Mood-Driven Risk Preference:
How Induced Mood Affects Risk-Sensitive Learning

Claire S. Lee

A senior thesis submitted in partial fulfillment
of the requirements for the degree of
Bachelor of Arts
in the Department of Neuroscience at
Princeton University

Adviser: Professor Yael Niv

May 2020

Abstract

How does mood influence one's preference for risk through experiential learning? Mood has been shown to color various aspects of cognition, including the processes of learning and decision-making. Previous work has revealed that these very processes are also highly sensitive to risk – the variance associated with an outcome. While many studies have investigated the relationship between mood and risk-taking tendencies, this relationship in the context of trial-and-error learning has been underexplored. Drawing from theoretical and experimental findings, we propose that mood affects risk-sensitive learning through nonlinear effects on the learning of probabilistic stimuli. To test this hypothesis, we recruited and tested 150 subjects on Amazon Mechanical Turk using a risk-sensitive reinforcement learning task containing experimental mood inductions (happy, sad, or neutral). We addressed the following research aims: (1) to examine mood's effects on the learning of deterministic vs. probabilistic stimuli, (2) to compare distinct computational cognitive models of risk-sensitive learning, and (3) to tease out the mechanism by which mood drives risk preferences within the framework of the best-fitting model. Our behavioral results demonstrate a significant link between mood and risk attitudes, with a happy induction, relative to a sad induction, predicting a greater preference for risk. While the specific mechanism by which mood modulates risk preference is unclear, our results suggest the possibility of a more nuanced, dynamic model with mood in interaction with asymmetric learning.

Contents

Listing of figures

Listing of tables

Acknowledgments

I would like to thank Daniel Bennett for being such an amazing mentor and for truly being there every step of the way. Whether in your office at PNI, at an Australian cafe in New York City, or through Zoom, thank you so much for your countless hours of dedication. I have grown so much as a logical thinker and I have always appreciated your warmth and energy.

I would like to thank Professor Yael Niv for being such a wonderful advisor, mentor, and professor. Thank you for supporting me in every way, both inside and outside the lab; you have always been so receptive of new ideas and I feel grateful to have had an advisor who cared for me not only as a student, but also as a person.

Thank you to everyone in the Niv Lab. Special shout-out to Mingyu Song - thank you for your feedback on my thesis and for always making my day with your contagious smiles.

I would also like to thank Sara Berger and Guillermo Cecchi at IBM Research for your generous mentorship - I have left every one of our meetings feeling inspired and excited to pursue research in computational psychiatry.

Thank you to my friends - you have made my Princeton experience so fun. I'll cherish all our Firestone study sessions, cozy movie nights, and endless conversations.

Mom, Dad, and Rachel - words cannot express my gratitude for your unconditional love and support. You are the best family ever. Thank you for truly everything.

1

Introduction

Imagine the following scenario. You are a frequent visitor to a certain restaurant, and you are deciding between two options. You know that the fish taco is delicious… but you know that sometimes it is bad. On the other hand, the beef taco is pretty good every time – not bad, not delicious – just reliably good.

Would you take the bet or would you play it safe? "It depends," you say. And it certainly does. Navigating an uncertain world, we make countless choices every day, learning from trial-and-error along the way. We have good days and bad days, and there is not a single moment during which we do not experience mood.

So, back to the scenario. Would you be more likely to order the fish taco on a good day or a bad day? We set forth to test this out, specifically with the following question in mind: how does mood influence our attitudes toward risk?

Through behavioral analyses and computational modeling, we sought to uncover the neurocomputational processes involved in mood-driven risky decision-making.

1.1 Risk-Sensitive Reinforcement Learning

Risky decision-making has been a topic of great interest in the fields of neuroscience[57], behavioral economics[20], and psychology[51]. From financial investors[58] to reckless teenagers[68] to patients with bipolar disorder[56], risk-takers abound in every corner of society. Within ourselves, the holistic influence of cognitive, emotional, and hormonal factors have been shown to modulate decision-making under risk[38], defined as the variance associated with an outcome.

The beef taco is always pretty good, with zero variance; its outcome is deterministic. The fish taco, on the other hand, is risky; its outcome is probabilistic.

Risk preference is defined as an individual's general susceptibility to or avoidance of risky prospects. In the hypothetical scenario above, a risk-seeking individual might prefer the fish taco, while a risk-averse individual might prefer the beef taco.

