Past Trials	4	528.26	2.2e-16
6	Participant + Current Trial	4	3227.8	2.2e-16
7	Participant + Past Trials	2	2.6	0.26

5.6.2 Reducing Computational Cost By Choosing Covariates Related to Past Trials and Related to Reward

We can reduce the computational complexity of our analysis by choosing a specific version of covariates related to the past and related to reward. This is exploratory and not statistically rigorous.
Past Trials

We have three versions (primacy, recency, no effect/ equal) for each of four covariates which summarize the past (past changes in emotional valence, past gambles, past rewards, and past reward prediction errors), which yields $3^4 = 81$ combinations of covariates. Instead of building models for all combinations of covariates, in this experiment, we build 12 logistic models to pick just one combination of covariates (for full methods, see Appendix G). An estimate of predictive ability of a logistic regression model trained on each version of the covariates is shown in Table 5.7.

Table 5.7: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss.

Model #	Model	ACC	AUC	Loss
1	Past Changes In Emotional Valence Primacy	0.500	0.391	3.97E-04
2	Past Changes In Emotional Valence Recency	0.500	0.402	3.97E-04
3	Past Changes In Emotional Valence Equal/ No Effect	0.500	0.418	3.97E-04
4	Past Gambles Primacy	0.500	0.391	3.99E-04
5	Past Gambles Recency	0.500	0.476	3.98E-04
6	Past Gambles Equal/ No Effect	0.545	0.567	3.92E-04
7	Past Rewards Primacy	0.518	0.541	3.94E-04
8	Past Rewards Recency	0.493	0.465	3.98E-04
9	Past Rewards Equal/ No Effect	0.501	0.426	3.97E-04
10	Past Reward Prediction Errors Primacy	0.499	0.447	3.97E-04
11	Past Reward Prediction Errors Recency	0.500	0.384	3.98E-04
12	Past Reward Prediction Errors Equal/ No Effect	0.497	0.427	3.97E-04

Based on this exploratory analysis, out of the models trained on different versions of each covariate, the best estimated predictive ability was for models trained with:

- Past Changes In Emotional Valence Primacy,
- Past Gambles Equal/ No Effect,
- Past Rewards Primacy,
- Past Reward Prediction Errors Recency.

Then, in our hypothesis test for whether emotional valence affects risk-taking, we rely on only this combination instead of modeling all 81 possible combinations. We chose past changes in emotional valence (primacy), past gambles (equal/ no effect), past rewards (primacy), and past reward prediction errors (recency) which yielded the best predictive performance (by virtue of numerical comparison instead of statistical testing). The cost of this reduction in computational cost is a reduction in statistical power, because our exploratory data analysis may not have indicated the best functional form of these covariates.

Rewards

We narrow down to using utility-transformed reward covariates after considering two possible reward covariates: utility function transformed reward covariates and raw reward covariates. In modern economic theory, people have utility functions which quantify customers' preferences over a set of choices [16]. Some utility functions are nonlinear, such that linear increases in rewards do not yield linear increases in utility, in which case looking at utility is not the same as looking at raw rewards. In this thesis, in our gambling task, it is possible people are making decisions based on utility instead of raw rewards. For full methods, see Appendix H, an estimate of the predictive ability of each version of the covariates is shown in Table 5.8.

Table 5.8: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss.

Model #	Model	ACC	AUC	Loss
1	Utility-Transformed Rewards	0.677	0.701	3.59E-04
2	Raw Rewards	0.673	0.705	3.58E-04

Based on this exploratory analysis, the best predictive ability was for the model which used utility-transformed rewards. Then, in our hypothesis test for whether emotional valence affects risk-taking, we can cut the number of models we train in half by using the utility-transformed reward covariates and not the raw reward covariates. The cost of this reduction in computational cost is a reduction in statistical power, because our exploratory data analysis may not have indicated the best functional form of these covariates.

Overall, to reduce computational cost at the expense of some statistical power in our hypothesis test for whether emotional valence affects risk-taking, we build models which use utility-transformed reward covariates and use the following specific versions of covariates which summarize the past: past changes in emotional valence (primacy), past gambles (equal/ no effect), past rewards (primacy), and past reward prediction errors (recency).

Chapter 6

Conditional Randomization Test

In this chapter, we rely on scientific literature and models informed by the previous chapter to build a conditional randomization test for deciding whether emotional valence affects risk-taking. Unlike propensity scores that also use a conditional distribution to perform inference on the conditional relationship between the response and a covariate conditional on all other covariates, the conditional randomization test does not rely on discreteness of covariates or experimental control of covariates [28], and thus does not have as many assumptions.

6.1 Scientifically Informed Joint and Conditional Distributions

We rely on scientific literature and our previous chapter to inform our models of the relationships among covariates of risk-taking, $X = X_1, ..., X_p$ and the conditional distribution of risk-taking given covariates $Y|X_1, ..., X_p$.

For our covariates of risk-taking $X_1, ..., X_p$, we will use all covariates related to the participant, the current trial, and the outcomes of past trials. Specifically, we will use utility-transformed reward covariates and we will use a specific combination of summaries of past trials (as specified in the previous chapter).

We will use neural networks and flexible logistic regressions to model risk-taking Y from the covariates $X_1, ..., X_p$. As tested in the exploratory data analysis, generative logistic regression and generative neural network models are more predictive of risk-taking than chance. As shown in the exploratory data analysis, we have some exploratory, not statistically rigorous evidence that both models do pick up on these flexible relationships between emotional valence and risk-taking and thus can detect three hypotheses about how emotional valence affects risk-taking: the mood-maintenance hypothesis [17], the affect-infusion model [18], and the reward processing hypothesis [19] (for the full model specifications, see Appendix I).

We will model emotional valence X_e and past changes in emotional valence X_{de} as linear regressions on all other covariates of risk-taking, $X_{-\{e,de\}}$. In [10], the mean of the conditional distribution of emotional valence was informed by exponentially weighted summaries of the past reward prediction errors, the past rewards from not gambling, and the past rewards from gambling, so we included all these covariates in our model. Specifically, in our models, we will have a covariate as summary of the past reward prediction errors. We will also have a covariate that is the combined summary of all past rewards, not separated into past rewards from gambling or not gambling (for the full model specifications, see Appendix I).

6.2 Hypotheses

We are interested in whether emotional valence affects risk-taking, which can be framed by the statistical concept of a Markov blanket. A Markov blanket of risk-taking, Y given its potential covariates $X_1, ..., X_p$ is the smallest set of covariates S such that if we condition on S, then Y is independent of all other covariates not in S. That is, for $X_i \notin S$:

$$Y \perp\!\!\!\perp X_i | S$$

Intuitively, a Markov blanket of Y given $X_1, ..., X_p$ is the smallest set of covariates which "fully" explains Y so that as long as we have access to those covariates, the other covariates are no longer informative and we can't explain any additional variation in Y. Of course, there are many unknown covariates that explain variation in Y, so it is impossible to fully explain Y, but out of a near comprehensive set of possible covariates, a Markov blanket is as close as we can get to "fully" explaining Y.

Our null hypothesis is that emotional valence does not affect risk-taking once we control for other covariates. The statistical claim of the null hypothesis is that (X_e, X_{de}) are not part of a Markov blanket of Y with respect to covariates $X_e, X_{de}, X_{-\{e,de\}}$. Our alternative hypothesis is that emotional valence does affect risk-taking even after we control for other covariates. The statistical claim of the alternative hypothesis is that (X_e, X_{de}) are not part of a Markov blanket of Y with respect to covariates $X_e, X_{de}, X_{-\{e,de\}}$.

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp (X_e, X_{de}))|X_{-\{e,de\}}$
  - (X_e, X_{de}) is not in a Markov blanket of Y with respect to $X_e, X_{de}, X_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\perp\!\!\!\perp (X_e, X_{de}))|X_{-\{e,de\}}$

– (X_e, X_{de}) is in a Markov blanket of Y with respect to $X_e, X_{de}, X_{-\{e,de\}}$

This Markov Blanket approach is not only broad enough to allow us to consider many covariates of risk-taking but also defined in a way that's narrow enough to make a claim of whether or not emotional valence affects risk-taking. A lack of consideration of many possible covariates could be one explanation for lack of consensus in the scientific literature on whether emotional valence affects risk-taking.

6.3 Test Statistic

This test statistic, T_9 is designed to detect unidirectional or bidirectional effects between emotional valence and risk-taking, along with moderator effects of emotional valence on the relationship between past gambles and risk-taking via the mood-maintenance hypothesis, the affect infusion model, or the reward processing hypothesis.

This test statistic T_9 is a composition of two test statistics: 1) T_5 (Appendix L) which detects for unidirectional or bidirectional effect between emotional valence and risk-taking and can correspond to detecting the mood-maintenance hypothesis and/ or the affect infusion model and 2) T_8 (Appendix M) which detects for moderator effects of emotional valence on the relationship between past gambles and risk-taking and can correspond to detecting the reward processing hypothesis:

$$
\begin{aligned}
T_9 &= t_9(D) \\
&= 1 - \min(p(T_5), p(T_8))
\end{aligned}
$$

for $p(T_5)$ and $p(T_8)$ corresponding to the p-values of the two test statistics (for full details of how we derive these p-values, see Appendix L and M respectively).

With this definition, we know that if any of T_5, T_8 are large, then the p-value of at least one is small and T_9 will be large too. Alternatively, if all of T_5, T_8 are small, then the p-values of both are large and T_9 will be small too. So this test statistic T_9 detects for whether emotional valence has a unidirectional or bidirectional effect on risk-taking or past changes in emotional valence has a moderator effect on the relationship between past gambles risk-taking under the stated assumptions about data generating processes.

Both test statistics T_5, T_8 are compositions of other test statistics which make all possible assumptions of joint and conditional distributions considered in this thesis, starting with whether the data generating processes of $Y|X_1, ..., X_p$ as either a neural network or a logistic regression, given that the data generating process of $X_e|X_{-\{e,de\}}$ as a linear regression, and the data generating process of $X_{de}|X_{-\{e,de\}}$ as a linear regression.

For test statistic T_9 and all test statistics underlying T_9, we listed the scientific hypotheses it detected, the data generating processes it assumed on risk-taking and emotional valence, the definition of the detector, the parameter of interest, the type of data randomization (A,B,C,D), and how to calculate the p-value in Table 6.1.

6.4 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp \{X_e, X_{de}\})|X_{-\{e,de\}}$, we can derive the empirical distribution of T_9 by computing our detector value over 100 test statistics T_5^*, T_8^*, evaluated over randomized data sets with conditional randomizations of $X_e|X_{-\{e,de\}}$ and $X_{de}|X_{-\{e,de\}}$ (see Appendix L and M for full details). Each T_9^* will require calculating the p-value of T_5^*, T_8^*. Calculating the p-value of T_5^* will require $10,100$ data sets of type A and $10,100$ data sets of type B while calculating the p-value of T_8^* will require $10,100$ data sets of type C and $10,100$ data sets of type D (see Appendix J for definitions of randomizations).

The resulting 100 detector values $\{T_9(i)^*\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_9 = t_9(D)$.

We rely from the following lemma, with statement and proof reproduced from Lemma 4.1 [28], to prove conditional distributions for this test statistic (and all test statistics used to construct this test statistic).

Lemma 1. *Let $D = (Z_1, Z_2, Y)$ be a triple of random variables and we can construct another triple that replaces Z_1 with a conditional randomization Z_1^*: $D^* = (Z_1^*, Z_2, Y)$ such that*

$$Z_1^*|(Z_2, Y) \overset{\mathrm{d}}{=} Z_1|Z_2$$

Then if we assume the null hypothesis $Y \perp\!\!\!\perp Z_1|Z_2$, then any test statistic $T = t(D) = t(Z_1, Z_2, Y)$ obeys

$$T|(Z_2, Y) \overset{\mathrm{d}}{=} T^*|(Z_2, Y)$$

where $T^* = t(D^*) = t(Z_1^*, Z_2, Y)$

Proof. We show that $Z_1^*|(Z_2, Y)$ has the same distribution as $Z_1|(Z_2, Y)$ under the null hypothesis, which enables us to prove that $T = t(D)$ has the same distribution as $T^* = t(D^*)$.

$$\begin{aligned} Z_1^*|(Z_2, Y) &\overset{\mathrm{d}}{=} Z_1|Z_2 \text{ by the definition of } Z_1^* \\ &\overset{\mathrm{d}}{=} Z_1|(Y, Z_2) \text{ because under the null } Z_1 \perp\!\!\!\perp Y|Z_2 \\ &\overset{\mathrm{d}}{=} Z_1|(Z_2, Y) \\ Z_1^*, Z_2, Y &\overset{\mathrm{d}}{=} Z_1, Z_2, Y \text{ by integrating over } Z_2, Y \end{aligned}$$

So we have shown that if we assume the null hypothesis $Z_1 \perp\!\!\!\perp Y|Z_2$, then

$$Z_1^*, Z_2, Y \overset{\mathrm{d}}{=} Z_1, Z_2, Y$$

,

Table 6.1: All the detectors. MMH stands for mood-maintenance hypothesis, AIM stands for affect infusion model, RPH stands for reward processing hypothesis. The data randomization type was determined by assumptions about the DGP of emotional valence. DGP stands for data generating process.

Detector	Scientific Hypotheses	Markov Blanket	DGP of Risk-Taking	DGP of Emotional Valence	Definition	Parameter of Interest	Data Randomization Type	P-value
1	MMH, AIM	X_e	$\eta_{(it)} = \sum_k \beta_k X_k$ $+\beta_{kk} X_k^2$ $+\beta_{kj} X_k \cdot X_j$ for $j \neq k$ and $X_k, X_j \in \{X_e, X^{\text{LR}}_{-\{e,de\}}\}$	$X_e \| X^{\text{LR}}_{-\{e,de\}}$ $\sim N(\mu, \epsilon)$ $\mu = \sum_i \beta_i X_i$ $for X_i \in X^{\text{LR}}_{-\{e,de\}}$	$\|\hat{\beta}_e\|$	$\|\beta_e\|$	A	$p(T_1) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_1^*(i) \geq T_1)}{100}$
2	MMH, AIM	X_e	$\eta_{(it)} = f(X_e, X^{\text{NN}}_{-\{e,de\}})$ for neural network f	$X_e \| X^{\text{NN}}_{-\{e,de\}}$ $\sim N(\mu, \epsilon)$ $\mu = \sum_i \beta_i X_i$ $for X_i \in X^{\text{NN}}_{-\{e,de\}}$	$\left\| \frac{1}{\|D\|} \sum_{it} \frac{\partial \hat{f}}{\partial X_e}(D_{it}) \right\|$	$\left\| \text{E}\left(\frac{\partial f}{\partial X_e} \right) \right\|$	B	$p(T_2) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_2^*(i) \geq T_2)}{100}$
3	MMH, AIM	X_e	$\eta_{(it)} = \sum_k \beta_k X_k$ $+\beta_{kk} X_k^2$ $+\beta_{kj} X_k \cdot X_j$ for $j \neq k$ and $X_k, X_j \in \{X_e, X^{\text{LR}}_{-\{e,de\}}\}$	$X_e \| X^{\text{LR}}_{-\{e,de\}}$ $\sim N(\mu, \epsilon)$ $\mu = \sum_i \beta_i X_i$ for $X_i \in X^{\text{LR}}_{-\{e,de\}}$	$\|\hat{\beta}_{ee}\|$	$\|\beta_{ee}\|$	A	$p(T_3) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_3^*(i) \geq T_3)}{100}$
4	MMH, AIM	X_e	$\eta_{(it)} = f(X_e, X^{\text{NN}}_{-\{e,de\}})$ for neural network f	$X_e \| X^{\text{NN}}_{-\{e,de\}}$ $\sim N(\mu, \epsilon)$ $\mu = \sum_i \beta_i X_i$ for $X_i \in X^{\text{NN}}_{-\{e,de\}}$	$\left\| \frac{1}{\|D\|} \sum_{it} \frac{\partial^2 \hat{f}}{\partial X_e^2}(D_{it}) \right\|$	$\left\| \text{E}\left(\frac{\partial^2 f}{\partial X_e^2} \right) \right\|$	B	$p(T_4) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_4^*(i) \geq T_4)}{100}$
5	MMH, AIM	X_e	-	-	$1 - \min(p(T_1), p(T_2), p(T_3), p(T_4))$	-	-	$p(T_5) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_5^*(i) \geq T_5)}{100}$
6	RPH	X_{de}	$\eta_{(it)} = \sum_k \beta_k X_k$ $+\beta_{kk} X_k^2$ $+\beta_{kj} X_k \cdot X_j$ for $j \neq k$ and $X_k, X_j \in \{X_{de}, X^{\text{LR}}_{-\{e,de\}}\}$	$X_{de} \| X^{\text{LR}}_{-\{e,de\}}$ $\sim N(\mu, \epsilon)$ $\mu = \sum_i \beta_i X_i$ for $X_i \in X^{\text{LR}}_{-\{e,de\}}$	$\|\hat{\beta}_{de,dg}\|$	$\|\beta_{de,dg}\|$	C	$p(T_6) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_6^*(i) \geq T_6)}{100}$
7	RPH	X_{de}	$\eta_{(it)} = f(X_{de}, X^{\text{NN}}_{-\{e,de\}})$ for neural network f	$X_{de} \| X^{\text{NN}}_{-\{e,de\}}$ $\sim N(\mu, \epsilon)$ $\mu = \sum_i \beta_i X_i$ for $X_i \in X^{\text{NN}}_{-\{e,de\}}$	$\left\| \frac{1}{\|D\|} \sum_{it} \frac{\partial^2 \hat{f}}{\partial X_{de} \partial X_{dg}}(D_{it}) \right\|$	$\left\| \text{E}\left(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}} \right) \right\|$	D	$p(T_7) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_7^*(i) \geq T_7)}{100}$
8	RPH	X_{de}	-	-	$1 - \min(p(T_6), p(T_7))$	-	-	$p(T_8) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_8^*(i) \geq T_8)}{100}$
9	MMH, AIM, RPH	X_e, X_{de}	-	-	$1 - \min(p(T_5), p(T_8))$	-	-	$p(T_9) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_9^*(i) \geq T_9)}{100}$

in which case we can apply the function t to both sides

$$t(Z_1^*, Z_2, Y) \overset{\text{d}}{=} t(Z_1, Z_2, Y)$$
$$t(D^*) \overset{\text{d}}{=} t(D)$$
$$T_1^* \overset{\text{d}}{=} T_1$$

so we have shown that under the null, T^* has the same distribution as T conditional on Z_2, Y. □

Based on the notation of lemma 1, let $Z_1 = \{X_e, X_{de}\}$, $Z_2 = X_{-e,de}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp \{X_e, X_{de}\}|X_{-e,de}$, by Lemma 1, these 100 detector values $\{T_9(i)^*\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_9 = t_9(D)$.

6.5 P-value

Then, a p-value for this detector would be the fraction of $\{T_9^*(i)\}_{i=1}^{100}$ which are greater than or equal to $t_9(D)$.

$$p(T_9) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_9^*(i) \geq T_9)}{100}$$

6.6 Scientific Intepretation of Hypothesis Test

If $p(T_9) > 0.05$, then we do not have enough evidence to reject our null hypothesis that emotional valence does not affect risk-taking. If $p(T_9) \leq 0.05$, then we do have enough evidence to reject our null hypothesis that emotional valence does not affect risk-taking, in favor of our alternative hypothesis that emotional valence does affect risk-taking.