While aspects of risky behavior have been well explored in experimental contexts where information about risk is explicitly provided[17,29,36,54], they have been less explored in the context of trial-and-error learning where the probability of an outcome is uncertain[28,30]. In the restaurant scenario, you do not know the exact probability with which the fish taco might be delicious, and this is true for most other instances in the real world where we are often not provided such knowledge. It is mostly through experience that we learn about different outcomes and their variance[25]. Experiential decision-making has been shown to depend on distinct cognitive processes from those evoked by explicit knowledge, suggesting a high degree of context-dependence for the emergence of risk attitudes[32].

Experiential learning and decision-making is the essence of reinforcement learning (RL), a field of machine learning that has strongly influenced neuroscience[45]. It is defined as an adaptive process in which an agent utilizes its previous experiences to improve the outcomes of future choices in an uncertain environment[66]. This framework has been instrumental in advancing our understanding of human learning and decision-making[10,45,13,35], all the more popular for its neurobiological basis. Dopaminergic neurons in the midbrain are believed to encode the hallmark of the RL model, called the prediction error, by firing more actively from positive prediction errors (due to better-than-expected outcomes) and less actively from negative prediction errors (due to less-than-expected outcomes)[62].

In a study by Niv (2012), the same neural correlates for RL were also found to be sensitive to experienced risk, suggesting that risk sensitivity plays a crucial role in learning[46]. Niv (2012) considered two main ways by which such risk sensitivity could result: nonlinear subjective utilities for uncertain outcomes[6] and nonlinear effects on learning for uncertain outcomes[43]. In a simple reinforcement learning framework[66], subjects use past experience to learn the values of different stimuli based on prediction errors (discrepancies between expected and actual outcomes). While this model has proven successful in explaining a wide range of behavioral and neural data, it cannot adequately model risk sensitivity because it focuses on estimating the mean outcome (and does not track variance). Particularly for cases in which there is no highest value option, the model's linear treatment of outcome valuation and learning prevents it from capturing risk attitudes. On the other hand, risk-sensitive models like the subjective utility model[6] and the asymmetric learning model[43] are able to capture risk preferences through their nonlinear effects on the processes of reinforcement learning. Niv (2012) found the asymmetric learning model, which treats unexpected positive outcomes differently from unexpected negative outcomes, to best capture risk-sensitive behavior (continued in Chapter 2).

1.2 Mood-Sensitive Reinforcement Learning

As humans, we are innately sensitive to the pervasive effects of mood, defined as a state of affect that typically lasts longer than emotions[16]. While good and bad outcomes have been known to affect mood[42,64], it was unclear whether mood also affected outcome valuations. Importantly, a study by Eldar & Niv (2015) found RL to be sensitive to mood[15]. Ultimately, interactions between mood and cognition are bidirectional; mood affects cognitive processes (like perception, attention, memory, and executive functions)[67] and these same processes also affect our mood[37].

The important relationship between mood and cognition is further spotlighted when we consider how things could go wrong; theoretical modeling showed that mood instability may result from a positive-feedback effect between mood and cognition[15,4], resulting in disturbances in both mood and cognition.

Mood is commonly thought to constitute abstract, subjective feelings that are difficult to grasp, even descriptively[21,52]. However, computational modeling has suggested that mood, as well as other subjective states, could be captured more objectively. Specifically, mood has been proposed as a "representation of momentum" – the cumulative impact of prediction errors, which then biases both the perception and learning of outcomes[16]. Emotional states were found to essentially feed back onto the perception of outcomes, biasing valuations in a mood-driven way.

1.3 Mood and Risky Behavior

How exactly does mood affect risky behavior?

A large bulk of the literature points to positive moods being linked to higher risk-taking tendencies[24] and negative moods being linked to lower risk-taking tendencies[77]. When we consider the manic episode (pathologically elevated mood) of bipolar disorder[49], increased risk-taking is even a diagnostic criterion for a manic episode[1]. One of the most prominent theories explaining these effects is the Affect Infusion Model (AIM)[18], which proposed that a positive mood would make individuals more risk-seeking because they would rely more on positive cues during the judgment process. In the hypothetical restaurant scenario, an elevated mood may make you more likely to order the fish taco by focusing your thoughts on memories of when the fish taco was delicious (and less on memories of the bad meals).