6.7 Computational Cost

Overall, finding the p-value of detector T_9 requires a total of $1,010,000$ data sets of type A, $1,010,000$ data sets of type B, $1,010,000$ data sets of type C, $1,010,000$ data sets of type D, $2,040,200$ logistic regressions, $2,040,200$ neural networks, and $40,400$ linear regressions. That is, to compute the detector T_9, we would compute the detectors T_5, T_8. Computing T_5 will require $10,100$ data sets of type A, $10,100$ data sets of type B, $10,201$ logistic regressions, $10,201$ neural networks, and 202 linear regressions. Computing T_8 will require $10,100$ data sets of type C, $10,100$ data sets of type D, $10,201$ logistic regressions, $10,201$ neural networks, and 202 linear regressions. Computing both T_5, T_8 will require a combined total of $10,100$ data sets of type A, $10,100$ data sets of type B, 10, 100 data sets of type C, 10, 100 data sets of type D, 20, 402 logistic regressions, 20, 402 neural networks, and 404 linear regressions. Then, we would repeat this 100 times process to compute an empirical null distribution for T_9. See Appendix K for more details on the individual computational costs for calculating p-values for each test statistic.

6.8 Limitations of Hypothesis Test

As shown in the previous chapter, the neural network model is not stable across regularization levels and for all generative models, do not seem to generalize well, indicating that this test may have lower statistical power when applied to a completely new data set.

Our test statistic only examines a subset of the null hypothesis. While our hypotheses are about a Markov blanket of Y with respect to $X_e, X_{de}, X_{-\{e,de\}}$, we rely on test statistics for Markov blankets of Y with respect to $X_e, X_{-\{e,de\}}$ or with respect to $X_{de}, X_{-\{e,de\}}$.

Stepping back, importantly, we are not necessarily achieving a Markov blanket S, but instead only an approximation. Moreover, we can only learn a restricted set of many mathematical relationships between covariates and risk-taking in our generative models, while we might miss other crucial mathematical relationships between covariates and risk-taking. As a result, a lack of mathematical flexibility in our generative models might force us to include more covariates in our $\hat{S}$ and thus overestimate the number of variables in the the Markov blanket. For example, it is possible the Markov blanket is $S = \{X_1, X_2\}$ but our models do not sufficiently learn the entire relationship among $\{X_1, X_2, Y\}$, so our generative models instead rely on X_3 to compensate, so our $\hat{S} = \{X_1, X_2, X_3\}$.

Our test statistic certainly isn't looking for any possibly relationship in any model between emotional valence and risk-taking, instead we are only looking at 6 relationships rooted in three scientific hypotheses. Despite not being comprehensive, our final detector, T_9 still is valuable because it is a composition of smaller detectors that search through many scientifically likely possibilities (first and second order relationships) in different types of models (logistic regression, neural network) which are flexible to find many relationships between emotional valence and risk-taking conditional on their covariates.

While our 9 detectors are a promising start to find a relationship between emotional valence and risk-taking conditional on their covariates, it is possible we have a high chance that the null hypothesis is false but we cannot reject the null hypothesis. Notably, we could improve the Type II error by testing more possible relationships between emotional valence and risk-taking conditional on their covariate and more functional forms of emotional valence; however, doing so would yield an exponential growth in computation time beyond the scope of this thesis. Additionally, the subset of relationships between the log odds of gambling and functional forms of emotional valence we did test for were scientifically informed and probable, so we posit that the subset of models and tests used in this study are sufficient for a first pass attempt to answer whether emotional valence affects risk-taking.

Chapter 7

Conclusions

Overall, scientific literature has found fascinating explanations for the relationship between emotional valence and gambling. We have already started to leverage computer science and statistics to incorporate some of existing literature to inform a test for whether emotional valence does affect risk-taking. In doing so, this thesis also began to reconcile competing explanations by using nonlinear modeling and considering dynamics of the task itself.

7.1 Computer Science and Statistics

Future work that builds on the computer science and statistics portion of this thesis can examine a less rigid, more powerful way of detecting moderator effects for continuous variables than relying on their interaction.

Future work can develop detectors for the mood-maintenance hypothesis, affect infusion model, and reward processing hypothesis with respect to random forest models of risk-taking decisions. For example, to examine whether emotional valence has a unidirectional effect on risk-taking, one could examine how a split on emotional valence changes the proportion of gambles, which could be a proxy for how emotional valence can either increase or decrease risk-taking.

Future work can focus on reducing the computational cost of the test for whether emotional valence affects risk-taking. One potential step is to turn some components of the hypothesis test into multiple hypothesis tests and rely on Bonferroni correction, so that even reducing one layer of nested detectors cuts the computational cost by a factor of 10^2 at the expense of a more conservative, less powerful test.

Future work can conduct this hypothesis test in a separate set of participants who perform the same gambling task. Concretely, we can use the test set of participants from this thesis, in addition to extra participants who completed the gambling task on Amazon MechanicalTurk.

Future work can account for how emotional valence and risk-taking dynamically affect each other throughout the task. In this thesis, we assume probability of risk-taking is affected by emotional valence and past changes in emotional valence. Additionally, we assume emotional valence is affected by past choices to gamble, which is influenced by risk-taking. Despite these assumptions, we model both in separate regressions. Future work can look for a more formal way to describe the two-way relationship between emotional valence and risk-taking.

Future work can rely on more inferential statistics theory to start exploring mechanisms for the effect of emotional valence on risk-taking behavior.

7.2 Psychology

Future work that builds on the psychology portion of this thesis can examine more simple and more complex experiments, can look beyond task dynamics and towards broader trends with real-world utility, and can tie in broader measurements of affect or even discrete emotions, despite how emotions are difficult to define and study.

We worked with a risk-taking task in which many variables are changing at once and still we are controlling for many variables too. Future studies may examine if trends found in this thesis hold in more simple or more complex gambling task which control for more and control for less variables respectfully.

Moreover, in this thesis, we closely examined task dynamics, but we still have to examine for behavior beyond a gambling task. Examining task dynamics enabled us to synthesize literature by better isolating whether emotional valence influences risk-taking. Without accounting for task-specific elements, we run the risk of falsely attributing changes in gambling to affect when it might be the case that another covariate drives the change. Still, if we as psychologists focus too exclusively on tasks, we are moving away from real-world utility. Documenting the dynamics of behavior during a task, then incorporating that knowledge into a model is the first step to understanding the more universal components of the relationship between emotional valence and risk-taking.

Future research could benefit from going beyond emotional valence and considering arousal, cognitive appraisal, and discrete emotions. One limitation of studying only emotional valence is that two emotions with the same emotional valence can have opposite effects on decision making behavior. For example, [22] showed that anger can lead to more risky decisions because people do not think things through while another emotion with negative affect, sadness, can lead to risk-avoidant decisions because people start to consider one's affective state as information. Alternatively, future research could investigate how other components of affect (appraisal and intensity). Already, [29] explain that discrete emotional states with similar valence and arousal still have different effects on risk perception, so we might benefit from going beyond examining only valence and arousal. Also, future research can study how discrete emotional states affect gambling. One barrier is that emotions are hard to define and study in the lab, which can hinder progress. B

Appendices

Appendix A

Evaluating Predictions

For prediction, in this thesis we report the results of leave-one-out between-subject cross validation. Between-subject cross-validation allows us to evaluate whether our models can predict decisions for an unseen subject.

We report three prediction metrics:

- **balanced accuracy**, which is the average of accuracy rate for predicting gambles and accuracy rate for predicting not-gambles. We chose balanced accuracy because we equally value good prediction on choices to gamble as choices not-to-gamble
- **area under the receiver operating characteristic curve (AUC)**, which indicates performance across probability thresholds
- **average cross entropy loss**, which rewards large predicted probabilities for the correct outcome and small predicted probabilities for the incorrect outcome.

$$\mathbb{L}\text{oss}(\hat{y}, y) = \begin{cases} -\log(\hat{y}) & y = 1 \\ -\log(1-\hat{y}) & y = 0 \end{cases}$$

Appendix B

Predicting Gambling Better Than Chance & With Many Covariates

Question: For this data set, is the signal-to-noise ratio high enough to predict gambling better than chance?
Answer: Yes.

- **Question**: Are participant characteristics predictive and informative of gambling probability? **Answer**: They are informative. We could not reject the null hypothesis that they are not predictive.
- **Question**: Is considering participant characteristics, trial parameters, and summaries of the outcomes of past trials predictive and informative of gambling probability? **Answer**: Yes.
- **Question**: Is considering all the covariates (participant characteristics, trial parameters, and summaries of the outcomes of past trials) predictive and informative of gambling probability? **Answer**: Yes.
- **Question**: Is considering all the covariates (participant characteristics, trial parameters, and summaries of the outcomes of past trials) more predictive and informative of gambling probability than using some of the covariates? **Answer**: Sometimes it is, other times, we could not reject the null hypothesis that they were equally predictive and informative.

Implications for Hypothesis Test: This experiment assures us that the signal-to-noise ratio in this data collected throughout the gambling task is high enough to predict risk-taking better than by chance. This experiment also shows that by using combinations of the covariates, we can build models which are significantly more predictive than trivial models (with no covariates). Thus, in our hypothesis test models, we will use combinations of the covariates to model risk-taking. Overall, when we perform statistical inference for our research question, it boosts the relevance of our results that we are working with models that are more predictive than trivial models (with no covariates).

We want to know if the signal-to-noise ratio is high enough in this data set to predict risk-taking behavior. In other words, we'd like to know if we can predict risk-taking in this data set. If we cannot predict better than chance, we would not proceed with the analysis.

The baseline model we are trying to beat is the global mean model which predicts the probability of gambling of every trial of every participant to be the global mean gambling rate, 0.55.

A.2 Methods

We attempt to beat the baseline model with logistic regression models to predict the probability of gambling for each subject for each trial.
More specifically, for subject i on trial t, probability of gambling, $\mathrm{P}(y_{(it)} = 1)$, is modeled as:

$$\mathrm{P}(y_{(it)} = 1) = \mathrm{logit}(\eta_{(it)})$$

such that $\eta_{(it)}$ is a linear combination of features. Then, to summarize the four models:

1. Global Mean Model

$$\eta_{(it)} = \frac{\sum_{n=1}^{N} \sum_{t=0}^{T} y_{(nt)}}{N \cdot T}$$

2. Participant Model

$$\eta_{(it)} = v_0 + v_1 \text{College}_{(i)} + v_2 \text{Gender}_{(i)} + v_3 \text{Diagnosis}_{(i)} + v_4 \text{Emotional Valence}_{(i)}$$

3. Current Trial Model

$$\eta_{(it)} = v_0 + v_1 \text{Expected Reward}_{(it)} + v_2 \text{Gamble Outcome Range}_{(it)}$$

4. Past Trials Model

$$\begin{aligned}\eta_{(it)} = & v_0 + v_1 \text{Past Rewards (Equal)}_{(it)} + v_2 \text{Past Reward Prediction Errors (Equal)}_{(it)} \\ & + v_3 \text{Past Gambles (Equal)}_{(it)} + v_4 \text{Past Changes in Emotional Valence (Equal)}_{(it)}\end{aligned}$$

5. Current Trial + Past Trials Model

$$\begin{aligned}\eta_{(it)} = & v_0 + v_1 \text{Expected Reward}_{(it)} + v_2 \text{Gamble Outcome Range}_{(it)} \\ & + v_3 \text{Past Rewards (Equal)}_{(it)} + v_4 \text{Past Reward Prediction Errors (Equal)}_{(it)} \\ & + v_5 \text{Past Gambles (Equal)}_{(it)} + v_6 \text{Past Changes in Emotional Valence (Equal)}_{(it)}\end{aligned}$$

6. Participant + Current Trial Model

$$\begin{aligned}\eta_{(it)} = & w_0 + w_1 \text{College}_{(i)} + w_2 \text{Gender}_{(i)} + w_3 \text{Diagnosis}_{(i)} + w_4 \text{Emotional Valence}_{(i)} \\ & + w_5 \text{Expected Reward}_{(it)} + w_6 \text{Gamble Outcome Range}_{(it)}\end{aligned}$$

7. Participant + Past Trials Model

$$\begin{aligned}\eta_{(it)} = & w_0 + w_1 \text{College}_{(i)} + w_2 \text{Gender}_{(i)} + w_3 \text{Diagnosis}_{(i)} + w_4 \text{Emotional Valence}_{(i)} \\ & w_5 \text{Past Rewards (Equal)}_{(it)} + w_6 \text{Past Reward Prediction Errors (Equal)}_{(it)} \\ & + w_7 \text{Past Gambles (Equal)}_{(it)} + w_8 \text{Past Changes in Emotional Valence (Equal)}_{(it)}\end{aligned}$$

8. Participant + Current Trial + Past Trials Model

$$\begin{aligned}\eta_{(it)} = & w_0 + w_1 \text{College}_{(i)} + w_2 \text{Gender}_{(i)} + w_3 \text{Diagnosis}_{(i)} + w_4 \text{Emotional Valence}_{(i)} \\ & + w_5 \text{Expected Reward}_{(it)} + w_6 \text{Gamble Outcome Range}_{(it)} \\ & + w_7 \text{Past Rewards (Equal)}_{(it)} + w_8 \text{Past Reward Prediction Errors (Equal)}_{(it)} \\ & + w_9 \text{Past Gambles (Equal)}_{(it)} + w_{10} \text{Past Changes in Emotional Valence (Equal)}_{(it)}\end{aligned}$$

To preprocess the covariates, we converted binary features (college, gender, diagnosis) to 0 or 1, and we scaled all other parameters to have a standard deviation of 1. To evaluate models we looked towards prediction and inference. For prediction, in this paper we report the results of leave-one-out between-subject cross validation (Appendix A) We performed McNemar tests to evaluate whether one model had a significantly smaller or higher rate of classification than another model.

For the purpose of inferring the set of useful variables to include in models, we used likelihood ratio tests to compare nested models. We are interested in which models generate predictions that yield significantly higher likelihood of the data than that of the global mean model. We are also interested in determining whether a combined model of all the covariates generates predictions that yield a significantly higher likelihood of the data than smaller models that rely on only a subset of the covariates.

A.2 Results

We found models for this data set that predicted decisions throughout the gambling task better than chance. Relative to the global mean model that uniformly predicted a 0.55 chance of risk-taking during the gambling task, we found promising predictive models (models 4,5,6,7,8, Table B.1) of risk-taking that satisfied the following properties

- accuracy was higher, AUC was higher, cross-entropy loss was lower in leave-one-subject-out cross validation
- rate of incorrect classification in leave-one-subject-out cross validation was significantly lower (McNemar Test, $p < 0.01$)
- likelihood of the data was significantly higher (likelihood ratio test, $p < 0.01$)

We can also focus on drawing conclusions about the individual sets of covariates (participant covariates, current trial covariates, and past trial covariates) and combinations of these sets based on their predictive and inferential performance.

We reject the null hypothesis that the participant characteristics are not predictive of gambling. We accept the alternative hypothesis that participant characteristics are informative of gambling. In a leave-one-subject-out cross validation, the participant model did not have higher accuracy than the global mean model (Table B.2). Relative to the global mean model, the additional participant specific parameters of college, gender, depression, and emotional valence significantly improved the likelihood of data (likelihood ratio test $p < 0.01$, Table B.5).

We accept the alternative hypotheses that considering both current trial parameters and summaries of past trials is both predictive and informative of gambling. In a leave-one-subject-out cross validation, the trial models had higher accuracy, higher AUC and lower loss than the global mean model (Table B.2). The models of past trials had significantly better prediction than the global mean model (McNemar test $p < 0.01$, Table B.3). Relative to the global mean model, adding the covariates corresponding to the current trial (current expected reward, current gambling range) did not significantly improve the likelihood of the data relative to predicted model parameters(likelihood ratio test $p > 0.01$, Table B.5); however, adding covariates corresponding to the current trial (current expected reward, current gambling range) and past trials (past choices to gamble, past changes in emotional valence, past rewards, past reward prediction errors) or just the covariates summarizing past trials did significantly improve the likelihood of the data relative to the predicted model parameters (likelihood ratio test $p < 0.01$, Table B.5). This indicates that the set of covariates related to the past trials or the combination of all covariates related to the current trial or past trials is informative of gambling.

We accept the alternative hypotheses that the set of all covariates is informative and predictive of gambling relative to having no covariates. We sometimes accept the alternative hypothesis that the set of all covariates is more informative and predictive of gambling than using only some of the covariates. In a leave-one-subject-out cross validation, the full model had higher accuracy, higher AUC, and lower cross entropy loss than the global mean model and some of the models that only used some a subset of the covariates (Table B.2). The full model that included all covariates from the participant, the current trial, and past trials had significantly better prediction than the global mean model (McNemar Test $p < 0.01$, Table B.3) and some of the models that only used some a subset of the covariates (McNemar Test $p < 0.01$, Table B.4). The full model also had significantly higher likelihood than some of the models that used only subset of the covariates (likelihood ratio test $p < 0.01$, Table B.6).

Table B.1: Summary of Models To Beat The Baseline For each model, we evaluated it is predictive ability during a leave-one-subject-out cross validation using accuracy (ACC), area under the receiver operating characteristic curve (AUC), and cross entropy loss (Loss). We compared these performance metrics to those of the global mean model. For each model, we also conducted a McNemar Test to test if the model misclassified trials significantly less than the global mean model ($p < 0.01$). For each model, we also conducted a likelihood ratio test (LRT) to test if the additional covariates were significantly informative of gambling ($p < 0.01$). We labeled the additional covariates as informative if the LRT had a significant p-value. For model n, we labeled the additional covariates as predictive if relative to the baseline, it had higher accuracy, higher AUC, lower Loss, and significantly fewer misclassifications (McNemar $p < 0.01$). 'Additional Covariates' columns are relative to the global mean model which had no covariates.

Model	greater accuracy	greater AUC	smaller loss	McNemar p <0.01	LRT p <0.01	additional covariates informative	additional covariates predictive
1	False	False	False	False	False	False	False
2	False	True	False	False	True	True	False
3	True	True	True	True	False	False	True
4	True	True	True	True	True	True	True
5	True	True	True	True	True	True	True
6	True	True	True	True	True	True	True
7	True	True	True	True	True	True	True
8	True	True	True	True	True	True	True

Table B.2: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss. Bolds indicated most desirable value among the models.

Model #	Model	ACC	AUC	Loss
1	Global Mean	0.500	0.374	4.0E-04
2	Participant	0.479	0.501	4.0E-04
3	Current Trial	0.668	0.702	3.6E-04
4	Past Trials	0.591	0.622	3.8E-04
5	Current Trial + Past Trials	**0.684**	**0.739**	**3.4E-04**
6	Participant + Current Trial	0.664	0.698	3.6E-04
7	Participant + Past Trials	0.579	0.622	3.8E-04
8	Participant + Current Trial + Past Trials	0.669	0.734	3.5E-04

Table B.3: McNemar Test To Beat The Baseline. We performed McNemar tests to see if models significantly out predicted risk-taking relative to the global mean model.

Model	Model Name	statistic	p-value
2	Participant	0.0	1.0
3	Current Trial	267.0	5.8e-15
4	Past Trials	178.0	4.9e-07
5	Current Trial + Past Trials	231.0	2.5e-20
6	Participant + Current Trial	266.0	2.4e-14
7	Participant + Past Trials	192.0	5.9e-05
8	Participant + Current Trial + Past Trials	242.0	1.1e-16

Table B.4: McNemar Test Against the Full Model. We performed McNemar tests to see if the full model using covariates from participants, current trial, and past trials significantly out predicted other models.

Model	Model Name	statistic	p-value
1	Global Mean	242.0	1.1e-16
2	Participant	223.0	9.6e-25
3	Current Trial	134.0	0.76
4	Past Trials	213.0	2.3e-06
5	Current Trial + Past Trials	38.0	0.017
6	Participant + Current Trial	126.0	0.50
7	Participant + Past Trials	200.0	1.4e-08

Table B.5: Likelihood Ratio Tests To Beat The Baseline. We performed likelihood ratio tests for nested models. Model numbers correspond to Table B.2. Dev. stands for deviance and p stands for p-value.