Yet other studies have found the opposite results, with positive mood states reducing real-world risk-taking behavior[33] and negative mood states contributing to greater risk-aversion in everyday decision-making[27]. A study evaluating age-related differences driving mood's influence on risk-taking tendencies found an asymmetrical effect of positive and negative mood on risk preference[9]. The Mood Maintenance Hypothesis (MMH)[31], predicts the opposite effect from AIM. According to this theory, for people in a positive mood, there is a greater desire to maintain their state of happiness. As a result, a positive mood is predicted to make people more cautious, and ultimately more risk-averse. If you are in a great mood, perhaps you would be more inclined to stick with the beef taco just in case the fish taco turns out to be bad. If you are having a bad day, maybe you would go ahead and order the fish taco, since it may make your day better.

As the literature shows, the strong relationship between mood and risk preference is undeniable, but conflicting results suggest a strong dependence on context. State-level vs. trait-level effects, experimental vs. real-world settings, normal vs. pathological mood states, financial vs. health-related risks, and risk for gains vs. losses are just a few examples of how differential risk attitudes may emerge as a function of not only mood, but also important mediating variables.

1.4 Mood-Driven Risk-Sensitive Reinforcement Learning

The studies by Niv (2012) and Eldar & Niv (2015) showed that our learning and decision-making processes are sensitive to both risk and mood. Taken together, these two studies provide potential evidence for a mood-driven risk-sensitive RL process that may result in distinct attitudes toward risk.

In our study, we contextualize our focus on mood-driven risk-taking under the framework of reinforcement learning. The interaction between mood and risk attitudes in the context of experiential learning remains unclear; a major objective of this study is to elucidate if and how mood modulates risk preferences in this novel framework. If mood is shown to modulate risk attitudes during trial-and-error learning, computational modeling may clarify how. The asymmetric learning model and the utility model represent distinct mechanisms by which choices within a risk-sensitive learning paradigm could be modeled. They suggest the following questions: what mechanism accounts for changes in risk attitudes as a function of mood? Do certain mood states directly modify the curvature of one's subjective utility function, which in turn modifies one's propensity to make risky decisions, or do they instead directly exacerbate asymmetry in learning for unexpected good vs. bad outcomes, which in turn modifies one's risk attitudes?

To test these questions, we designed an experiment containing a variant of the choice task used by Niv (2012). The biggest difference was that we experimentally manipulated mood as subjects engaged in a risk-sensitive RL task. Our research aims can be summarized as follows: (1) to examine mood's effects on the learning of deterministic vs. probabilistic stimuli, (2) to compare distinct computational cognitive models of risk-sensitive learning, and (3) to tease out the mechanism by which mood drives risk preferences within the framework of the best-fitting model.

1.5 Hypotheses and Predictions

(1) We hypothesized mood to distinctly affect the learning and decision-making processes for probabilistic outcomes, and not for deterministic outcomes. Based on the bulk of literature implicating positive mood in higher risk-taking tendencies, we predicted a happy mood induction to increase preferences for risk and a sad mood induction to decrease preferences for risk. (2) Given the computational modeling results by Niv (2012) as well as the results of a small pilot experiment, we hypothesized the asymmetric learning model to best capture

risk-sensitive behavior (relative to other candidate models). (3) Assuming the asymmetric learning model to be the best-fitting model, we hypothesized that mood would affect risk preferences by modulating the learning, rather than the utility, of the probabilistic stimulus (specifically, exacerbating the nonlinear affects of learning to drive distinct risk attitudes).

All life is an experiment. The more experiments you make the better.

Ralph Waldo Emerson

2

Methods

2.1 Experimental Design

2.1.1 Participants

Through Amazon Mechanical Turk (MTurk), 150 individuals (eighty-four male; age, 20-71; mean, 39 years) from across the United States participated in our experiment. MTurk is an online crowdsourcing platform where participants get paid to complete web-based tasks for money [11]. In order to control for the testing environment as much as possible, we imposed a series of measures including audio checks and comprehension checks.

Participants were excluded from analysis if they failed to learn adequately during the choice task (overall proportion correct on deterministic trials was less than a binomial probability of 0.59), failed to respond correctly to at least four out of six comprehension checks, or failed to choose a stimulus on every trial of the task. In total, thirteen participants met the exclusion criteria, leaving 137 subjects for analysis. Prior to this study, an additional 20 participants were tested in a shortened pilot study containing just the choice task. All participants were compensated for their time with a base payment of $4.50 and an additional payment of up to $1.50 based on performance. Study procedures were approved by the Princeton University Institutional Review Board.