Model	Model Name	DF	Dev.	p
2	Participant	4	1588.5	2.2e-16
3	Current Trial	2	0.624	0.73
4	Past Trials	4	4285.9	2.2e-16
5	Current Trial + Past Trials	6	4288.2	2.2e-16
6	Participant + Current Trial	6	4288.2	2.2e-16
7	Participant + Past Trials	8	4813.8	2.2e-16
8	Participant + Current Trial + Past Trials	10	4816.5	2.2e-16

Table B.6: Likelihood Ratio Tests Against the Full Model. We performed likelihood ratio tests for nested models. Model numbers correspond to Table B.2. Dev. stands for deviance and p stands for p-value.

Model	Model Name	DF	Dev.	p
1	Global Mean	10	4816.5	2.2e-16
2	Participant	6	3227.9	2.2e-16
3	Current Trial	8	4815.8	0.73
4	Past Trials	6	530.53	2.2e-16
5	Current Trial + Past Trials	4	528.26	2.2e-16
6	Participant + Current Trial	4	3227.8	2.2e-16
7	Participant + Past Trials	2	2.6	0.26

A.2 Discussion

We found models that significantly outperformed guessing by chance. This is promising to indicate that we can answer our research question with respect to models that predict better than chance.

Additionally, our models revealed trends similar to scientific literature about risk-taking. We didn't draw any conclusions about whether participant characteristics were predictive of gambling, but our data rejected a null hypothesis in thus supported an alternative hypothesis that participant characteristics were informative of gambling, which is consistent with literature [12] [11]. We did find that considering current and past trial parameters was predictive and informative of gambling, which is consistent with some scientific literature about the task itself affecting risk-taking [18] [17] [10].

We also found evidence that a combination of all the covariates is predictive and informative of gambling relative to no information and is sometimes predictive and sometimes informative relative to only some covariates. In our hypothesis test, we will rely on models which use a combination of covariates to model risk-taking.

Appendix C

Scientifically Supported Hypotheses About Directions of How Covariates Affect Emotional Valence and Risk-Taking In Our Data Set

Question: Is the data inconsistent with scientifically established null hypotheses about how covariates affect emotional valence and risk-taking?

Answer: No, we do not have evidence that our data is inconsistent with scientifically established null hypotheses. Moreover, using our data, we successfully failed to reject the scientifically informed null hypothesis about how covariates affect emotional valence and risk-taking.

Implications for Hypothesis Test: Based on our testing, we did not find evidence that our data is inconsistent with the presented scientific hypotheses about emotional valence and risk-taking.

Scientific literature has already proposed how some covariates affect emotional valence and risk-taking. In this experiment, we were testing

- if a statistical test on our data supported we reject these established null hypotheses. If so, then we detected significant deviations in our data from null hypotheses of existing literature, which is an indication that our data is inconsistent with scientific literature.
- if a statistical test on our our data failed to reject these established null hypotheses. If so, then we detected not significant deviations in our data

from null hypotheses of claims in existing literature. From these tests alone, we wouldn't have evidence to say our data is inconsistent with hypotheses in scientific literature.

A.3 Methods

We wanted to test if our data suggested we reject a scientifically established null hypothesis that people with depression have lower emotional valence than people without depression. We calculated an estimated difference and a confidence interval of the difference, then we used a t-test to see if the observed test statistic suggests we reject the scientifically informed null hypothesis. Using a t-test isn't exactly correct because average emotional valence is bounded between $[0, 1]$ and not normally distributed. Also, average emotional valence is not the same for all members each group. Still, this test is useful to see if the data indicates we reject or fail to reject an established, scientifically informed null hypothesis.

We wanted to test if our data suggested we reject scientifically established null hypothesis about the odds ratio of gambling between two groups with different characteristics. We calculated an estimate of the odds ratio and confidence intervals for the odds ratio, then used a Fisher's exact test to see if the observed test statistic suggests we reject the scientifically informed null hypothesis. Using the Fisher's exact test is incorrect because many subject characteristics influence gambling rate so not all participants in the same group have a gambling rate drawn from the same distribution. Still, this test is useful to see if the data indicates we reject or fail to reject three established, scientifically informed null hypotheses.

A.3 Results

Our data did not indicate we should reject scientifically informed null hypotheses about how participant characteristics inform average emotional valence (depression) and risk-taking rates (age, gender, depression diagnosis).

Emotional Valence

According to [12], people with depression have lower average emotional valence than healthy controls. Using a t-test in our data set, we successfully fail to reject this null hypothesis. The null hypothesis was that the difference in average emotional valence of healthy participants and depressed participants > 0. In our data, participants with depression were closer to unhappy on average (average = 0.46; Figure C.1) than healthy controls (average = 0.63; Figure C.1) and we successfully failed to reject the null hypothesis that people with depression have lower average emotional valence than healthy controls (95% CI of difference $[-\infty, 0.267]$, estimated difference = 0.175, t = 3.28, p-value = 0.99).

Risk-Taking

According to [12], older people gamble at lower rates. Using a Fisher's

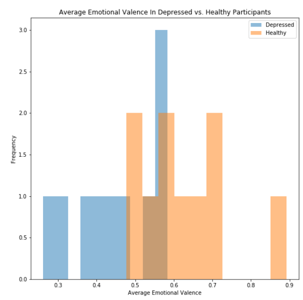

Figure C.1: The average emotional valence was centered at 0.54 with range $0.26, 0.89$. Participants with depression had a lower average emotional valence than healthy participants $(0.46 < 0.63)$

exact test in our data, we successfully failed to reject this null hypothesis. The null hypothesis corresponds to the odds ratio of gambling in college students to gambling in adolescents ≤ 1. In our data, older participants gambled less than younger patients and we successfully failed to reject the null hypothesis that older participants gambled less than younger participants (college students' gambling rate = 0.51; adolescents' gambling rate = 0.59; 95% CI of odds ratio $[0.61, \infty]$, odds ratio estimate = 0.72, p-value =0.9997; Figure C.2a).

According to [11], men gamble more than women. Using a Fisher's exact test, we successfully failed to reject this null hypothesis. The null hypothesis corresponded to the odds ratio of gambling in men to gambling in women≥ 1. In our data, men gambled at a similar rate as women, and we successfully could not reject the null hypothesis that men gamble more than women (male gambling rate = 0.547; female gambling rate = 0.553; 95% CI of odds ratio $[0.82, \infty]$, odds ratio estimate = 0.97, p-value = 0.625; Figure C.2b).

According to [13] [14] [15], we are unsure of the relationship between depression and risk-taking. Using a Fisher's exact test, we successfully fail to reject this null hypothesis. The null hypothesis corresponds to the odds ratio of gambling in participants with depression to gambling in healthy participants = 1. In our data, patients with depression gambled at a similar rate as healthy participants and we successfully could not reject the null hypothesis that participants with depression and healthy participants gambled at the same rate (depression gambling rate = 0.552; healthy gambling rate = 0.551; 95% CI of odds ratio [0.83, 1.22], odds ratio estimate = 1.01, p-value $\simeq 1$; Figure C.2c).

A.3 Discussion

Overall, our data does not indicate we should reject scientifically validated null hypotheses of how age [12], gender [11], and depression [13] [14] [15] affect risk-taking. This is not the only way our data could indicate trends which are

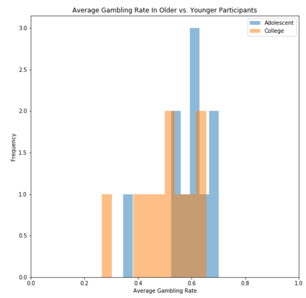

(a) By Age

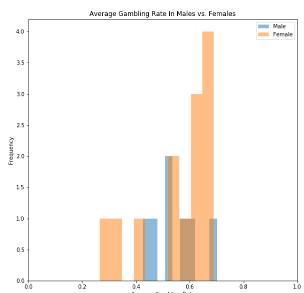

(b) By Gender

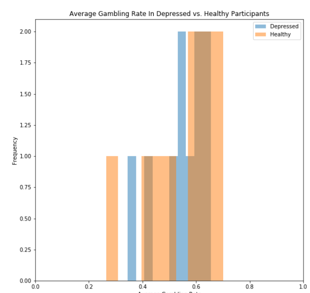

(c) By Depression

Figure C.2: Participant's gambling rate varied from 26 % to 70 %, with an average gambling rate of 55 %. Adolescent participants gambled more than college students (college students' gambling rate = 0.51, adolescents' gambling rate = 0.59). Females and males gambled at a similar rate (female gambling rate = 0.547, male gambling rate = 0.553). Participants with depression and healthy participants gambled at a similar rate (depression gambling rate = 0.552, healthy gambling rate = 0.551).

inconsistent with scientific literature, future studies could examine other ways our data supports trends inconsistent with scientific literature.

Appendix D

Model Behavior Across Regularization and Data Resamplings

Question: Are the estimated directions of the effects of covariates on risk-taking unstable across different levels of regularization or across different resamplings of the subjects?

Answer: No and sometimes. The estimated directions of the effects of covariates on risk-taking are the same even over a large grid of L1 regularization values. For some of the covariates, the estimated direction of the covariate's effect on risk-taking stayed the same across resamplings. For other covariates, the estimated direction of the covariate's effect on risk-taking did change across resamplings. Later in this part of the thesis, we will try to use other models with more flexible effects which are scientifically informed, which might better capture true effects and better generalize to different data sets. This experiment is not statistically rigorous but is instead meant to be exploratory.

Implications for Hypothesis Test: Our exploratory data analysis indicate that in our fitted prediction functions, the estimated directions of effects of covariates in on the response do not change despite different levels of regularization. The estimated directions of effects of covariates on the response do change signs when we build prediction functions on different resamplings of the data. We are not doing inference on the direction of the effect of emotional valence on risk-taking; however, these results suggest we need to be careful about generalizing models to other data sets.

In this experiment, we wanted to examine if the directions of estimated effects drastically changed in models built under different levels of regularization or built based on different resamplings of the subjects. If either were the case, we would be cautious to claim these directional effects were generalize to other data sets. This experiment is not statistically rigorous but is instead meant to be exploratory.

A.4 Methods

We were interested how the direction of estimated effects in a logistic regression of the probability of gambling changed depending on our level of L1 regularization and depending on the resampling of the dataset we used to fit the prediction function. We based our methods on a paper about stability selection [30]. We used the following logistic regression model that included all the covariates related to the participant, the current trial, and previous trials

for subject i on trial t, we modelled the log-odds of the probability of gambling, $\eta_{(it)}$ as a linear combination of features:

$$\begin{aligned}\eta_{(it)} =& w_0 + w_1 \text{College}_{(i)} + w_2 \text{Gender}_{(i)} + w_3 \text{Diagnosis}_{(i)} + w_4 \text{Emotional Valence}_{(i)} \\ &+ w_5 \text{Expected Reward}_{(it)} + w_6 \text{Gamble Outcome Range}_{(it)} \\ &+ w_7 \text{Past Rewards (Equal)}_{(it)} + w_8 \text{Past Reward Prediction Errors (Equal)}_{(it)} \\ &+ w_9 \text{Past Gambles (Equal)}_{(it)} + w_{10} \text{Past Changes in Emotional Valence (Equal)}_{(it)}\end{aligned}$$

A.4.1 Regularization Path

We were interested in a L1 regularization grid $\lambda \in \{e^{\frac{i}{2}}\}$ for $i = -8, -7, ..., 19, 20$ which corresponds to approximately $\lambda \in [10^{-2}, 10^{-4}]$. For each L1 regularization level λ, we fit a logistic regression prediction function and recorded the beta weight assigned to each covariate. We plotted all these beta weights for all these covariates for all $\lambda \in \{e^{\frac{i}{2}}\}$ for $i = -8, -7, ..., 19, 20$.

A.4.2 Stability Path

We calculated a stability for a covariate X_j in a logistic regression model trained on all covariates subject to λ level L1 regularization. We created many subsampled data sets, by sampling $n/2$ out of n participants in the train set. There are n-choose-(n/2) (ie $O(n^2)$) possible data sets, but we only sampled 100 (ie $O(n)$) at random to save computational cost. For each subsampled data set, we fit a L1-regularized logistic regression at strength λ. We chose L1 regularization because it shrinks coefficients to zero for variables with low predictive value. For a given covariate X_j, we checked the signs (+1, 0, or -1) of its coefficient across subsamples to see if the signs were the same across subsamples. We defined stability+ as the percentage of positive signed beta weights of X_j in the 100 models fit to the 100 resamplings of data. We defined stability- as the percentage of negative signed beta weights of X_j in the 100 models fit to the 100 resamplings of data.

A.4 Results

From our exploratory data analysis, our models' estimated directions of effects were not sensitive to this range of regularization strengths. On the other hand, for only some covariates, the estimated direction of the effect of the covariates on risk-taking was the same across resamplings of the data and stayed the same across the grid of regularization strengths.

A.4.1 Regularization Path

For each covariate, the estimates of the effect of the covariate on risk-taking $\text{sign}(\hat{\beta}_j^\lambda)$ in L1 regularized logistic regression prediction functions were either all ≥ 0 or all ≤ 0, so they did not change sign throughout the grid of regularization strengths (Figure D.1). This indicates that our models' estimated directions of effects were not sensitive to this range of regularization strengths.

A.4.2 Stability Path

For some of the covariates, such as expected reward, depression, past gambles, and college, either stability+ or stability- was at 100% until both were at 0% for large enough regularization λ (Figure D.2), which indicates that for these covariates, the estimated direction of the effect of the covariates on risk-taking was the same across resamplings of the data and stayed the same (or the beta weight was estimated to be 0) across the grid of regularization strengths.

On the other hand, for some covariates, the stability+ and stability- were lower than 100% and greater than 0% at many levels of regularization, which indicates that the estimated direction of the effect of the covariates on risk-taking varied across resamplings of the data (Figure D.2).

A.4 Discussion

It is promising that our models estimate the same direction of an effect for all covariates across a grid of regularization strengths. When we resample our data, some of the estimated directions of the effect of a covariate on risk-taking change, which could indicate that something about our models does not generalize to different resamplings. Scientific literature already suggests that the same covariate can either increase or decrease risk-taking behavior [18]. We are motivated to examine more flexible effects (see later experiments) than the unidirectional effects we tested in this simple logistic regression model.

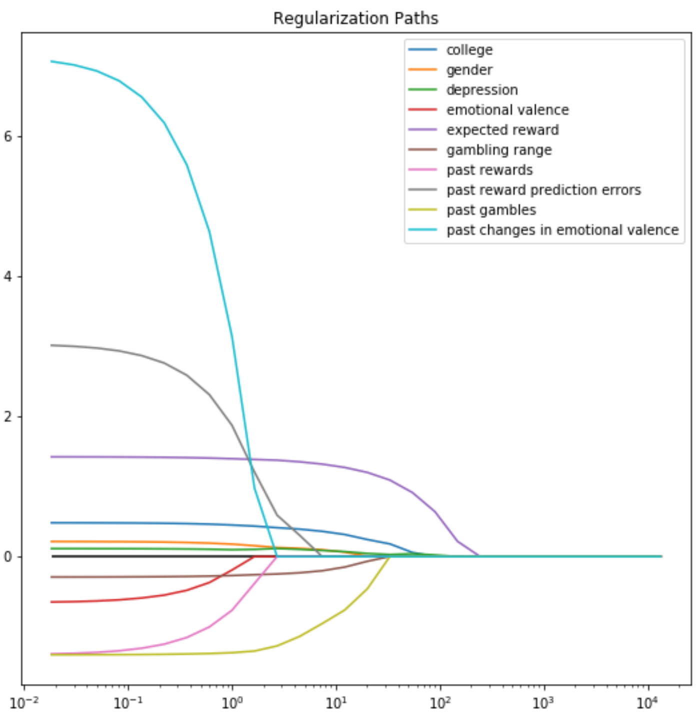

Figure D.1: The regularization path of all covariates. The x-axis is a log scale of regularization strengths, λ. The y-axis is the estimated beta weight $\hat{\beta}_j^{\lambda}$ for each covariate X_j at regularization strength λ. The black line indicates $\hat{\beta}_j^{\lambda} = 0$.

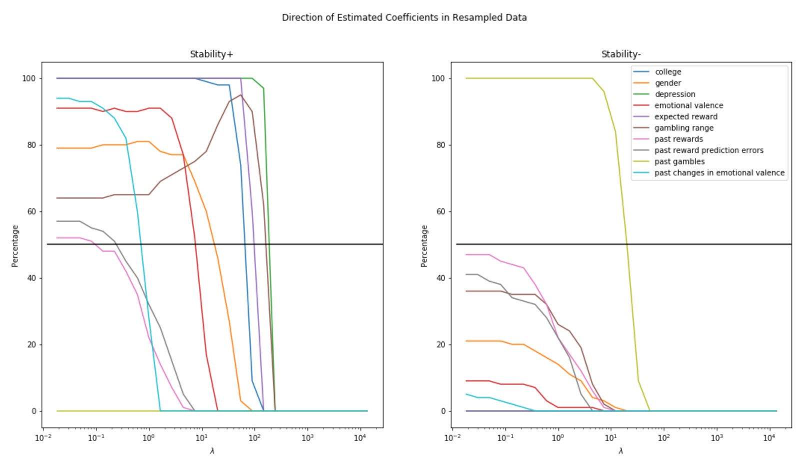

Figure D.2: The stability+ and stability- path of all covariates. The x-axis is a log scale of regularization strengths, λ. The y-axis is stability+ and stability- for each covariate X_j at regularization strength λ. The black line indicates 50%.

Appendix E

Methods for Fitting Neural Networks

here

A.5.1 Prediction Function

We built the neural networks in Python using the tensorflow package. Let η_i be a 2 layer, 5-node per layer, fully connected neural network. Each layer was fully connected with a sigmoid activation function. We tuned weights with the Adams optimizer set at learning rate= 0.001. We used adversarial noise [31] for $\epsilon \in [10, 5, 1, 1e-1, 1e-2, 1e-3, 1e-4]$. For every data point we would augment our data set with an adversarial data point which was formed using a small but worst-case perturbation on the original data point. Specifically, as in previous literature [31], we created an adversarial example with covariate value uniformly picked $[0, \epsilon]$ away from the original data point in the direction of the gradient of the original data point, with a response value the opposite of the original data point. We used l1 regularization for $\lambda \in [1e-1, 1e-2, 1e-3, 1e-4]$. We selected hyperparameter pairs (ϵ, λ) by searching the grid of possibilities and picking the pair which yielded weights that were 'stable' (defined below) and the highest leave-one-participant-out cross validated accuracy. We examined the 'stability' of weights as follows. For each tuple, we drew 4 bootstrap samples from the original data set, building 5 models (1 for the original data set, 1 for each of 4 bootstrap samples), then we took examined for each covariate X_i, the mean $\frac{\partial \eta_i}{\partial X_i}$ evaluated at each data point. These means indicated the estimated direction of unidirectional effect the covariate had on risk-taking. We considered the model 'stable' under (ϵ, λ) if 4 out of the 5 models had the same estimated direction (i.e. ≥ 4 +'s or ≥ 4 -'s) for all the covariates.

A.5.2 Smoothing The Gradient Vector and Hessian Matrix

For the neural network, we used the automatic differentiator in the tensorflow package and evaluated the estimated gradient vector and estimated Hessian matrices of η_{it} with respect to all covariates at each observation of participant i in trial t for all i, t, then averaged across observations. For the neural network, we wanted to smooth out gradient estimates and Hessian matrix element estimates with smooth gradient [27] so our estimates were less sensitive to randomness in the tuning process. We used smooth gradient at strengths $\epsilon = 0.2$, $n = 50$. We also tried $\epsilon = 1.0, 5.0$ but noticed this zeroed our average gradient values. That is, for each observed data point we directed calculated the gradients and Hessian matrices at our observed data point and at n points within a $\epsilon-$radius sphere of our observed data point, then averaged the $n + 1$ values to estimate the gradient and Hessian at that data point. By choosing $\epsilon = 0.2$ assume that the ratio of signal to noise is 1: 0.2, which is recommended by the literature [27]. We also chose to examine 50 additional points which is also recommended by literature. [27].

Appendix F

Models with Flexible Relationships Between Emotional Valence and Risk-Taking

Question: Do models learn flexible relationships between emotional valence and risk-taking?

Answer: We have some evidence that they do. This experiment is exploratory instead of statistically rigorous.

Question: Does considering other more flexible relationships between emotional valence and risk-taking compromise predictive ability?

Answer: No.

Implications for Hypothesis Test: It is promising that our generative models which are designed to detect flexible and scientifically informed relationships might actually detect nonzero flexible relationships between emotional valence and risk-taking without compromising predictive ability.

We will examine if models which consider unidirectional, bidirectional, and moderator effects between emotional valence and risk-taking predict risk-taking any worse than models which only consider unidirectional effects. All the models used the same structure of generative models that predicted the log odds of gambling at any trial for any participant as η_i. We have nonzero evidence that models learned some of these effects with no worse prediction than models which did not consider these effects. This experiment is exploratory instead of statistically rigorous.

A.6 Methods

We build three models, one that does not detect flexible relationships between emotional valence and risk-taking, and two that do.

- Model 1: Logistic Regression with Explicitly Coded unidirectional Effects. Set η_i as a linear combination of the covariates below:
  - college
  - gender
  - depression
  - emotional valence
  - expected reward
  - gambling range
  - past rewards
  - past reward prediction errors
  - past gambles
  - past changes in emotional valence
- Model 2: Flexible Logistic Regression with Explicitly Coded unidirectional, bidirectional, and Moderator Effects. Set η_i as a linear combination of the below covariates X_j and their second order forms $X_j X_i$ for all $i \neq j$ and X_j^2
  - college
  - gender
  - depression
  - emotional valence
  - expected reward
  - gambling range
  - past rewards
  - past reward prediction errors
  - past gambles
  - past changes in emotional valence
- Model 3: Neural Network with Implicitly Coded unidirectional, bidirectional, and Moderator effects. Set η_i as a 2 layer, 5-node per layer, fully connected neural network with the same covariates as model 1 and 2, except used the nonbinary form when available (age instead of college and depression severity instead of depression)

– age
– gender
– depression severity
– emotional valence
– expected reward
– gambling range
– past rewards
– past reward prediction errors
– past gambles
– past changes in emotional valence

We built the logistic regressions in Python using sklearn. For the flexible logistic regression, we used elastic net regularization at levels l1-ratio in $[0.0, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75,$

$0.8, 0.85, 0.9, 0.95]$ and α in $[1e-07, 1e-06, 1e-05, 0.0001, 0.001, 0.01, 0.03, 0.05, 0.07, 0.09, 0.11, 0.13, 0.15, 0.17, 0.19]$. We decided on a l1-ratio, α pair based on which achieved the highest balanced accuracy in leave-one-out cross validation.

We built the neural networks as explained in Appendix E

For predictive performance of each of the models, we report three measures of success in a leave-one-out between-subject cross-validation (Appendix A).

A.6 Results

We have nonzero evidence that both models picked up on flexible relationships between emotional valence and risk-taking.

A.6.1 Flexible Logistic Regression

We have some exploratory, not statistically rigorous evidence that the flexible logistic regression found flexible unidirectional, bidirectional, and moderator effects of covariates on risk-taking. We chose l1-ratio, $\alpha = 0.65, 0.001$ which maximized predictive performance.

Unidirectional Effects X_j

For any logistic regression model of probability of gambling, one indication that a covariate X_j had a unidirectional effect on the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$ is if the gradient entry of X_j: $\frac{\partial \eta}{\partial X_j}$ is nonzero. We can estimate this quantity with the estimated beta-weight for X_j: $\hat{\beta}_j$. For the flexible logistic regression model, many values had nonzero $\hat{\beta}_j$ with the largest $\hat{\beta}_j = 1.8$ (Figure F.1), indicating that the flexible logistic regression model might have picked up on unidirectional effects. This is completely exploratory, because we didn't set up a hypothesis test for $\hat{\beta}_j \neq 0$.

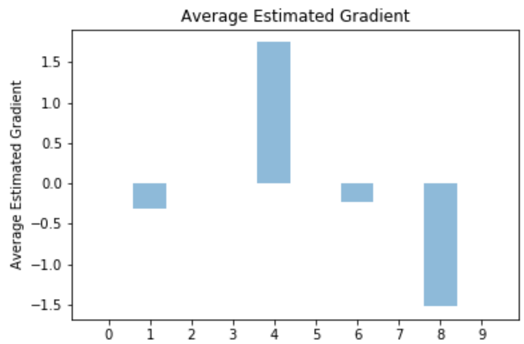

Figure F.1: The estimated gradient of η with respect to covariates. The numbered covariates corresponded to 0 = 'college', 1 = 'gender', 2 = 'depression', 3= 'emotional valence', 4 = 'expected reward', 5= 'gambling range', 6= 'past rewards', 7= 'past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'

bidirectional Effects X_j^2

One indication that a covariate X_j had a bidirectional effect on the probability of risk-taking $\text{P}(Y_{it} = 1) = \text{logit}(\eta_{it})$ is if the Hessian matrix diagonal entry of X_j: $\frac{\partial^2 \eta}{\partial X_j^2}$ is nonzero. For the logistic regression, we can estimate this quantity with the estimated beta-weight for X_j^2: $\hat{\beta}_{jj}$. For the flexible logistic regression model, many values had nonzero $\hat{\beta}_{jj}$ with the largest $\hat{\beta}_{jj} = 0.40$ (Figure F.2), indicating that the flexible logistic regression model might have picked up on bidirectional effects. This is completely exploratory, because we didn't set up a hypothesis test for $\hat{\beta}_{jj} \neq 0$.

Moderator Effects $X_j \cdot X_i$

One indication that a covariate X_j had a moderator effect on the probability of risk-taking $\text{P}(Y_{it} = 1) = \text{logit}(\eta_{it})$ is if the Hessian matrix off-diagonal entry of X_{ji}: $E(\frac{\partial^2 \eta}{\partial X_j \partial X_i})$ is nonzero. For the logistic regression, we can estimate this quantity with the estimated beta-weight for $X_j \cdot X_i$: $\hat{\beta}_{ji}$. For the flexible logistic regression model, many values had nonzero $\hat{\beta}_{ji}$ with the largest $\hat{\beta}_{ji} = 2.12$ (Figure F.2), indicating that the flexible logistic regression model might have picked up on moderator effects. This is completely exploratory, because we didn't set up a hypothesis test for $\hat{\beta}_{ji} \neq 0$.

Overall, the estimated gradient vector and estimated Hessian matrix corresponded to estimated β weights assigned to X_i and $X_i \cdot X_j$ respectively and since many of the entries of the estimated gradient and estimated Hessian matrix values were nonzero. These results are exploratory (not statistically rigorous) evidence that the logistic regression did learn unidirectional, bidirectional and moderator effects between covariates and risk-taking.

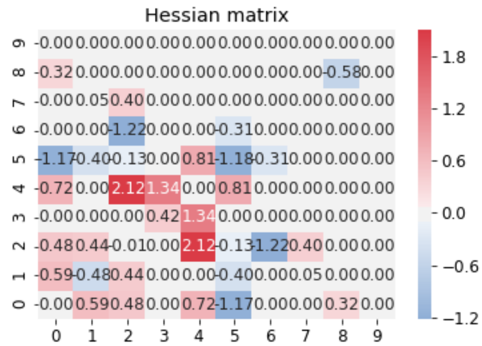

Figure F.2: The estimated Hessian matrix of η with respect to covariates. The numbered covariates corresponded to 0 = 'college', 1 = 'gender', 2 = 'depression', 3= 'emotional valence', 4 = 'expected reward', 5= 'gambling range', 6= 'past rewards', 7= 'past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'

A.6.2 Neural Network

We settled on the hyperparameters $\lambda = 0.01$ for l1 regularization and $\epsilon = 0.0001$ for adversarial noise. We have some exploratory evidence that the neural network detected some nonzero flexible on-direction, bidirectional, and moderator effects.

Unidirectional Effects X_j

We can test for unidirectional effects X_j by looking at the average gradient of η with respect to each covariate. While it is hard to say what it would mean for a neural network to detect a unidirectional effect, we can examine the average gradient of η (subject to a smoothing procedure [27]) with respect to a covariate as an indication that the neural network detected a unidirectional effect on risk-taking. If the average gradient across is large, this is some exploratory evidence that the neural network has learned a unidirectional effect. For the neural network model we trained, some average estimated gradients were nonzero (Figure F.3a), with the largest average gradient at 1.25 indicating that our neural network may have learned some unidirectional effects of the covariates on risk-taking.

bidirectional Effects X_j^2

We can test for moderator effects X_j^2 by looking at the values of diagonal Hessian entries of η with respect to covariates X_j for all j. This is because each Hessian matrix entry $\frac{\partial^2 \eta}{\partial X_j^2}$ is an indication for how a function of risk-taking η changes with respect to one covariate X_j^2.

In our data, we can evaluate a Hessian matrix for each data point, subject to a smoothing procedure [27]. It is hard to say what a bidirectional effect looks like with respect to the Hessian matrices for all data points. We posit that if the average of the Hessian matrices over all data points is nonzero on diagonal entries, then we have some evidence that the neural network is learning some

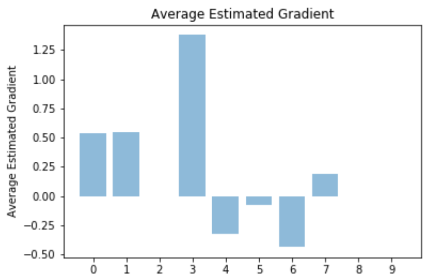

(a) Average Gradient

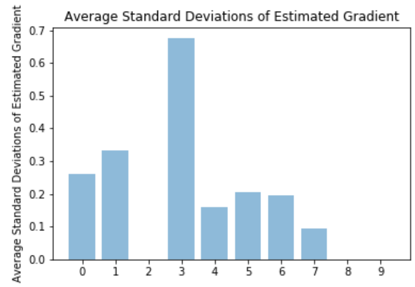

(b) Standard Deviation of Gradient

Figure F.3: The average and standard deviation of the estimated gradient of η with respect to a covariate evaluated at all data points corresponding to all participants and all trials. The numbered indices correspond to 0 = 'gender', 1 = 'depression severity', 2= 'emotional valence', 3 = 'expected reward', 4 = 'gambling range', 5 = 'age', 6 ='past rewards', 7 ='past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'.

bidirectional effects of X_j on Y. Indeed, we found some nonzero diagonal entries on the averaged Hessian, with magnitude up to 0.02 (Figure F.4a). Also, we posit that if the majority of entries are all positive or all negative, then we have some evidence that the neural network is learning some bidirectional effects of X_j on Y. To quantify this concept, we define prevalence of j as the difference in the ratio of jj^{th} Hessian matrix entries which are positive and the ratio of jj^{th} Hessian matrix entries which are negative, so that larger magnitude of prevalence indicates that many of the Hessian matrix entries were of the same sign. In our neural network, some of the prevalence values of some of the covariates were nonzero (Figure F.4b), with magnitude up to 0.13 which provides some evidence that the neural networks are learning bidirectional relationships between the covariates and the probability of gambling.

Moderator Effects $X_j \cdot X_i$

We can test for moderator effects $X_j \cdot X_i$ by looking at the off-diagonal Hessian value of η with respect to covariates X_j, X_i for $i \neq j$. This is because the Hessian entries $\frac{\partial^2 \eta}{\partial X_i \partial X_j}$ are a proxy for how a function of risk-taking η changes with respect to two covariates X_j, X_i.

In our data, we can evaluate a Hessian matrix for each data point. It is hard to say what a moderator effect looks like with respect to the ij^{th} entries of the Hessian matrices for all data points. We posit that if the average of the Hessian matrices over all data points is nonzero on off-diagonal entries, then we have some evidence that the neural network is learning some moderator effects of X_j on the relationship between X_i and Y. Indeed, we found some nonzero diagonal entries on the averaged Hessian, with magnitude up to 0.04 (Figure F.4a). Also, we posit that if the majority of entries are all positive or all negative, then this is some evidence that the neural network is learning some moderator effects of $X_j \cdot X_i$ on Y. To quantify this concept, we examine prevalence of ij as the difference in the ratio of ij^{th} Hessian matrix entries which are positive and the ratio of ij^{th} Hessian matrix entries which are negative, so that larger magnitude of prevalence indicates that many of the Hessian matrix entries were of the same sign. In our neural network, some of the prevalence values were nonzero (Figure F.4b), with magnitude up to 0.26, which provides some evidence that the neural networks are learning moderator relationships between the covariates and the probability of gambling.

(a) Average Hessian

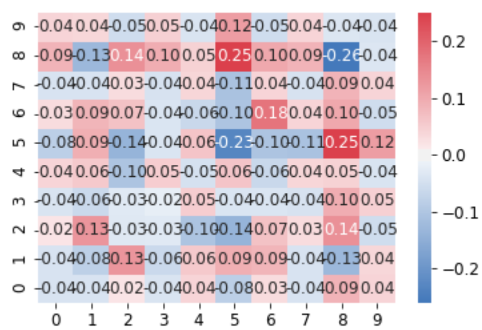

(b) Prevalence

Figure F.4: The prevalence values, positive ratio, and negative ratio of each entry of the Hessian matrix of η with respect to two covariates. The numbered indices correspond to 0 = 'gender', 1 = 'depression severity', 2= 'emotional valence', 3 = 'expected reward', 4 = 'gambling range', 5 = 'age', 6 ='past rewards', 7 ='past reward prediction errors', 8= 'past gambles', 9 ='past changes in emotional valence'.

A.6.3 Prediction

The models which can learn nonlinearities did no worse at prediction than the model that can't.

Table F.1: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss. The neural network was trained with $\lambda = 1e-3, \epsilon = 1e-2$

Model #	Model	ACC	AUC	Loss
1	Logistic Regression	0.66	0.712	3.57E-04
2	Flexible Logistic Regression (All Effects)	0.670	0.708	3.62E-04
3	Neural Network (All Effects)	0.683	0.683	3.36E-04

Appendix G

Types of Serial Position Effects of Covariates For Prediction

Question: How do we build models which summarize the past trials?
Answer: We can reduce the computational complexity of our analysis by choosing a specific version of covariates related to the past. We chose past changes in emotional valence (primacy), past gambles (equal/ no effect), past rewards (primacy), and past reward prediction errors (recency) which yielded the best predictive performance (by virtue of numerical comparison instead of statistical testing).
Implications for Hypothesis Test : In the hypothesis test, we build models which use the following specific versions of covariates which summarize the past: past changes in emotional valence (primacy), past gambles (equal/ no effect), past rewards (primacy), and past reward prediction errors (recency).

We have three versions (primacy, recency, no effect/ equal) for each of four covariates which summarize the past (past changes in emotional valence, past gambles, past rewards, and past reward prediction errors) which yields $3^4 = 81$ combinations of covariates. Instead of building models for all combinations of covariates, in this experiment, we build 12 models to pick just one combination of covariates. Then, in our hypothesis test we rely on only this combination instead of modeling all 81 possible combinations.

A.7 Methods

We relied on 12 logistic regressions models that each predicted the probability of gambling based on only one covariate. The covariate was one of the three serial position effects (primacy, recency, no effect/ equal) for each of the four covariates which summarize the past (past changes in emotional valence, past gambles, past rewards, and past reward prediction errors). Moreover, for subject i on trial t, we modeled the log odds of gambling, η_{it}, as is a linear transformation of that one covariate.

- Model 1: $\eta_{(it)} = w_0 + w_1$Past Changes in Emotional Valence (Primacy)$_{(it)}$
- Model 2: $\eta_{(it)} = w_0 + w_1$Past Changes in Emotional Valence (Recency)$_{(it)}$
- Model 3: $\eta_{(it)} = w_0 + w_1$Past Changes in Emotional Valence (Equal)$_{(it)}$
- Model 4: $\eta_{(it)} = w_0 + w_1$Past Gambles (Primacy)$_{(it)}$
- Model 5: $\eta_{(it)} = w_0 + w_1$Past Gambles (Recency)$_{(it)}$
- Model 6: $\eta_{(it)} = w_0 + w_1$Past Gambles (Equal)$_{(it)}$
- Model 7: $\eta_{(it)} = w_0 + w_1$Past Rewards (Primacy)$_{(it)}$
- Model 8: $\eta_{(it)} = w_0 + w_1$Past Rewards (Recency)$_{(it)}$
- Model 9: $\eta_{(it)} = w_0 + w_1$Past Rewards (Equal)$_{(it)}$
- Model 10: $\eta_{(it)} = w_0 + w_1$Past Reward Prediction Errors (Primacy)$_{(it)}$
- Model 11: $\eta_{(it)} = w_0 + w_1$Past Reward Prediction Errors (Recency)$_{(it)}$
- Model 12: $\eta_{(it)} = w_0 + w_1$Past Reward Prediction Errors (Equal)$_{(it)}$

To preprocess the covariates, we converted binary features (college, gender, diagnosis) to 0 or 1, and we scaled all other parameters to have a standard deviation of 1.

For prediction, we report three measures of success in a leave-one-out between-subject cross-validation (Appendix A).

For each of the four covariates (past changes in emotional valence, past gambles, past rewards, and past reward prediction errors), we selected the best representation of the covariate for the hypothesis test based on which representation of the covariates yielded a model with the best predictive ability (numerically, not with a statistical test).

A.7 Results

An estimate of predictive ability of a model trained on each version of the covariates is shown in Table H.1.

Out of the models trained on different versions of each covariate, the best estimated predictive ability was for models trained with:

Table G.1: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss.

Model #	Model	ACC	AUC	Loss
1	Past Changes In Emotional Valence Primacy	0.500	0.391	3.97E-04
2	Past Changes In Emotional Valence Recency	0.500	0.402	3.97E-04
3	Past Changes In Emotional Valence Equal/ No Effect	0.500	0.418	3.97E-04
4	Past Gambles Primacy	0.500	0.391	3.99E-04
5	Past Gambles Recency	0.500	0.476	3.98E-04
6	Past Gambles Equal/ No Effect	0.545	0.567	3.92E-04
7	Past Rewards Primacy	0.518	0.541	3.94E-04
8	Past Rewards Recency	0.493	0.465	3.98E-04
9	Past Rewards Equal/ No Effect	0.501	0.426	3.97E-04
10	Past Reward Prediction Errors Primacy	0.499	0.447	3.97E-04
11	Past Reward Prediction Errors Recency	0.500	0.384	3.98E-04
12	Past Reward Prediction Errors Equal/ No Effect	0.497	0.427	3.97E-04

- Past Changes In Emotional Valence Primacy
- Past Gambles Equal/ No Effect
- Past Rewards Primacy
- Past Reward Prediction Errors Recency

A.7 Discussion

In our hypothesis test, we rely on the chosen combination of features that summarize the past, with the best predictive performance in this experiment. This experiment is not a statistically rigorous way of choosing the best combination of covariates, but does significantly reduce our computational complexity from building models for 81 combinations of covariates to building models for just one combination of covariates.

Appendix H

Types of Rewards Covariates For Prediction

Question: How do we build models which account for rewards?
Answer: We can reduce the computational complexity of our analysis by choosing a specific version of covariates related to rewards. We chose utility-transformed reward covariates which yielded the best predictive performance (by virtue of numerical comparison instead of statistical testing)
Implications for Hypothesis Test: In the hypothesis test, we build models which use utility-transformed covariates.