2.2 Overview

To explore the interactions between mood and risk preference during learning, we utilized a sequential trial-and-error choice task consisting of stimuli differing in outcome variance. Our task was a modified version of the task used by Niv (2012)[46].

The choice task was framed as an Orchard Harvesting Game in which subjects were motivated (by extra financial gain) to maximize reward by harvesting as many apples as possible. Importantly, subjects were not told the values or outcome probabilities of each orchard, and instead had to learn from experience. Throughout the task, subjects were shown short mood-inducing video clips, and various self-report probes were used to assess mood changes as well as subjective valuations of stimuli. A questionnaire was given to subjects (before or after the choice task, with random assignment) to extract individual-difference measures (state-

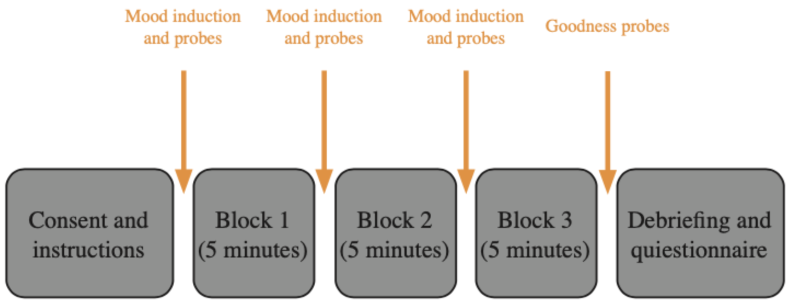

Figure 2.1: Schematic of Experiment. After consent was obtained, the MTurk study started with a display of instructions. The choice task contained three blocks, each preceded by mood-inducing videos. Before and after each mood induction, self-reported measures of mood were collected. After the choice task, goodness ratings of the experimental stimuli were collected. A questionnaire was given to subjects before or after the choice task, with random assignment.

level positive and negative affect and trait-level manic and depressive tendencies). Figure 2.1 shows the overall schematic of the study.

2.3 Risk-Sensitive Choice Task

A binary choice task was presented in which the goal was to harvest as many apples as possible. Subjects were familiarized with the task through on-screen instructions: (1) *On each trial, you will be asked to choose one of two orchards to harvest. Each orchard is represented by a unique flag.* (2) *After you harvest an orchard, you will receive either ZERO apples, ONE apple, or TWO apples.* (3) *Some orchards are better than others, in the sense that they will give you more apples when you harvest them. Your goal is to harvest as many apples as you can, and your bonus payment for this HIT will reflect the number of apples that you harvest.* (4) *Lastly, sometimes you will see only a single orchard available to choose. When this happens, please choose the single orchard. Subjects were required to pass an instructions comprehension check before moving on to the task.*

Since subjects were told in advance that they could receive either zero, one, or two apples, the expected value for any stimulus before experience was assumed to be 1 apple. Thus, in our computational modeling, we set the expected value of each stimulus to 1 since we had provided this a priori knowledge.

There were four possible stimuli (orchards) in the choice task: $S1$, $S2$, $S3$, and $S4$ (Figure 2.2). $S1$, $S2$, and $S3$ were the deterministic stimuli, while $S4$ was the probabilistic stimulus. $S1$ always produced zero apples (m=0, p=100%), $S2$ always produced one apple (m=1, p=100%), $S3$ always produced two apples (m=2, p=100%), and $S4$ produced either zero apples or two apples, by a 50/50 payout (m=2, p=50%). Given that $S2$ and $S4$ predicted rewards of equal mean but different variance, these two stimuli were of most interest. The probabilistic reward scheme was designed to capture risk preference throughout the task with presentations of this particular stimulus pair. The stimulus-value pairings were randomized for each subject to control for possible differences arising from the particular pattern used for the risky and safe stimuli.

Two types of trials were presented, intermixed randomly: free trials and forced trials. In free trials, subjects were shown two stimuli on the screen and were instructed to choose one (by pressing the left or right keyboard arrow). In forced trials, subjects were shown only one stimulus and were instructed to choose it. Following each choice (whether free or forced), the screen displayed the number of apples the chosen stimulus had given. The forced trials were included to make sure subjects adequately sampled each stimulus (so as to prevent choice bias). This was important for computational modeling, since biased sampling could allow for risk aversion to arise implicitly from even the most simple, risk-neutral RL model[47]. Thus, this was an essential way of mitigating potential biases that could result from an interaction between choice and learning. Additionally, given our focus on intersubject differences, forced exploration allowed for more accurate comparisons of risk attitudes, since all subjects would have had comparable exposures to the stimuli.