In modern economic theory, people have utility functions which quantify customers' preferences over a set of choices [16]. Some utility functions are nonlinear, such that linear increases in rewards do not yield linear increases in utility, in which case looking at utility is not the same as looking at raw rewards. In this thesis, in our gambling task, it is possible people are making decisions based on utility instead of raw rewards, so in this analysis we consider two possible reward covariates: utility function transformed reward covariates and raw reward covariates.

Then, for our hypothesis test, we can cut the number of models we train in half by choosing one set of reward covariates instead of building models for both utility-transformed and raw reward covariates. We decided which set based on which set of covariates had the best predictive ability in the following experiment (only based on numerical estimates instead of rigorous statistics).

A.8 Methods

We built two models which predicted the probability of gambling based on only covariates related to reward. One model used utility-transformed covariates and the other used raw covariates. Moreover, for subject i on trial t, we modeled the log odds of gambling, $\eta_{(it)}$ as a linear combination of covariates.

- Model 1: Utility-Transformed Rewards

$$\begin{aligned}\eta_{(it)} =& w_0 + w_1\text{Expected Utility-Transformed Reward}_{(it)}\\ &+ w_2\text{Gamble Outcome Range Utility-Transformed }_{(it)}\\ &+ w_3\text{Past Utility-Transformed Rewards (Primacy)}_{(it)}\\ &+ w_4\text{Past Utility-Transformed Reward Prediction Errors (Recency)}_{(it)}\end{aligned}$$

- Model 2: Raw Rewards

$$\begin{aligned}\eta_{(it)} =& w_0 + w_1\text{Expected Reward}_{(it)} + w_2\text{Gamble Outcome Range}_{(it)}\\ &+ w_3\text{Past Rewards (Primacy)}_{(it)}\\ &+ w_4\text{Past Reward Prediction Errors (Recency)}_{(it)}\end{aligned}$$

To preprocess the covariates, we converted binary features (college, gender, diagnosis) to 0 or 1, and we scaled all other parameters to have a standard deviation of 1. To utility-transform the rewards, we applied a standard utility function ($f(x) = \frac{e^x}{e^x+1} - 0.5$) to the raw point values and then proceeded with the rest of our standard covariate construction (Figure 4.1).

For prediction, we report three measures of success in a leave-one-out between-subject cross-validation (Appendix A).

We selected the best representation of the reward covariates for the hypothesis test based on which representation of the covariates yielded a model with the best predictive ability.

A.8 Results

An estimate of the predictive ability of each version of the covariates is shown in Table H.1.

Table H.1: Predictive Modeling Results of Between-Subject Cross-Validation. For each model, we use leave one out cross validation on $n = 20$ participants and evaluate predictions. Better models have higher accuracy, higher area under the receiver operating characteristic curve, and lower cross-entropy loss. ACC stands for accuracy. AUC stands for area under receiver operating characteristic curve. Loss indicates cross-entropy loss.

Model #	Model	ACC	AUC	Loss
1	Utility-Transformed Rewards	0.677	0.701	3.59E-04
2	Raw Rewards	0.673	0.705	3.58E-04

The best predictive ability was for the model which used utility-transformed rewards.

A.8 Discussion

In our hypothesis test, we rely on the utility-transformed reward covariates. This experiment is not a statistically rigorous way of choosing the best combination of covariates, but does reduce our computational complexity from building models for 2 sets of reward covariates to building models for just one set of rewards covariates.

Appendix I

Models of Risk-Taking, Emotional Valence, & Past Changes in Emotional Valence

A.9 Models of Risk-Taking

We will use the following covariates for risk-taking according to our theoretical developments and exploratory data analysis.

- age
- college indicator
- gender
- depression indicator
- depression severity
- emotional valence
- expected reward utility-transformed
- gambling range utility-transformed
- past rewards utility-transformed primacy
- past reward prediction errors utility-transformed recency
- past gambles equal/ no effect
- past changes in emotional valence primacy

To preprocess the data, we will encode all binary covariates as 0's and 1's, we will standardize reward covariates to have standard deviation 1, we will center and standardize all remaining covariates.

As already tested in the exploratory data analysis, we will build two flexible and interpretable generative models of risk-taking designed to pick up on scientifically informed relationships between emotional valence and risk-taking. The generative models will use a latent variable enable an interpretation that mimics human decision making so that the choice for participant i to gamble or not at trial t is Y_{it} such that:

$$Y_{it} = \begin{cases} 1 & Z_{it} \geq 0 \\ 0 & Z_{it} < 0 \end{cases}$$

and the latent variable Z_{it} is a combination of observed variables X_j and unobserved variables ϵ_{it}. We define η_{it} to be the function of our known covariates of risk-taking and ϵ_{it} to be a function of unobserved variables that affect Z_{it} so that $Z_{it} = \eta_{it} + \epsilon_{it}$

And if we assume the unobserved variables follow a logistic distribution, then

$$\epsilon_{it} \sim \text{Logistic}(\mu = 0, s = 1)$$
$$\text{E}(Z_{it}) = \text{E}(\eta_{it}) + \text{E}(\epsilon_{it}) = \eta_{it}$$

$$\begin{aligned} \text{P}(Y_{it} = 1) &= \text{P}(Z_{it} \geq 0) \\ &= \text{P}(\eta_{it} + \epsilon_{it} \geq 0) \\ &= \text{P}(\epsilon_{it} \geq -\eta it) \\ &= \text{logit}(\eta_{it}) \end{aligned}$$

such that covariates effects on η_{it} mirror their effect on the probability of risk-taking .

We will pick two generative models, a logistic regression and a neural network. We will train both on specific covariates of risk-taking decided in the exploratory data analysis. We will explicitly encode the logistic regression model to learn unidirectional, bidirectional, and moderator effects between the covariates and the response, so that each model can detect three hypotheses about how emotional valence affects risk-taking: the mood-maintenance hypothesis [17], the affect-infusion model [18], and the reward processing hypothesis [19]. As shown in the exploratory data analysis, we have some exploratory, not statistically rigorous evidence that both models do pick up on these flexible relationships between emotional valence and risk-taking.

We will build logistic regressions in Python using sklearn using elastic-net regularization using l1-ratio, $\alpha = 0.65, 0.001$ as tuned in the exploratory data analysis to maximize predictive performance.

We will build neural networks in Python using the tensorflow package. The neural networks will be 2 layer, 5-node per layer, and fully connected with sigmoid activation functions. We will tune weights with the Adams optimizer set at learning rate= 0.001. We used adversarial noise [31] for $\epsilon = 0.0001$ and l1 regularization for $\lambda = 0.01$, as tuned in the exploratory data analysis to achieve baseline stability of estimated gradient vectors and to achieve the highest predictive accuracy in leave-one-subject-out cross validation across a grid of (λ, ϵ) pairs.

A.9 Models of Emotional Valence

For each of the above models of risk-taking, we will also model the conditional distribution of emotional valence given the other covariates using a linear regression of the other covariates, while modeling the error term as normally distributed. In [10], the mean of the conditional distribution of emotional valence was informed by exponentially weighted summaries of the past reward prediction errors, the past rewards from not gambling, and the past rewards from gambling, so we included all these covariates in our model. Specifically, in our model, we will have a covariate as summary of the past reward prediction errors. We will also have a covariate that is the combined summary of all past rewards, not separated into past rewards from gambling or not gambling.

When we are examining and controlling for the effect of covariates in the flexible logistic regression model of risk-taking, we will model emotional valence to have the above data generating process, X_e^{LR}, using the following covariates (all covariates which are not derived from emotional valence and not collinear):

- college indicator
- gender
- depression
- expected reward utility-transformed
- gambling range utility-transformed
- past rewards utility-transformed primacy
- past reward prediction errors utility-transformed recency
- past gambles equal/ no effect

When we are examining and controlling for the effect of covariates in the neural network model of risk-taking, we will model emotional valence to have the above data generating process X_e^{NN} using the following covariates (all covariates which are not derived from emotional valence and not collinear):

- age
- gender

- depression severity
- expected reward utility-transformed
- gambling range utility-transformed
- past rewards utility-transformed primacy
- past reward prediction errors utility-transformed recency
- past gambles equal/ no effect

We will build these models in Python using sklearn. This is a flawed model because emotional valence is a measurement in $[0, 1]$ while we are modeling average emotional valence conditional on the other covariates as normally distributed. While flawed, this model is still a useful approximation for our statistical tests.

A.9 Models of Past Changes in Emotional Valence

For each of the above models of risk-taking, we will also model the conditional distribution of past changes in emotional valence X_{de} given the other covariates using a linear regression of the other covariates, while modeling the error term as normally distributed. We rely on the same covariates as in our models of emotional valence to create models $\hat{X_{de}}^{\text{LR}}$ and $\hat{X_{de}}^{\text{NN}}$. While a better approach to modeling past changes in emotional valence is to modify our covariates and search through existing literature about changes in emotional valence, this linear regression model is probably one step in the right direction because the covariates of our models potentially inform emotional valence and past changes in emotional valence is a function of emotional valence. We build these models in Python using sklearn. This is still a flawed model because past changes in emotional valence is bounded in $[0, 1]$ while we are modeling its average to be normally distributed which has a support of $[-\infty, \infty]$.

Appendix J

Hypothesis Test: Types of Data Randomization

The specific types of data randomization correspond to our assumption about the data generating process of emotional valence:

- Data Randomization Type A: $X_e \sim \mathrm{N}(\mu, \epsilon), \mu = \sum_i \beta_i X_i$, for $X_i \in X^{\mathrm{LR}}_{-\{e,de\}}$
- Data Randomization Type B: $X_e \sim \mathrm{N}(\mu, \epsilon), \mu = \sum_i \beta_i X_i$, for $X_i \in X^{\mathrm{NN}}_{-\{e,de\}}$
- Data Randomization Type C: $X_{de} \sim \mathrm{N}(\mu, \epsilon), \mu = \sum_i \beta_i X_i$, for $X_i \in X^{\mathrm{LR}}_{-\{e,de\}}$
- Data Randomization Type D: $X_{de} \sim \mathrm{N}(\mu, \epsilon), \mu = \sum_i \beta_i X_i$, for $X_i \in X^{\mathrm{NN}}_{-\{e,de\}}$

Appendix K

Hypothesis Test: Computational Costs

The computational cost of finding the p-value for each detector is summarized below:

- T_1: 100 data sets of type A, 101 logistic regressions, and 1 linear regression
- T_2: 100 data sets of type B, 101 neural networks, and 1 linear regression
- T_3: 100 data sets of type A, 101 logistic regressions, and 1 linear regression
- T_4: 100 data sets of type B, 101 neural networks, and 1 linear regression
- T_5: $10,100$ data sets of type A, $10,100$ data sets of type B, $10,201$ logistic regressions, $10,201$ neural networks, and 202 linear regressions
- T_6: 100 data sets of type C, 101 logistic regressions, and 1 linear regression
- T_7: 100 data sets of type D, 101 neural networks, and 1 linear regression
- T_8: $10,100$ data sets of type C, $10,100$ data sets of type D, $10,201$ logistic regressions, $10,201$ neural networks, and 202 linear regressions
- T_9: $1,010,000$ data sets of type A, $1,010,000$ data sets of type B, $1,010,000$ data sets of type C, $1,010,000$ data sets of type D, $2,040,200$ logistic regressions, $2,040,200$ neural networks, and $40,400$ linear regressions

Appendix L

Hypothesis Test: Detecting the Mood-Maintenance Hypothesis and Affect-Infusion Model

We will design detectors to detect whether emotional valence affects risk-taking through the mood-maintenance hypothesis or the affect infusion model.

Then, we will design detectors T_1, T_2, T_3, T_4 for these two hypotheses under two different assumptions for the data generating processes of risk-taking and emotional valence, then summarize these detectors with detector T_5. Detector T_5 will be used in a final detector T_9 to test for all three scientific hypotheses about how emotional valence affects risk-taking (the mood-maintenance hypothesis, the affect infusion model, and the reward processing hypothesis).

A.12 Null and Alternative Hypotheses

The mood-maintenance hypothesis and affect infusion model explain how positive emotional valence decreases and increases risk-taking, respectively. Specifically, if our mathematical models of risk-taking have nonzero $\frac{\partial \eta}{\partial X_e}$, then larger emotional valence either increases risk-taking or decreases risk-taking , which means this model captured to either the affect-infusion model [18] or the mood-maintenance hypothesis [17] respectively. Also, if our mathematical models of risk-taking have nonzero $(\frac{\partial^2 \eta}{\partial X_e^2})$ then, the relationship between emotional valence and risk-taking has two directions, which means this model captured a synthesis of the affect-infusion model [18] or the mood-maintenance hypothesis [17], such that higher emotional valence can both increase and decrease risk-taking.

Since these hypotheses deal with the variable of emotional valence X_e and not the variable of past changes in emotional valence X_{de}, we will restrict our detectors to only a subset of the null hypothesis dealing with X_e, or in other words the markov blanket of Y with respect to $X_e | X_{-\{e,de\}}$:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_e) | X_{-\{e,de\}}$
  - X_e is not in the Markov blanket of Y with respect to $X_e, X_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\!\perp\!\!\!\perp X_e) | X_{-\{e,de\}}$
  - X_e is in the Markov blanket of Y with respect to $X_e, X_{-\{e,de\}}$

Each detector will make additional assumptions about the data generating processes for risk-taking and emotional valence, which help us concretely detect a unidirectional and a bidirectional effect of emotional valence on risk-taking.

A.12 Detector 1: Emotional Valence Has a unidirectional Effect on Risk-Taking (Logistic Regression)

This detector is designed to detect a subset of the alternative hypothesis that emotional valence has a unidirectional effect on risk-taking via the mood-maintenance hypothesis or the affect-infusion model.

A.12.1 Detector Assumptions/ Optimal Conditions

This detector makes two assumptions about the data generating process of risk-taking and the data generating process of emotional valence. Let $X^{\text{LR}}_{-\{e,de\}}$ be a subset of $X_{-\{e,de\}}$ such that
Let $X^{\text{LR}}_{-\{e,de\}}$:

- college indicator
- gender
- depression
- expected reward utility-transformed
- gambling range utility-transformed
- past rewards utility-transformed primacy
- past reward prediction errors utility-transformed recency
- past gambles equal/ no effect

then, this detector assumes that risk-taking follows a data generating process which depends on $X_e, X^{\text{LR}}_{-\{e,de\}}$ while emotional valence depends on only $X^{\text{LR}}_{-\{e,de\}}$. Specifically, if the data generating process for risk-taking is:

$$\begin{aligned}\eta_{(it)} =& \beta_0 + \beta_1 \text{College}_{(i)} + \beta_2 \text{Gender}_{(i)} + \beta_3 \text{Diagnosis}_{(i)} + \beta_4 \text{Emotional Valence}_{(i)} \\ &+ \beta_5 \text{Expected Reward Utility-Transformed}_{(it)} \\ &+ \beta_6 \text{Gamble Outcome Range Utility-Transformed}_{(it)} \\ &+ \beta_7 \text{Past Rewards Utility-Transformed (Primacy)}_{(it)} \\ &+ \beta_8 \text{Past Reward Prediction Errors Utility-Transformed (Recency)}_{(it)} \\ &+ \beta_9 \text{Past Gambles (Equal)}_{(it)} + \\ &+ \text{a weighted sum of all second order versions of the covariates:} \\ &\beta_{jj} X_j^2 + \beta_{jk} X_j \cdot X_k \text{ for } j \neq k \text{ and } X_k, X_j \in X_e, X^{\text{LR}}_{-e,de}\end{aligned}$$

and the data generating process for emotional valence is:

$$\begin{aligned}\text{E}(X_e^{\text{LR}}) =& \beta_0 + \beta_1 \text{College}_{(i)} + \beta_2 \text{Gender}_{(i)} + \beta_3 \text{Depression}_{(i)} \\ &+ \beta_4 \text{Expected Reward Utility-Transformed}_{(it)} \\ &+ \beta_5 \text{Gamble Outcome Range Utility-Transformed}_{(it)} \\ &+ \beta_6 \text{Past Rewards Utility-Transformed (Primacy)}_{(it)} \\ &+ \beta_7 \text{Past Reward Prediction Errors Utility-Transformed (Recency)}_{(it)} \\ &+ \beta_8 \text{Past Gambles (Equal)}_{(it)} +\end{aligned}$$

then our detector is more powerful than if the data generating processes were not as above.

Assuming risk-taking follows the flexible logistic regression model and emotional valence follows the linear regression model as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_e) | X^{\text{LR}}_{-\{e,de\}}$
  - X_e is not in the Markov blanket of Y with respect to $X_e, X^{\text{LR}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\!\perp\!\!\!\perp X_e) | X^{\text{LR}}_{-\{e,de\}}$
  - X_e is in the Markov blanket of Y with respect to $X_e, X^{\text{LR}}_{-\{e,de\}}$

A.12.2 Detector Definition

Under those assumptions about the data generating processes, this detector is designed to test whether emotional valence has a unidirectional effect on risk-taking, as consistent with the scientifically informed mood-maintenance hypothesis or affect-infusion model.

We that if know the beta coefficient of emotional valence is nonzero, that is $\beta_e \neq 0$, then emotional valence has a unidirectional effect on risk-taking. For example if $\beta_e > 0$, all other covariates held constant, then a participant with larger emotional valence during one of the trials has a larger probability of risk-taking, consistent with the affect-infusion model. Alternatively, if $\beta_e < 0$, all other covariates held constant, then a participant with larger emotional valence during one of the trials has a smaller probability of risk-taking, consistent with the mood-maintenance model. Because $\beta_e \neq 0$ is indicative of a unidirectional effect of emotional-valence on risk-taking, then we build our detector to be larger when $|\beta_4|$ is larger.

That is we define detector 1, T_1, evaluated on our data set D:

$$\begin{aligned} T_1 &= t_1(D) \\ &= |\hat{\beta}_e| \end{aligned}$$

for $\hat{\beta}_e$ as the estimated beta weight of the emotional valence covariate in an elastic net regularized logistic regression model $\hat{\eta}_{it}$ we tuned in our data set. Importantly, elastic net regularized logistic regression yields beta estimates $\hat{\beta}$ in $\hat{\eta}_{it}$ which are not always consistent with β in η_{it} [32]. So while our detector is not necessarily consistent with our parameter of interest β_e, we can still say that in some cases, this detector is large when $|\beta_e|$ is large and thus $\beta_e \neq 0$ in which case emotional valence has a unidirectional effect on risk-taking if risk-taking is generated according to the flexible logistic regression. So this detector T_1 detects for whether emotional valence has a unidirectional effect on risk-taking under the stated assumptions about the data generating process of risk-taking and emotional valence.

A.12.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_e)|X^{\text{LR}}_{-\{e,de\}}$, we can derive the empirical distribution of T_1 by computing our detector value over 100 additional randomized data sets.

We will compute the detector value on randomized data sets. First, we can use our assumed data generating process for emotional valence $X_e|X^{\text{LR}}_{-\{e,de\}}$ to generate 100 conditional randomizations of emotional valence X^*_e. Then, we can create 100 additional data sets $\{D^*_{Ai}\}_{i=1}^{100}$ such that each D^*_{Ai} is the same as our original data set D, except X_e is replaced with a conditional randomization

X_e^*. Then, we can evaluate our detector function on each of the data sets to generate 100 detector values $\{T_1(i)^*\}_{i=1}^{100} = \{t_1(D_{Ai}^*)\}_{i=1}^{100}$.

The 100 detector values $\{T_1(i)^*\}_{i=1}^{100} = \{t_1(D_{Ai}^*)\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_1 = t_1(D)$. Based on the notation of lemma 1, let $Z_1 = X_e$, $Z_2 = X_{-e,de}^{\text{LR}}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_e | X_{-e,de}^{\text{LR}}$, by Lemma 1, these 100 detector values $\{T_1(i)^*\}_{i=1}^{100} = \{t_1(D_{Ai}^*)\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_1 = t_1(D)$.

A.12.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{t_1(D_{Ai}^*)\}_{i=1}^{100}$ which are greater than or equal to $t_1(D)$.

$$p(T_1) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_1^*(i) \geq T_1)}{100}$$

A.12.5 Computational Cost

Overall, finding the p-value of detector T_1 requires 100 additional data sets, 101 logistic regressions, and 1 linear regression. To achieve the detector T_1, we would fit 1 logistic regression on our original data set. To achieve 100 samples of the empirical null distribution of detector T_1, we would need to fit 100 logistic regressions of risk-taking on 100 additional data sets generated with 1 linear regression of emotional valence.

Now we have constructed detector for a unidirectional effect of emotional valence on risk-taking in which risk-taking is generated by a flexible logistic regression of $X_e, X_{-\{e,de\}}^{\text{LR}}$ and emotional valence is generated by a linear regression of $X_{-\{e,de\}}^{\text{LR}}$. Next, we construct a detector for when the risk-taking and emotional valence data follows a data generating process aligned with a specific neural network structure of $X_e, X_{-\{e,de\}}^{\text{NN}}$ and a specific linear regression of $X_{-\{e,de\}}^{\text{NN}}$ respectfully.

A.12 Detector 2: Emotional Valence Has a unidirectional Effect on Risk-Taking (Neural Network)

This detector is designed to detect a subset of the alternative hypothesis that emotional valence has a unidirectional effect on risk-taking via the mood-maintenance hypothesis or the affect-infusion model.

A.12.1 Detector Assumptions/ Optimal Conditions

This detector makes two assumptions about the data generating process of risk-taking and the data generating process of emotional valence. Let $X^{\text{NN}}_{-\{e,de\}}$ be a subset of $X_{-\{e,de\}}$ such that
Let $X^{\text{NN}}_{-\{e,de\}}$:

- age
- gender
- depression severity
- expected reward utility-transformed
- gambling range utility-transformed
- past rewards utility-transformed primacy
- past reward prediction errors utility-transformed recency
- past gambles equal/ no effect

Note that $X^{\text{NN}}_{-\{e,de\}}$ is the same as $X^{\text{NN}}_{-\{e,de\}}$ except the college indicator and depression indicator are swapped for variables which measure approximately the same thing except have more variation, age and depression severity. Then, this detector assumes that risk-taking follows a data generating process which depends on $X_e, X^{\text{NN}}_{-\{e,de\}}$ while emotional valence depends on only $X^{\text{NN}}_{-\{e,de\}}$. Specifically, if the data generating process for risk-taking such that the log odds of gambling is a neural network:

$$\eta_{(it)} = f(X_e, X^{\text{NN}}_{-\{e,de\}})$$

where $f(X_e, X^{\text{NN}}_{-\{e,de\}})$ is a 2 layer, 5-node per layer neural network trained on the covariates $X_e, X^{\text{NN}}_{-\{e,de\}}$ and with fully connected layers each with a sigmoid activation function and the data generating process for emotional valence is such that:

$$\begin{aligned}
\text{E}(X^{\text{NN}}_e) = &\beta_0 + \beta_1 \text{Age}_{(i)} + \beta_2 \text{Gender}_{(i)} + \beta_3 \text{Depression Severity}_{(i)} \\
&+ \beta_4 \text{Expected Reward Utility-Transformed}_{(it)} \\
&+ \beta_5 \text{Gamble Outcome Range Utility-Transformed}_{(it)} \\
&+ \beta_6 \text{Past Rewards Utility-Transformed (Primacy)}_{(it)} \\
&+ \beta_7 \text{Past Reward Prediction Errors Utility-Transformed (Recency)}_{(it)} \\
&+ \beta_8 \text{Past Gambles (Equal)}_{(it)} +
\end{aligned}$$

then our detector is more powerful than if the data generating processes were not as above.

Assuming risk-taking follows the neural network model and emotional valence follows the linear regression model as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_e)|X^{\text{NN}}_{-\{e,de\}}$
  - X_e is not in the Markov blanket of Y with respect to $X_e, X^{\text{NN}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\!\perp\!\!\!\perp X_e)|X^{\text{NN}}_{-\{e,de\}}$
  - X_e is in the Markov blanket of Y with respect to $X_e, X^{\text{NN}}_{-\{e,de\}}$

A.12.2 Detector Definition

Under those assumptions about the data generating processes, this detector is designed to test whether emotional valence has a unidirectional effect on risk-taking, as consistent with the scientifically informed mood-maintenance hypothesis or affect-infusion model.

We know that if $\mathrm{E}\Big(\frac{\partial f}{\partial X_e}\Big) \neq 0$, then emotional valence has a unidirectional effect on risk-taking.

For example if $\mathrm{E}\Big(\frac{\partial f}{\partial X_e}\Big) > 0$, all other covariates held constant, then a participant with larger emotional valence during one of the trials has a larger probability of risk-taking, consistent with the affect-infusion model. Alternatively, if $\mathrm{E}\Big(\frac{\partial f}{\partial X_e}\Big) < 0$, all other covariates held constant, then a participant with larger emotional valence during one of the trials has a smaller probability of risk-taking, consistent with the mood-maintenance model. Because $\mathrm{E}\Big(\frac{\partial f}{\partial X_e}\Big) \neq 0$ is indicative of a unidirectional effect of emotional-valence on risk-taking, then we build our detector to be larger when $|\mathrm{E}\Big(\frac{\partial f}{\partial X_e}\Big)|$ is larger.

That is we define detector 2, T_2, evaluated on our data set D:

$$\begin{aligned} T_2 &= t_2(D) \\ &= \Big|\frac{1}{|D|}\sum_{it}\frac{\partial \hat{f}}{\partial X_e}(D_{it})\Big| \end{aligned}$$

for $\frac{\partial \hat{f}}{\partial X_e}(D_{it})$ as the estimated change in the log odds of gambling as emotional valence changes in a fitted l1-regularized neural network model of the log odds of gambling $\hat{f}$ which we tuned in our data set.

Importantly, fitting l1-regularized neural networks of sigmoid activation functions yields beta an estimate of $\hat{f}$ which are not always consistent with f [32]. So while our detector is not necessarily consistent with our parameter of interest $\left|\mathrm{E}\left(\frac{\partial f}{\partial X_e}\right)\right|$, we can still say that in some cases, this detector is large when $\left|\mathrm{E}\left(\frac{\partial \hat{f}}{\partial X_e}\right)\right|$ is large and thus $\mathrm{E}\left(\frac{\partial \hat{f}}{\partial X_e}\right) \neq 0$ in which case emotional valence has a unidirectional effect on risk-taking if risk-taking is generated according to the neural network. So this detector T_2 detects for whether emotional valence has a unidirectional effect on risk-taking under the stated assumptions about the data generating process of risk-taking and emotional valence.

A.12.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_e)|X^{\mathrm{NN}}_{-\{e,de\}}$, we can derive the empirical distribution of T_2 by computing our detector value over 100 additional randomized data sets.

We will compute the detector value on randomized data sets. First, we can use our assumed data generating process for emotional valence $X_e|X^{\mathrm{NN}}_{-\{e,de\}}$ to generate 100 conditional randomizations of emotional valence X^*_e. Then, we can create 100 additional data sets $\{D^*_{Bi}\}_{i=1}^{100}$ such that each D^*_{Bi} is the same as our original data set D, except X_e is replaced with a conditional randomization X^*_e. Then, we can evaluate our detector function on each of the data sets to generate 100 detector values $\{T_2(i)^*\}_{i=1}^{100} = \{t_2(D^*_{Bi})\}_{i=1}^{100}$.

The 100 detector values $\{T_2(i)^*\}_{i=1}^{100} = \{t_2(D^*_{Bi})\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_2 = t_2(D)$. Based on the notation of lemma 1, let $Z_1 = X_e$, $Z_2 = X^{\mathrm{NN}}_{-e,de}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_e|X^{\mathrm{NN}}_{-e,de}$, by Lemma 1, these 100 detector values $\{T_2(i)^*\}_{i=1}^{100} = \{t_2(D^*_{Bi})\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_2 = t_2(D)$.

A.12.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{t_2(D^*_{Bi})\}_{i=1}^{100}$ which are greater than or equal to $t_2(D)$.

$$p(T_2) = \frac{\sum_{i=1}^{100} \mathbb{I}(T^*_2(i) \geq T_2)}{100}$$

A.12.5 Computational Cost

Overall, finding the p-value of detector T_2 requires 100 additional data sets, 101 neural networks, and 1 linear regression. To achieve the detector T_2, we would fit 1 neural network on our original data set. To achieve 100 samples of the empirical null distribution of detector T_2, we would need to fit 100 neural networks of risk-taking on 100 additional data sets generated with 1 linear regression of emotional valence.

Now we have constructed two detectors for a unidirectional effect of emotional valence on risk-taking while assuming the data generating processes for risk-taking and emotional valence are

- Detector T_1: a flexible logistic regression of $X_e, X^{\text{LR}}_{-\{e,de\}}$ and a linear regression of $X^{\text{LR}}_{-\{e,de\}}$
- Detector T_2: a specific neural network structure of $X_e, X^{\text{NN}}_{-\{e,de\}}$ and a linear regression of $X^{\text{NN}}_{-\{e,de\}}$

Next, we construct two analogous detectors for a bidirectional effect of emotional valence on risk-taking.

A.12 Detector 3: Emotional Valence Has a bidirectional Effect on Risk-Taking (Logistic Regression)

This detector is designed to detect a subset of the alternative hypothesis that emotional valence has a bidirectional effect on risk-taking via the mood-maintenance hypothesis or the affect-infusion model, that is that higher emotional valence can yield both increases and decreases in risk-taking.

A.12.1 Detector Assumptions/ Optimal Conditions

This detector makes two assumptions about the data generating process of risk-taking and the data generating process of emotional valence, the same assumptions as Detector T_1. Specifically, this detector has the highest statistical power when risk-taking follows a data generating process as a logistic regression of first order and second order versions if the covariates $X_e, X^{\text{LR}}_{-\{e,de\}}$ while emotional valence is a linear regression of $X^{\text{LR}}_{-\{e,de\}}$.

Assuming risk-taking follows the flexible logistic regression model and emotional valence follows the linear regression model as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_e)|X^{\text{LR}}_{-\{e,de\}}$
  - X_e is not in the Markov blanket of Y with respect to $X_e, X^{\text{LR}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\perp\!\!\!\perp X_e)|X^{\text{LR}}_{-\{e,de\}}$
  - X_e is in the Markov blanket of Y with respect to $X_e, X^{\text{LR}}_{-\{e,de\}}$

A.12.2 Detector Definition

Under those assumptions about the data generating processes, this detector is designed to test whether emotional valence has a bidirectional effect on risk-taking, as consistent with the scientifically informed mood-maintenance hypothesis or affect-infusion model.

We that if know the beta coefficient of emotional valence raised to the second power is nonzero, that is $\beta_{ee} \neq 0$, then emotional valence has a bidirectional effect on risk-taking. For example, if $\beta_{ee} > 0$, all other covariates held constant, then, as emotional valence is larger, then the effect of emotional valence on the probability of risk-taking $\mathrm{P}(Y_{it} = 1) = \mathrm{logit}(\eta_{it})$ is more positive, so while at a smaller smaller, an increase in emotional valence will yield a decrease in the probability of risk-taking $\mathrm{P}(Y_{it} = 1)$ which is consistent with the mood-maintenance hypothesis, one the other hand at a larger emotional valence, an increase in X_j will yield a increase in the probability of risk-taking $\mathrm{P}(Y_{it} = 1)$ which is consistent with the affect infusion model. Because $\beta_{ee} \neq 0$ captures a bidirectional effect of emotional valence on risk-taking, we build our detector to be larger when $|\beta_{ee}|$ is larger.

That is we define detector 3, T_3, evaluated on our data set D:

$$
\begin{aligned}
T_3 &= t_3(D) \\
&= |\hat{\beta}_{ee}|
\end{aligned}
$$

for $\hat{\beta}_{ee}$ as the estimated beta weight of the emotional valence covariate raised to the second power in an elastic net regularized logistic regression model $\hat{\eta}_{it}$ we tuned in our data set. Importantly, elastic net regularized logistic regression yields beta estimates $\hat{\beta}$ in $\hat{\eta}_{it}$ which are not always consistent with β in η_{it} [32]. So while our detector is not necessarily consistent with our parameter of interest β_{ee}, we can still say that in some cases, this detector is large when $|\beta_{ee}|$ is large and thus $\beta_{ee} \neq 0$ in which case emotional valence has a bidirectional effect on risk-taking if risk-taking is generated according to the flexible logistic regression. So this detector T_3 detects for whether emotional valence has a bidirectional effect on risk-taking under the stated assumptions about the data generating process of risk-taking and emotional valence.

A.12.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_e)|X^{\mathrm{LR}}_{-\{e,de\}}$, we can derive the empirical distribution of T_3 by computing our detector value over 100 additional randomized data sets. We can rely on the same data sets $\{D^*_{Ai}\}_{i=1}^{100}$ as generated for detector T_1, which made the same assumptions about the data generating process for emotional valence $X_e|X^{\mathrm{LR}}_{-\{e,de\}}$. Then, we can evaluate our detector function on each of the data sets to generate 100 detector values $\{T_3(i)^*\}_{i=1}^{100} = \{t_3(D^*_{Ai})\}_{i=1}^{100}$.

The 100 detector values $\{T_3(i)^*\}_{i=1}^{100} = \{t_3(D_{Ai}^*)\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_3 = t_3(D)$. Based on the notation of lemma 1, let $Z_1 = X_e$, $Z_2 = X_{-e,de}^{\text{LR}}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_e | X_{-e,de}^{\text{LR}}$, by Lemma 1, these 100 detector values $\{T_3(i)^*\}_{i=1}^{100} = \{t_3(D_{Ai}^*)\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_3 = t_3(D)$.

A.12.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{t_3(D_{Ai}^*)\}_{i=1}^{100}$ which are greater than or equal to $t_3(D)$.

$$p(T_3) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_3^*(i) \geq T_3)}{100}$$

A.12.5 Computational Cost

Overall, finding the p-value of detector T_3 requires 100 additional data sets, 101 logistic regressions, and 1 linear regression. To achieve the detector T_3, we would fit 1 logistic regression on our original data set. To achieve 100 samples of the empirical null distribution of detector T_1, we would need to fit 100 logistic regressions of risk-taking on 100 additional data sets generated with 1 linear regression of emotional valence.

Now we have constructed detector for a unidirectional effect of emotional valence on risk-taking in which risk-taking is generated by a flexible logistic regression of $X_e, X_{-\{e,de\}}^{\text{LR}}$ and emotional valence is generated by a linear regression of $X_{-\{e,de\}}^{\text{LR}}$. Next, we construct a detector for when the risk-taking and emotional valence data follows a data generating process aligned with a specific neural network structure of $X_e, X_{-\{e,de\}}^{\text{NN}}$ and a specific linear regression of $X_{-\{e,de\}}^{\text{NN}}$ respectfully.

A.12 Detector 4: Emotional Valence Has a bidirectional Effect on Risk-Taking (Neural Network)

This detector is designed to detect a subset of the alternative hypothesis that emotional valence has a bidirectional effect on risk-taking via the mood-maintenance hypothesis or the affect-infusion model.

A.12.1 Detector Assumptions/ Optimal Conditions

This detector makes two assumptions about the data generating process of risk-taking and the data generating process of emotional valence, the same assumptions as Detector T_2. Specifically, this detector has the highest power when risk-taking follows a data generating process as a neural network of covariates $X_e, X^{\text{NN}}_{-\{e,de\}}$ while emotional valence is a linear regression of $X^{\text{NN}}_{-\{e,de\}}$.

Assuming risk-taking follows the neural network model and emotional valence follows the linear regression model as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_e)|X^{\text{NN}}_{-\{e,de\}}$
  - X_e is not in the Markov blanket of Y with respect to $X_e, X^{\text{NN}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\perp\!\!\!\perp X_e)|X^{\text{NN}}_{-\{e,de\}}$
  - X_e is in the Markov blanket of Y with respect to $X_e, X^{\text{NN}}_{-\{e,de\}}$

A.12.2 Detector Definition

Under those assumptions about the data generating processes, this detector is designed to test whether emotional valence has a bidirectional effect on risk-taking, as consistent with the scientifically informed mood-maintenance hypothesis or affect-infusion model.

We know that if $\text{E}\left(\frac{\partial^2 f}{\partial X_e^2}\right) \neq 0$, then emotional valence has a bidirectional effect on risk-taking.

For example, if $\text{E}\left(\frac{\partial^2 f}{\partial X_e^2}\right) > 0$, all other covariates held constant, then, as emotional valence is larger, then the effect of emotional valence on the probability of risk-taking $\text{P}(Y_{it} = 1) = \text{logit}(\eta_{it})$ is more positive, so while at a smaller smaller, an increase in emotional valence will yield a decrease in the probability of risk-taking $\text{P}(Y_{it} = 1)$ which is consistent with the mood-maintenance hypothesis, one the other hand at a larger emotional valence, an increase in X_j will yield a increase in the probability of risk-taking $\text{P}(Y_{it} = 1)$ which is consistent with the affect infusion model. Because $\text{E}\left(\frac{\partial^2 f}{\partial X_e^2}\right) \neq 0$ captures a bidirectional effect of emotional valence on risk-taking, we build our detector to be larger when $|\text{E}\left(\frac{\partial^2 f}{\partial X_e^2}\right)|$ is larger.

That is we define detector 4, T_4, evaluated on our data set D:

$$
\begin{aligned}
T_4 &= t_4(D) \\
&= \Big| \frac{1}{|D|} \sum_{it} \frac{\partial^2 \hat{f}}{\partial X_e^2}(D_{it}) \Big|
\end{aligned}
$$

for $\frac{\partial^2 \hat{f}}{\partial X_e^2}(D_{it})$ as the estimated change in the log odds of gambling as emotional valence changes in a fitted l1-regularized neural network model of the log odds of gambling $\hat{f}$ which we tuned in our data set.

Because we made reasonable assumptions under which our data is Markovian, by the strong law of large numbers for Markov chains,

$$\lim_{|D|\to\infty}\left|\frac{1}{|D|}\sum_{it}\frac{\partial^2 f}{\partial X_e^2}(D_{it})\right| \to \left|\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_e^2}\Big)\right|$$

Importantly, fitting l1-regularized neural networks of sigmoid activation functions yields beta an estimate of $\hat{f}$ which are not always consistent with f [32]. So while our detector is not necessarily consistent with our parameter of interest $\left|\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_e^2}\Big)\right|$, we can still say that in some cases, this detector is large when $\left|\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_e^2}\Big)\right|$ is large and thus $\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_e^2}\Big) \neq 0$ in which case emotional valence has a bidirectional effect on risk-taking if risk-taking is generated according to the neural network. So this detector T_4 detects for whether emotional valence has a bidirectional effect on risk-taking under the stated assumptions about the data generating process of risk-taking and emotional valence.

A.12.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_e)|X^{\mathrm{NN}}_{-\{e,de\}}$, we can derive the empirical distribution of T_4 by computing our detector value over 100 additional randomized data sets. We can rely on the same data sets $\{D^*_{Bi}\}_{i=1}^{100}$ as generated for detector T_2, which made the same assumptions about the data generating process for emotional valence $X_e|X^{\mathrm{NN}}_{-\{e,de\}}$. Then, we can evaluate our detector function on each of the data sets to generate 100 detector values $\{T_4(i)^*\}_{i=1}^{100} = \{t_3(D^*_{Bi})\}_{i=1}^{100}$.