The task consisted of 156 trials (three blocks of 52 trials each, with mood-inducing videos in between). The task continued sequentially from one block to the next, with no change in any of the experimental conditions (stimuli continued to have the same values and probabilities and no new stimuli were presented). The trials comprised of (1) 36 "risk" trials involving a choice between $S2$ (m=1, p=100%) and $S4$ (m=2, p=50%); (2) 84 "test" trials involving choices containing a stimulus with a higher expected value (correct stimulus); and (3) 36 forced trials involving each of the stimuli. The proportion of times a subject chose $S4$ across all "risk" trials was used as the behavioral measure of risk preference. The proportion of times a subject chose the correct stimulus across all "test" trials was used to assess overall learning (objective performance). Figure 2.2 displays an example sequence of trials and choices.

2.4 Mood Inductions

With random assignment, participants either received sad, neutral, or happy mood inductions. The mood inductions were in the form of short 90-second movie clips; example scenes are shown in Figure 2.3. The neutral and sad videos were taken from a validated film set for the induction of basic emotions[26] and a database of emotion-eliciting films[59]. One of the happy videos (involving an Olympic skater) was taken from a study on positive emotional disturbance[22]. The other two happy videos were validated in other experiments in the Niv Laboratory at Princeton University. Each subject was shown three movie clips total from the same video condition (sad, neutral, or happy). Each block was preceded by one of these