The 100 detector values $\{T_4(i)^*\}_{i=1}^{100} = \{t_4(D^*_{Bi})\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_4 = t_4(D)$. Based on the notation of lemma 1, let $Z_1 = X_e$, $Z_2 = X^{\mathrm{NN}}_{-e,de}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_e|X^{\mathrm{NN}}_{-e,de}$, by Lemma 1, these 100 detector values $\{T_4(i)^*\}_{i=1}^{100} = \{t_4(D^*_{Bi})\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_4 = t_4(D)$.

A.12.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{t_4(D^*_{Bi})\}_{i=1}^{100}$ which are greater than or equal to $t_4(D)$.

$$p(T_4) = \frac{\sum_{i=1}^{100}\mathbb{I}(T_4^*(i) \geq T_4)}{100}$$

A.12.5 Computational Cost

Overall, finding the p-value of detector T_4 requires 100 additional data sets, 101 neural networks, and 1 linear regression. To achieve the detector T_4, we would fit 1 neural network on our original data set. To achieve 100 samples of the empirical null distribution of detector T_4, we would need to fit 100 neural networks of risk-taking on 100 additional data sets generated with 1 linear regression of emotional valence.

Now we have constructed four detectors for a unidirectional effect and bidirectional effect of emotional valence on risk-taking while assuming the data generating processes for risk-taking and emotional valence are

- Detector T_1, T_3: a flexible logistic regression of $X_e, X^{\mathrm{LR}}_{-\{e,de\}}$ and a linear regression of $X^{\mathrm{LR}}_{-\{e,de\}}$
- Detector T_2, T_4: a specific neural network structure of $X_e, X^{\mathrm{NN}}_{-\{e,de\}}$ and a linear regression of $X^{\mathrm{NN}}_{-\{e,de\}}$

Next, we construct detector 5 to indicate if any of these hypotheses were detected.

A.12 Detector 5: Emotional Valence Affects Risk-Taking via the Mood-Maintenance Hypothesis or the Affect Infusion Model

This detector is designed to detect any unidirectional or bidirectional effect of emotional valence on risk-taking via the mood-maintenance hypothesis and/or the affect-infusion model.

A.12.1 Detector Assumptions/ Optimal Conditions

This detector makes two possible assumptions about the data generating process of risk-taking and emotional valence:

- Detector T_1, T_3: a flexible logistic regression of $X_e, X^{\mathrm{LR}}_{-\{e,de\}}$ and a linear regression of $X^{\mathrm{LR}}_{-\{e,de\}}$
- Detector T_2, T_4: a specific neural network structure of $X_e, X^{\mathrm{NN}}_{-\{e,de\}}$ and a linear regression of $X^{\mathrm{NN}}_{-\{e,de\}}$

Assuming risk-taking and emotional valence follow the models as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_e)|(X^{\mathrm{LR}}_{-\{e,de\}} \cup X^{\mathrm{NN}}_{-\{e,de\}})$
  - X_e is not in the Markov blanket of Y with respect to $X_e, X^{\mathrm{LR}}_{-\{e,de\}} \cup X^{\mathrm{NN}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking

– $(Y \not\perp\!\!\!\perp X_e)|(X^{\mathrm{LR}}_{-\{e,de\}} \cup X^{\mathrm{NN}}_{-\{e,de\}})$
– X_e is in the Markov blanket of Y with respect to $X_e, X^{\mathrm{LR}}_{-\{e,de\}} \cup X^{\mathrm{NN}}_{-\{e,de\}}$

A.12.2 Detector Definition

We want to design our summary detector T_5 so that if our detector is large, then we have evidence to reject the null hypothesis and if it is small, then we have little evidence to reject the null hypothesis.

Specifically, when any of detectors T_1, T_2, T_3, T_4 are large, then we have some evidence to reject the null hypothesis, in which case we want detector T_5 to be large too. Alternatively, when all of the detectors T_1, T_2, T_3, T_4 are small, then we have little evidence to reject the null hypothesis, in which case we want detector T_5 to be small too.

If we define this detector so that,

$$\begin{aligned} T_5 &= t_5(D) \\ &= 1 - \min(p(T_1), p(T_2), p(T_3), p(T_4)) \end{aligned}$$

then, we know that if any of T_1, T_2, T_3, T_4 are large, then the p-value is small and T_5 will be large too. Alternatively, if all of T_1, T_2, T_3, T_4 are small, then the p-values are large and T_5 will be small too. So this detector T_5 detects for whether emotional valence has a unidirectional or bidirectional effect on risk-taking under the stated assumptions about the data generating process of risk-taking and emotional valence.

A.12.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_e)|X^{\mathrm{LR}}_{-\{e,de\}} \cup X^{\mathrm{NN}}_{-\{e,de\}}$, we can derive the empirical distribution of T_5 by computing our detector value over 100 randomized data sets. Each T_5^* will require calculating $T_1^*, T_2^*, T_3^*, T_4^*$, which will require 100 data sets $\{D^*_{Ai}\}_{i=1}^{100}$ generated assuming emotional valence is a linear regression of $X_{-\{e,de\}^{\mathrm{LR}}}$ and 100 data sets $\{D^*_{Bi}\}_{i=1}^{100}$ generated assuming emotional valence is a linear regression of $X_{-\{e,de\}^{\mathrm{NN}}}$.

The 100 detector values $\{T_5(i)^*\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_5 = t_5(D)$. Based on the notation of lemma 1, let $Z_1 = X_e$, $Z_2 = X^{\mathrm{LR}}_{-e,de} \cup X^{\mathrm{NN}}_{-e,de}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_e|(X^{\mathrm{LR}}_{-e,de} \cup X^{\mathrm{NN}}_{-e,de})$, by Lemma 1, these 100 detector values $\{T_5(i)^*\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_5 = t_5(D)$.

A.12.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{T_5^*(i)\}_{i=1}^{100}$ which are greater than or equal to $t_5(D)$.

$$p(T_5) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_5^*(i) \geq T_5)}{100}$$

A.12.5 Computational Cost

Overall, finding the p-value of detector T_5 requires a total of $10,100$ data sets of type A, $10,100$ data sets of type B, $10,201$ logistic regressions, $10,201$ neural networks, and 202 linear regressions. That is, to achieve the detector T_5, we would find the p-values of T_1, T_2, T_3, T_4 with 10^2 data sets of type A, 10^2 data sets of type B, 101 logistic regressions, 101 neural networks, and 2 linear regressions. Then, to compute the empirical null distribution of detector T_5, we would repeat the process 100 additional times.

Now, we have created a detector T_5 such that if T_5 is large and has small p-value, then we have evidence to reject the null hypothesis that emotional valence does not affect risk-taking. We have created a test that incorporates two scientific hypotheses about how emotional valence affects risk-taking: the mood-maintenance hypothesis and the affect infusion model. In our next chapter, we build two additional detectors T_6, T_7, summarized in another detector T_8, to also test for a final theory about how emotional valence affects risk-taking: the reward processing hypothesis. Then, in our next next chapter, we will create a detector T_9 that incorporates all three hypotheses about how emotional valence affects risk-taking: the mood-maintenance hypothesis, the affect infusion model, and the reward processing hypothesis.

Appendix M

Hypothesis Test: Detecting the Reward Processing Hypothesis

We will design detectors to detect whether emotional valence affects risk-taking through the reward processing hypothesis.

Then, we will design detectors T_6, T_7 for the reward processing hypothesis under two different assumptions for the data generating processes of risk-taking and emotional valence, then summarize these detectors with detector T_8. Detector T_8 will be used in a final detector T_9 to test for all three scientific hypotheses about how emotional valence affects risk-taking (the mood-maintanence hypothesis, the affect infusion model, and the reward processing hypothesis).

A.13 Null and Alternative Hypotheses

The reward processing hypothesis explains how decreasing emotional valence is correlated with less risk-taking and increasing emotional valence is correlated with more risk-taking. Specifically, if we consider X_{de} as past changes in emotional valence and X_{dg} as past decisions to gamble and if our mathematical models of risk-taking have positive $\frac{\partial^2 \eta}{\partial X_{de} \partial X_{dg}}$, then our models have detected the reward processing hypothesis (see theoretical developments and exploratory data analysis for casework details).

Since these hypotheses regard the variable of past changes in emotional valence X_{de} and not the variable of emotional valence X_e, we will restrict our detectors to only a subset of the null hypothesis dealing with X_{de}, or in other words the markov blanket of Y with respect to $X_{de}|X_{-\{e,de\}}$:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_{de})|X_{-\{e,de\}}$

- X_{de} is not in the Markov blanket of Y with respect to $X_{de}, X_{-\{e,de\}}$

- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\perp\!\!\!\perp X_{de})|X_{-\{e,de\}}$
  - X_{de} is in the Markov blanket of Y with respect to $X_{de}, X_{-\{e,de\}}$

Each detector will make additional assumptions about the data generating processes for risk-taking and emotional valence, which help us concretely detect a unidirectional and a bidirectional effect of emotional valence on risk-taking.

A.13 Detector 6: Past Changes in Emotional Valence Has a Moderator Effect on How Changes in Gambling Decision Affects Risk-Taking (Logistic Regression)

This detector is designed to detect a subset of the alternative hypothesis that past changes in emotional valence has a moderator effect the relationship between past gambles and risk-taking via the reward processing hypothesis.

A.13.1 Detector Assumptions/ Optimal Conditions

This detector makes two assumptions about the data generating process of risk-taking and the data generating process of past changes in emotional valence. Let $X^{\mathrm{LR}}_{-\{e,de\}}$ be a subset of $X_{-\{e,de\}}$ as defined for detectors T_1, T_3.
Let $X^{\mathrm{LR}}_{-\{e,de\}}$:

- college indicator
- gender
- depression
- expected reward utility-transformed
- gambling range utility-transformed
- past rewards utility-transformed primacy
- past reward prediction errors utility-transformed recency
- past gambles equal/ no effect

then, this detector assumes that risk-taking follows a data generating process which depends on $X_{de}, X^{\mathrm{LR}}_{-\{e,de\}}$ while past changes in emotional valence depends on only $X^{\mathrm{LR}}_{-\{e,de\}}$. Specifically, if the data generating process for risk-taking is:

$$
\begin{aligned}
\eta_{(it)} =& \beta_0 + \beta_1 \text{College}_{(i)} + \beta_2 \text{Gender}_{(i)} + \beta_3 \text{Diagnosis}_{(i)} + \beta_4 \text{Past Changes in Emotional Valence}_{(i)} \\
&+ \beta_5 \text{Expected Reward Utility-Transformed}_{(it)} \\
&+ \beta_6 \text{Gamble Outcome Range Utility-Transformed}_{(it)} \\
&+ \beta_7 \text{Past Rewards Utility-Transformed (Primacy)}_{(it)} \\
&+ \beta_8 \text{Past Reward Prediction Errors Utility-Transformed (Recency)}_{(it)} \\
&+ \beta_9 \text{Past Gambles (Equal)}_{(it)} + \\
&+ \text{a weighted sum of all second order versions of the covariates:} \\
&\beta_{jj} X_j^2 + \beta_{jk} X_j \cdot X_k \text{ for } j \neq k \text{ and } X_k, X_j \in X_e, X^{\text{LR}}_{-e,de}
\end{aligned}
$$

and the data generating process for past changes in emotional valence is:

$$
\begin{aligned}
\text{E}(X^{\text{LR}}_{de}) =& \beta_0 + \beta_1 \text{College}_{(i)} + \beta_2 \text{Gender}_{(i)} + \beta_3 \text{Depression}_{(i)} \\
&+ \beta_4 \text{Expected Reward Utility-Transformed}_{(it)} \\
&+ \beta_5 \text{Gamble Outcome Range Utility-Transformed}_{(it)} \\
&+ \beta_6 \text{Past Rewards Utility-Transformed (Primacy)}_{(it)} \\
&+ \beta_7 \text{Past Reward Prediction Errors Utility-Transformed (Recency)}_{(it)} \\
&+ \beta_8 \text{Past Gambles (Equal)}_{(it)} +
\end{aligned}
$$

then our detector is more powerful than if the data generating processes were not as above.

Assuming risk-taking follows the flexible logistic regression model and past changes in emotional valence follows the linear regression model as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_{de}) | X^{\text{LR}}_{-\{e,de\}}$
  - X_{de} is not in the Markov blanket of Y with respect to $X_{de}, X^{\text{LR}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\perp\!\!\!\perp X_{de}) | X^{\text{LR}}_{-\{e,de\}}$
  - X_{de} is in the Markov blanket of Y with respect to $X_{de}, X^{\text{LR}}_{-\{e,de\}}$

A.13.2 Detector Definition

Under those assumptions about the data generating processes, this detector is designed to test if past changes in emotional valence has a moderator effect on the relationship between past gambles and risk-taking, as consistent with the scientifically informed reward process sing hypothesis.

We that if know the beta coefficient of the interaction between past changes in emotional valence and past gambles is nonzero, that is $\beta_{de,dg} \neq 0$, then past changes in emotional valence has a moderator effect on the relationship between past gambles and risk-taking. For example if $\beta_{de,dg} > 0$, if past changes in emotional valence X_{de} is positive and past decisions to gamble X_{dg} is positive, then this corresponds to the scenario where the participant has gambled and seen increases in emotional valence, so the choice to gamble is a beneficial decision and under the reward processing hypothesis, the participant's gambling probability for the next trial should be higher and indeed our model mirrors that because $X_{de} \cdot X_{dg}$ is positive, $\beta_{de,dg}$ is positive, and this term adds a positive contribution to the probability of risk-taking Y, which is consistent with the reward processing hypothesis. See the theoretical developments and exploratory data analysis section for the full casework. Because $\beta_{de,dg} \neq 0$ is indicative that past changes in emotional valence has a moderator effect on the relationship between past gambles and risk-taking, then we build our detector to be larger when $|\beta_{de,dg}|$ is larger.

That is we define detector 6, T_6, evaluated on our data set D:

$$\begin{aligned} T_6 &= t_6(D) \\ &= |\hat{\beta}_{de,dg}| \end{aligned}$$

for $\hat{\beta}_{de,dg}$ as the estimated beta weight of the emotional valence covariate in an elastic net regularized logistic regression model $\hat{\eta}_{it}$ we tuned in our data set. Importantly, elastic net regularized logistic regression yields beta estimates $\hat{\beta}$ in $\hat{\eta}_{it}$ which are not always consistent with β in η_{it} [32]. So while our detector is not necessarily consistent with our parameter of interest $\beta_{de,dg}$, we can still say that in some cases, this detector is large when $|\beta_{de,dg}|$ is large and thus $\beta_{de,dg} \neq 0$ in which case past changes in emotional valence has a moderator effect on the relationship between past gambles and risk-taking. So this detector T_6 detects for if past changes in emotional valence has a moderator effect on the relationship between past gambles and risk-taking under the stated assumptions about the data generating process of risk-taking and emotional valence.

A.13.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_{de})|X^{\text{LR}}_{-\{e,de\}}$, we can derive the empirical distribution of T_6 by computing our detector value over 100 additional randomized data sets.

We will compute the detector value on randomized data sets. First, we can use our assumed data generating process for past changes in emotional valence $X_{de}|X^{\text{LR}}_{-\{e,de\}}$ to generate 100 conditional randomizations of emotional valence

X^*_{de}. Then, we can create 100 additional data sets $\{D^*_{Ci}\}_{i=1}^{100}$ such that each D^*_{Ci} is the same as our original data set D, except X_{de} is replaced with a conditional randomization X^*_{de}. Then, we can evaluate our detector function on each of the data sets to generate 100 detector values $\{T_6(i)^*\}_{i=1}^{100} = \{t_6(D^*_{Ci})\}_{i=1}^{100}$.

The 100 detector values $\{T_6(i)^*\}_{i=1}^{100} = \{t_6(D^*_{Ci})\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_6 = t_6(D)$. Based on the notation of lemma 1, let $Z_1 = X_{de}$, $Z_2 = X^{\text{LR}}_{-e,de}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_{de}|X^{\text{LR}}_{-e,de}$, by Lemma 1, these 100 detector values $\{T_6(i)^*\}_{i=1}^{100} = \{t_6(D^*_{Ci})\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_6 = t_6(D)$.

A.13.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{t_6(D^*_{Ci})\}_{i=1}^{100}$ which are greater than or equal to $t_6(D)$.

$$p(T_6) = \frac{\sum_{i=1}^{100} \mathbb{I}(T^*_6(i) \geq T_6)}{100}$$

A.13.5 Computational Cost

Overall, finding the p-value of detector T_6 requires 100 additional data sets, 101 logistic regressions, and 1 linear regression. To achieve the detector T_6, we would fit 1 logistic regression on our original data set. To achieve 100 samples of the empirical null distribution of detector T_6, we would need to fit 100 logistic regressions of risk-taking on 100 additional data sets generated with 1 linear regression of past changes in emotional valence.

Now we have constructed detector for a moderator effect of past changes in emotional valence on the relationship between past gambles and risk-taking under the assumptions that risk-taking is generated by a flexible logistic regression of $X_{de}, X^{\text{LR}}_{-\{e,de\}}$ and past changes in emotional valence is generated by a linear regression of $X^{\text{LR}}_{-\{e,de\}}$. Next, we construct a detector for when the risk-taking and past changes emotional valence data follows a data generating process aligned with a specific neural network structure of $X_{de}, X^{\text{NN}}_{-\{e,de\}}$ and a specific linear regression of $X^{\text{NN}}_{-\{e,de\}}$ respectfully.

A.13 Detector 7: Past Changes in Emotional Valence Has a Moderator Effect on How Changes in Gambling Decision Affects Risk-Taking (Neural Network)

This detector is designed to detect a subset of the alternative hypothesis that past changes in emotional valence has a moderator effect on the relationship between past gambles and risk-taking via the reward processing hypothesis.

A.13.1 Detector Assumptions/ Optimal Conditions

This detector makes two assumptions about the data generating process of risk-taking and the data generating process of past changes in emotional valence. Like in detectors T_2, T_4, let $X^{\text{NN}}_{-\{e,de\}}$ be a subset of $X_{-\{e,de\}}$ such that
Let $X^{\text{NN}}_{-\{e,de\}}$:

- age
- gender
- depression severity
- expected reward utility-transformed
- gambling range utility-transformed
- past rewards utility-transformed primacy
- past reward prediction errors utility-transformed recency
- past gambles equal/ no effect

Then, this detector assumes that risk-taking follows a data generating process which depends on $X_{de}, X^{\text{NN}}_{-\{e,de\}}$ while past changes in emotional valence depends on only $X^{\text{NN}}_{-\{e,de\}}$. Specifically, if the data generating process for risk-taking such that the log odds of gambling is a neural network:

$$\eta_{(it)} = f(X_{de}, X^{\text{NN}}_{-\{e,de\}})$$

where $f(X_{de}, X^{\text{NN}}_{-\{e,de\}})$ is a 2 layer, 5-node per layer neural network trained on the covariates $X_{de}, X^{\text{NN}}_{-\{e,de\}}$ and with fully connected layers each with a sigmoid activation function and the data generating process for past changes in emotional valence is:

$$\begin{aligned}
\text{E}(X^{\text{NN}}_{de}) =& \beta_0 + \beta_1 \text{Age}_{(i)} + \beta_2 \text{Gender}_{(i)} + \beta_3 \text{Depression Severity}_{(i)} \\
&+ \beta_4 \text{Expected Reward Utility-Transformed}_{(it)} \\
&+ \beta_5 \text{Gamble Outcome Range Utility-Transformed}_{(it)} \\
&+ \beta_6 \text{Past Rewards Utility-Transformed (Primacy)}_{(it)} \\
&+ \beta_7 \text{Past Reward Prediction Errors Utility-Transformed (Recency)}_{(it)} \\
&+ \beta_8 \text{Past Gambles (Equal)}_{(it)} +
\end{aligned}$$

then our detector is more powerful than if the data generating processes were not as above.

Assuming risk-taking follows the neural network model and past changes in emotional valence follows the linear regression model as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_{de})|X^{\text{NN}}_{-\{e,de\}}$
  - X_{de} is not in the Markov blanket of Y with respect to $X_{de}, X^{\text{NN}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\perp\!\!\!\perp X_{de})|X^{\text{NN}}_{-\{e,de\}}$
  - X_e is in the Markov blanket of Y with respect to $X_{de}, X^{\text{NN}}_{-\{e,de\}}$

A.13.2 Detector Definition

Under those assumptions about the data generating processes, this detector is designed to test if past changes in emotional valence has a moderator effect on the relationship between past gambles and risk-taking, as consistent with the scientifically informed reward processing hypothesis.

We know that if $\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}}\Big) \neq 0$, then emotional valence has a moderator effect on the relationship between past gambles and risk-taking. For example if $\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}}\Big) > 0$, past changes in emotional valence X_{de} is positive, and past decisions to gamble X_{dg} is negative, then this corresponds to the scenario where the participant has chosen not-to-gamble and seen increases in emotional valence, so the choice not-to-gamble is a beneficial decision and under the reward processing hypothesis, the participant's gambling probability for the next trial should be lower and indeed our model mirrors that because $X_{de} \cdot X_{dg}$ is negative, $\text{sign}(\frac{\partial^2 \eta}{\partial X_{de} \partial X_{dg}})$ is positive, and this term adds a negative contribution to the probability of risk-taking Y, which captures the reward processing hypothesis. See the theoretical developments and exploratory data analysis section for the full casework. Because $\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}}\Big) \neq 0$ is indicative of a moderator effect of past changes in emotional valence on the relationship between past gambles and risk-taking, then we build our detector to be larger when $|\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}}\Big)|$ is larger.

That is we define detector 7, T_7, evaluated on our data set D:

$$
\begin{aligned}
T_7 &= t_7(D) \\
&= \Big| \frac{1}{|D|} \sum_{it} \frac{\partial^2 \hat{f}}{\partial X_{de} \partial X_{dg}}(D_{it}) \Big|
\end{aligned}
$$

for $\frac{\partial^2 \hat{f}}{\partial X_{de} \partial X_{dg}}(D_{it})$ as the estimated change in the log odds of gambling as emotional valence changes in a fitted l1-regularized neural network model of the log odds of gambling $\hat{f}$ which we tuned in our data set.

Because we made reasonable assumptions under which our data is Markovian, by the strong law of large numbers for Markov chains,

$$\lim_{|D|\to\infty} \Big| \frac{1}{|D|} \sum_{it} \frac{\partial^2 \hat{f}}{\partial X_{de} \partial X_{dg}}(D_{it}) \Big| \to \Big| \mathrm{E}\Big(\frac{\partial^2 \hat{f}}{\partial X_{de} \partial X_{dg}} \Big) \Big|$$

Importantly, fitting l1-regularized neural networks of sigmoid activation functions yields beta an estimate of $\hat{f}$ which are not always consistent with f [32]. So while our detector is not necessarily consistent with our parameter of interest $\Big|\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}}\Big)\Big|$, we can still say that in some cases, this detector is large when $\Big|\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}}\Big)\Big|$ is large and thus $\mathrm{E}\Big(\frac{\partial^2 f}{\partial X_{de} \partial X_{dg}}\Big) \neq 0$ in which case emotional valence has a moderator effect on the relationship between past gambles and risk-taking. So this detector T_7 detects for whether emotional valence has a moderator effect on the relationship between past gambles and risk-taking under the stated assumptions about the data generating process of risk-taking and past changes in emotional valence.

A.13.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_{de})|X^{\mathrm{NN}}_{-\{e,de\}}$, we can derive the empirical distribution of T_7 by computing our detector value over 100 additional randomized data sets.

We will compute the detector value on randomized data sets. First, we can use our assumed data generating process for past changes in emotional valence $X_{de}|X^{\mathrm{NN}}_{-\{e,de\}}$ to generate 100 conditional randomizations of emotional valence X^*_{de}. Then, we can create 100 additional data sets $\{D^*_{Di}\}_{i=1}^{100}$ such that each D^*_{Di} is the same as our original data set D, except X_{de} is replaced with a conditional randomization X^*_{de}. Then, we can evaluate our detector function on each of the data sets to generate 100 detector values $\{T_7(i)^*\}_{i=1}^{100} = \{t_7(D^*_{Di})\}_{i=1}^{100}$.

The 100 detector values $\{T_7(i)^*\}_{i=1}^{100} = \{t_7(D^*_{Di})\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_7 = t_7(D)$. Based on the notation of lemma 1, let $Z_1 = X_{de}$, $Z_2 = X^{\mathrm{NN}}_{-e,de}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_{de}|X^{\mathrm{NN}}_{-e,de}$, by Lemma 1, these 100 detector values $\{T_7(i)^*\}_{i=1}^{100} = \{t_7(D^*_{Di})\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_7 = t_7(D)$.

A.13.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{t_7(D^*_{Di})\}_{i=1}^{100}$ which are greater than or equal to $t_7(D)$.

$$p(T_7) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_7^*(i) \geq T_7)}{100}$$

A.13.5 Computational Cost

Overall, finding the p-value of detector T_7 requires 100 additional data sets, 101 neural networks, and 1 linear regression. To achieve the detector T_7, we would fit 1 neural network on our original data set. To achieve 100 samples of the empirical null distribution of detector T_7, we would need to fit 100 neural networks of risk-taking on 100 additional data sets generated with 1 linear regression of past changes in emotional valence.

Now we have constructed two detectors for a moderator effect of emotional valence on the relationship between past gambles and risk-taking while assuming the data generating processes for risk-taking and past changes in emotional valence are

- Detector T_6: a flexible logistic regression of $X_{de}, X_{-\{e,de\}}^{\text{LR}}$ and a linear regression of $X_{-\{e,de\}}^{\text{LR}}$
- Detector T_7: a specific neural network structure of $X_{de}, X_{-\{e,de\}}^{\text{NN}}$ and a linear regression of $X_{-\{e,de\}}^{\text{NN}}$

Next, we construct detector T_8 to indicate if any of these two detectors T_6, T_7 corresponding to the reward processing hypothesis were detected.

A.13 Detector 8: Emotional Valence Affects Risk-Taking via the Reward Processing Hypothesis

This detector is designed to detect any moderator of emotional valence on the relationship between past gambles and risk-taking via the reward processing hypothesis.

A.13.1 Detector Assumptions/ Optimal Conditions

This detector makes two possible assumptions about the data generating process of risk-taking and past changes in emotional valence:

- Detector T_6: a flexible logistic regression of $X_{de}, X_{-\{e,de\}}^{\text{LR}}$ and a linear regression of $X_{-\{e,de\}}^{\text{LR}}$
- Detector T_7: a specific neural network structure of $X_{de}, X_{-\{e,de\}}^{\text{NN}}$ and a linear regression of $X_{-\{e,de\}}^{\text{NN}}$

Assuming risk-taking and past changes in emotional valence follow the models as above, then our null and alternative hypotheses take the following form:

- H_0 : Emotional valence does not affect risk-taking
  - $(Y \perp\!\!\!\perp X_{de})|(X^{\text{LR}}_{-\{e,de\}} \cup X^{\text{NN}}_{-\{e,de\}})$
  - X_{de} is not in the Markov blanket of Y with respect to $X_{de}, X^{\text{LR}}_{-\{e,de\}} \cup X^{\text{NN}}_{-\{e,de\}}$
- H_1 : Emotional valence does affect risk-taking
  - $(Y \not\perp\!\!\!\perp X_{de})|(X^{\text{LR}}_{-\{e,de\}} \cup X^{\text{NN}}_{-\{e,de\}})$
  - X_{de} is in the Markov blanket of Y with respect to $X_{de}, X^{\text{LR}}_{-\{e,de\}} \cup X^{\text{NN}}_{-\{e,de\}}$

A.13.2 Detector Definition

We want to design our summary detector T_8 so that if our detector is large, then we have evidence to reject the null hypothesis and if it is small, then we have little evidence to reject the null hypothesis.

Specifically, when any of detectors T_6, T_7 are large, then we have some evidence to reject the null hypothesis, in which case we want detector T_8 to be large too. Alternatively, when all of the detectors T_6, T_7 are small, then we have little evidence to reject the null hypothesis, in which case we want detector T_8 to be small too.

If we define this detector so that,

$$
\begin{aligned}
T_8 &= t_8(D) \\
&= 1 - \min(p(T_6), p(T_7))
\end{aligned}
$$

then, we know that if any of T_6, T_7 are large, then the p-value of at least one is small and T_8 will be large too. Alternatively, if all of T_6, T_7 are small, then the p-values are large and T_8 will be small too. So this detector T_8 detects for if past changes in emotional valence has a moderator effect on the relationship between past gambles risk-taking under the stated assumptions about the data generating process of risk-taking and past changes in emotional valence.

A.13.3 Empirical Null Distribution

Under the null hypothesis $(Y \perp\!\!\!\perp X_{de})|X^{\text{LR}}_{-\{e,de\}}$ and $(Y \perp\!\!\!\perp X_{de})|X^{\text{NN}}_{-\{e,de\}}$, we can derive the empirical distribution of T_8 by computing our detector value over 100 randomized data sets. Each T_8^* will require calculating T_6^*, T_7^*, which will require 100 data sets $\{D^*_{Ci}\}_{i=1}^{100}$ generated assuming past changes emotional valence is a linear regression of $X_{-\{e,de\}^{\text{LR}}}$ and 100 data sets $\{D^*_{Di}\}_{i=1}^{100}$ generated assuming past changes in emotional valence is a linear regression of $X_{-\{e,de\}^{\text{NN}}}$.

The 100 detector values $\{T_8(i)^*\}_{i=1}^{100}$ are the empirical distribution of our original detector value $T_8 = t_8(D)$. Based on the notation of lemma 1, let $Z_1 = X_{de}$, $Z_2 = X^{\text{LR}}_{-e,de} \cup X^{\text{NN}}_{-e,de}$, and $Y = Y$. Then, under the null hypothesis $Y \perp\!\!\!\perp X_{de}|(X^{\text{LR}}_{-e,de} \cup X^{\text{NN}}_{-e,de})$, by Lemma 1, these 100 detector values $\{T_8(i)^*\}_{i=1}^{100}$ are the empirical distribution of our original detector value evaluated in the original data set $T_8 = t_8(D)$.

A.13.4 Detector p-value

Then, a p-value for this detector would be the fraction of $\{T_8^*(i)\}_{i=1}^{100}$ which are greater than or equal to $t_8(D)$.

$$p(T_8) = \frac{\sum_{i=1}^{100} \mathbb{I}(T_8^*(i) \geq T_8)}{100}$$

A.13.5 Computational Cost

Overall, finding the p-value of detector T_8 requires a total of 10, 100 data sets of type C, 10, 100 data sets of type D, 10, 201 logistic regressions, 10, 201 neural networks, and 202 linear regressions. That is, to achieve the detector T_8, we would find the p-values of T_6, T_7 with 10^2 data sets of type C, 10^2 data sets of type D, 101 logistic regressions, 101 neural networks, and 2 linear regressions. Then, to compute the empirical null distribution of detector T_8, we would repeat the process 100 additional times.

Now, we have created a detector T_8 such that if T_8 is large and has small p-value, then we have evidence to reject the null hypothesis that emotional valence does not affect risk-taking. We have created a test that incorporates one scientific hypotheses about how emotional valence affects risk-taking: the reward processing hypothesis.

In the main text, we created a detector T_9 that incorporates all three hypotheses about how emotional valence affects risk-taking: the mood-maintenance hypothesis, the affect infusion model, and the reward processing hypothesis.

Bibliography

[1] J. Lerner, Y. Li, P. Valdesolo, and K. Kassam, "Emotion and decision making," eng, 2015. [Online]. Available: `http://www.escholarship.org/uc/item/8552m415`.

[2] L. Somerville, "Emotional development in adolescence," eng, in *Handbook of Emotions, Fourth Edition*, L. F. Barrett, M. Lewis, and J. M. Haviland-Jones, Eds., Fourth edition, Guilford Publications Inc. M.U.A., 2016, ch. 20, pp. 350–365, ISBN: 1462525342.

[3] K. R. Scherer, "What are emotions? and how can they be measured?" eng, *Social Science Information*, vol. 44, no. 4, pp. 695–729, 2005, ISSN: 0539-0184.

[4] D. Grühn and S. Scheibe, "Age-related differences in valence and arousal ratings of pictures from the international affective picture system (iaps): Do ratings become more extreme with age?" eng, *Behavior Research Methods*, vol. 40, no. 2, pp. 512–521, 2008, ISSN: 1554-351X.

[5] N. Klein S, M. Smolka M, F. Wrase J, F. Gruesser S, F. Mann K, F. Braus D, and F. Heinz A, "The influence of gender and emotional valence of visual cues on fmri activation in humans," eng, *Pharmacopsychiatry*, vol. 36, pp. 191–194, 2003, ISSN: 0936-9528.

[6] M. T. Treadway and D. H. Zald, "Reconsidering anhedonia in depression: Lessons from translational neuroscience," eng, vol. 35, no. 3, pp. 537–555, 2011, ISSN: 0149-7634.

[7] W. Schultz, P. Dayan, and P. R. Montague, "A neural substrate of prediction and reward.(special section: Cognitive neuroscience)(cover story)," *Science*, vol. 275, no. 5306, p. 1593, 1997, ISSN: 0036-8075.

[8] J. Y. Cohen, S. Haesler, L. Vong, B. B. Lowell, and N. Uchida, "Neuron-type-specific signals for reward and punishment in the ventral tegmental area," *Nature*, vol. 482, no. 7383, p. 85, 2012, ISSN: 0028-0836.

[9] K. A. Zaghloul, J. A. Blanco, C. T. Weidemann, K. Mcgill, J. L. Jaggi, G. H. Baltuch, and M. J. Kahana, "Human substantia nigra neurons encode unexpected financial rewards," eng, vol. 323, no. 5920, p. 1496, 2009, ISSN: 0036-8075.

[10] R. B. Rutledge, N. Skandali, P. Dayan, and R. J. Dolan, "A computational and neural model of momentary subjective well-being," eng, *Proceedings of the National Academy of Sciences*, vol. 111, no. 33, p. 12 252, 2014, ISSN: 0027-8424. [Online]. Available: `http://www.pnas.org/content/111/33/12252.abstract`.

[11] P. Slovic, "Risk-taking in children: Age and sex differences," eng, *Child Development*, vol. 37, no. 1, pp. 169–176, 1966, ISSN: 00093920.

[12] J. DEAKIN, M. AITKEN, T. ROBBINS, and B. J. SAHAKIAN, "Risk taking during decision-making in normal volunteers changes with age," eng, *Journal of the International Neuropsychological Society*, vol. 10, no. 4, pp. 590–598, 2004, ISSN: 1355-6177.

[13] L. Spear, "The adolescent brain and age-related behavioral manifestations," eng, *Neuroscience and Biobehavioral Reviews*, vol. 24, no. 4, pp. 417–463, 2000, ISSN: 0149-7634.

[14] J. P. Roiser, A. Mclean, A. D. Ogilvie, A. D. Blackwell, D. J. Bamber, I. Goodyer, P. B. Jones, and B. J. Sahakian, "The subjective and cognitive effects of acute phenylalanine and tyrosine depletion in patients recovered from depression," *Neuropsychopharmacology*, vol. 30, no. 4, p. 775, 2005, ISSN: 0893-133X.

[15] A. Rawal, S. Collishaw, A. Thapar, and F. Rice, "
the risks of playing it safe': A prospective longitudinal study of response to reward in the adolescent offspring of depressed parents," vol. 43, no. 1, pp. 27–38, 2013, ISSN: 0033-2917.

[16] G. Debreu, *Mathematical economics : twenty papers of Gerard Debreu ; introduction by Werner Hildenbrand*, eng, ser. Econometric Society monographs in pure theory ; no. 4. Cambridge [Cambridgeshire] ; New York: Cambridge University Press, 1983, ISBN: 052123736X.

[17] J. Juergensen, J. S. Weaver, C. N. May, and H. A. Demaree, "More than money: Experienced positive affect reduces risk-taking behavior on a real-world gambling task.(report)," English, *Frontiers in Psychology*, vol. 9, 2018, ISSN: 1664-1078.

[18] C. Cahir and K. Thomas, "Asymmetric effects of positive and negative affect on decision making," English, *Psychological Reports*, vol. 106, no. 1, pp. 193–204, 2010, ISSN: 0033-2941.

[19] N. Ressler, "Rewards and punishments, goal-directed behavior and consciousness," eng, *Neuroscience and Biobehavioral Reviews*, vol. 28, no. 1, pp. 27–39, 2004, ISSN: 0149-7634.

[20] L. F. Cummins, M. R. Nadorff, and A. E. Kelly, "Winning and positive affect can lead to reckless gambling," eng, *Psychology of Addictive Behaviors*, vol. 23, no. 2, pp. 287–294, 2009, ISSN: 0893-164X.

[21] B. Figner, R. J. Mackinlay, F. Wilkening, and E. U. Weber, "Affective and deliberative processes in risky choice: Age differences in risk taking in the columbia card task," eng, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 35, no. 3, pp. 709–730, 2009, ISSN: 0278-7393.

[22] K. Pezza Leith and R. F. Baumeister, "Why do bad moods increase self-defeating behavior? emotion, risk taking, and self-regulation," eng, *Journal of Personality and Social Psychology*, vol. 71, no. 6, pp. 1250–1267, 1996, ISSN: 0022-3514.

[23] R. Ameli, D. A. Luckenbaugh, N. F. Gould, M. K. Holmes, N. Lally, E. D. Ballard, and C. A. Zarate, "Shaps-c: The snaith-hamilton pleasure scale modified for clinician administration.," eng, *PeerJ*, vol. 2, no. 1, e429–e429, 2014, ISSN: 2167-8359. [Online]. Available: `http://search.proquest.com/docview/1545420701/`.

[24] G. Amin, S. Shah, and G. Vankar, "The prevelance and recognition of depression in primary care," *Indian Journal of Psychiatry*, vol. 40, no. 4, pp. 364–369, 1998, ISSN: 0019-5545. [Online]. Available: `http://www.indianjpsychiatry.org/article.asp?issn=0019-5545;year=1998;volume=40;issue=4;spage=364;epage=369;aulast=Amin;type=0`.

[25] B. B. Murdock, "The serial position effect of free recall," eng, *Journal of Experimental Psychology*, vol. 64, no. 5, pp. 482–488, 1962, ISSN: 0022-1015.

[26] T. G. Dietterich, "Approximate statistical tests for comparing supervised classification learning algorithms," eng, vol. 10, no. 7, pp. 1895–1923, 1998, ISSN: 0899-7667.

[27] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, and K. Been, "Sanity checks for saliency maps," eng, *arXiv.org*, 2018. [Online]. Available: `http://search.proquest.com/docview/2117280370/`.

[28] E. Candès, Y. Fan, L. Janson, and J. Lv, "Panning for gold: model-x' knockoffs for high dimensional controlled variable selection," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 80, no. 3, pp. 551–577, 2018, ISSN: 1369-7412.

[29] J. S. Lerner and D. Keltner, "Fear, anger, and risk," eng, *Journal of Personality and Social Psychology*, vol. 81, no. 1, pp. 146–159, 2001, ISSN: 0022-3514.

[30] N. Meinshausen and P. Bühlmann, "Stability selection," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 72, no. 4, pp. 417–473, 2010, ISSN: 1369-7412.

[31] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," 2014.

[32] F. Bunea, "Honest variable selection in linear and logistic regression models via ℓ_1 and $\ell_1+\ell_2$ penalization," eng, *arXiv.org*, vol. 2, no. 0, 2008, ISSN: 19357524. [Online]. Available: `http://search.proquest.com/docview/2090780889/`.