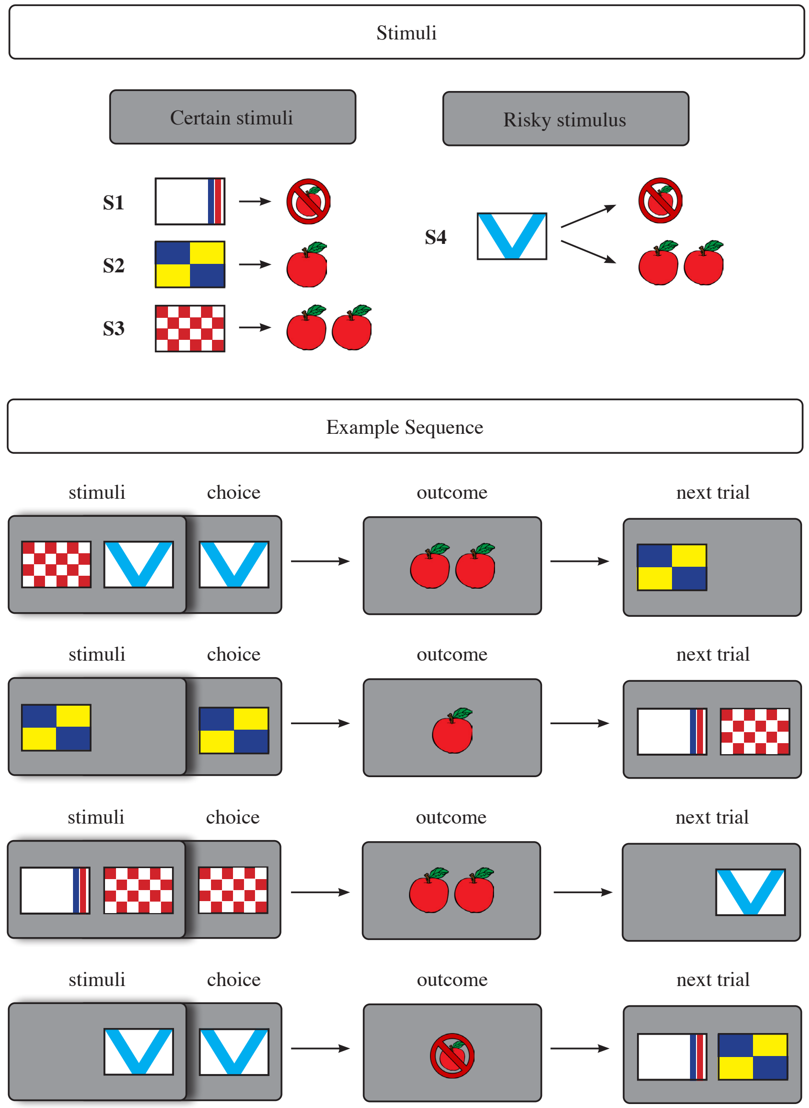

Figure 2.2: Stimuli & Example Sequence of Choice Task. There were four possible stimuli (orchards). $S1$, $S2$, and $S3$ were the deterministic stimuli, while $S4$ was the probabilistic stimulus. $S1$ always produced zero apples, $S2$ always produced one apple, $S3$ always produced two apples, and $S4$ produced either zero apples (50% of the time) or two apples (50% of the time). Subjects were not told this information (they were only told before-hand that they could receive either zero, one, or two apples and that some orchards produce more apples than others). The reinforcement learning choice task was designed to assess the effects of mood on learning and risk-sensitive choice behavior. In each trial, one or two orchards differing in color and design were presented on the left or right side of the screen. Subjects were instructed to choose one of the orchards on each trial (by pressing the left or right keyboard arrow). During forced trials, subjects were instructed to choose the single orchard that appeared on the screen. Each choice was followed by an image indicating the number of apples received. A sample trial/choice sequence is shown.

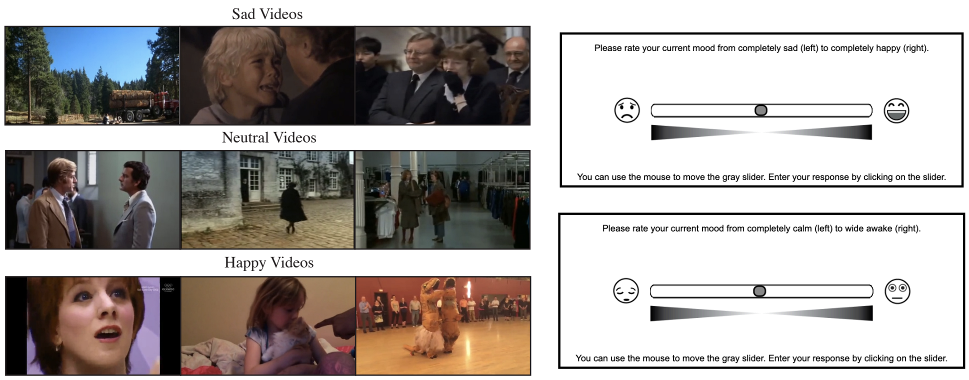

Figure 2.3: Mood-Inducing Videos and Self-Reported Mood Scales. Each subject was shown three movie clips total from the same video condition. The top, middle, and bottom rows display scenes from the sad, neutral, and happy video conditions, respectively. Before and after watching each 90-second video clip, subjects were prompted to report their valence and arousal. Valence and arousal were reported a total of six times each throughout the task, once before and after each of the three mood inductions.

three clips. To test whether subjects had paid attention to the videos, a series of image and comprehension checks followed each clip.

2.5 Self-Reported Measures

2.5.1 Mood Scales

A circumplex model of affect proposes that all affective states arise from cognitive interpretations of neural sensations that are thought to result from two independent neurophysiological systems[53]. Specific emotions are thought to arise out of patterns of activation within these two systems, together with cognitive interpretations. While the valence dimension captures hedonic tone (the degree to which an emotion is pleasant or unpleasant), the arousal dimension describes the degree to which an emotion is associated with high or low energy.

We asked subjects to rate their valence and their arousal on two separate scales, before and after each mood induction, in order to capture both dimensions of affect (Figure 2.3). We collected these self-reports to confirm the efficacy of the mood inductions and to see how they affected ratings of valence and arousal. Given important interactions between valence and arousal in affecting cognition, it was important to probe for both dimensions of mood to examine the videos' effects.

Different subjects may have interpreted levels of the mood scales differently, so to account for individual variability in the use of the mood scales, we z-scored all mood measurements by subject and then averaged the difference between post-induction and pre-induction reports to derive estimates for mood change. This was done for both valence and arousal before adding them to regressions. When calculating correlations, however, we did not z-score the mood ratings, since doing so would not have allowed us to compare across subjects. As a general note, we used Spearman's rank correlations, which are non-parametric and non-assuming of normality; given the non-normal distributions of our data, this method was appropriate.

2.5.2 Goodness Scales

Subjects were asked to provide their hedonic subjective valuations of each stimulus on a continuous scale of goodness (Figure 2.4). We included these probes to compare self-reports with behavioral findings. The difference between a subject's goodness ratings for the risky stimulus ($S4$) and safe stimulus ($S2$) provide a measure of risk preference. We were interested in visualizing the correlation between the self-reported and behavioral measures of risk preference to see the extent to which the conscious subjective report aligned with behavior. A high correlation might indicate greater awareness of one's risk preference and may suggest higher reliability in using self-reported goodness measurements to derive measures of true risk attitudes.

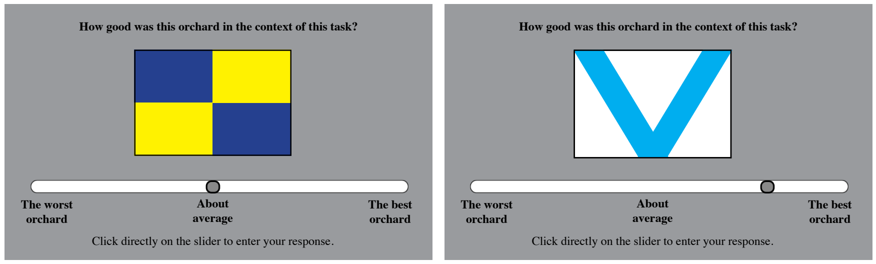

Figure 2.4: Goodness Scales. After the choice task, subjects were asked to rate the goodness of each of the stimuli on a continuous scale ranging from "The worse orchard" to "About average" to "The best orchard". This figure depicts the ratings for our stimuli of most interest. The normalized difference between a subject's goodness ratings for the risky stimulus ($S4$) and safe stimulus ($S2$) provide a measure of risk preference. In the sample ratings above, the hypothetical subject's self-reports indicate a higher preference for risk, since a higher rating was given for the risky stimulus ($S4$).

2.6 Questionnaire

A questionnaire was presented to subjects at the beginning or end of the experiment, with random assignment. Questions were taken from the Positive and Negative Affect Schedule (PANAS)[71] for measures of state-level positive affect and state-level negative affect and the 7 Up 7 Down Inventory[76] for measures of trait-level manic tendency and trait-level depressive tendency. PANAS is comprised of two 10-item mood scales consisting of words that describe different feelings and emotions. Subjects ranked each word on a five-point scale ranging from *Not At All* to *Extremely*. For both positive affect and negative affect, a higher score represents a higher level of affect.[71]. The 7 Up 7 Down inventory is a brief, 14-item version of the General Behavior Inventory that is specifically designed to identify dimensions of mania and depression. Higher scores represent higher tendencies. Mania and depression characterize symptoms of bipolar disorder, and given our study's focus on mood and risky decision-making (both of which are strongly implicated in bipolar disorder), these individual-difference measures were of particular interest.

2.7 Approaches to Computational Modeling

We were interested in finding the best-fitting model of trial-by-trial choices in our risk-sensitive RL task. Based on the findings by Niv (2012), whose study had used a very similar choice task, we decided to narrow our focus on the two risk-sensitive models that were explored in this study: utility model and asymmetric learning model. To compare against a baseline model as well as a more complex model, we compared four models in total: standard RL model, utility model, asymmetric learning model, and super model (containing both utility and asymmetric learning mechanisms).

2.7.1 Standard RL Model

Simple RL[66] offers a general framework for modeling sequential decisions arising from trial and error. Subjects use past experience to learn the values of different stimuli based on discrepancies between expected and actual outcomes. After an outcome (reward) is received (r_t), the expected value of the associated stimulus k is updated (at timepoint $t+1$):

$$V^k_{t+1} = V^k_t + \eta(r_t - V^k_t)$$

The prediction error represents the difference between the actual and expected rewards ($r_t - V^k_t$). The learning rate, η, takes a value between 0 and 1 and captures the extent to which the prediction error updates the expected value of the stimulus. The expected value of each stimulus (for all models) was initialized to 1, since subjects had been told before starting the task that orchards could produce either zero, one, or two apples. The model assumes that subjects use stimulus outcomes to guide their decisions, while occasionally exploring by choosing lower-value options. The softmax choice rule has these properties. It selects stimulus k with probability:

$$p_t^k = \frac{exp(\beta V_t^k)}{\sum_{i=1}^{K} exp(\beta V_t^i)}$$

β is the inverse temperature parameter that controls the level of stochasticity in the model's choices. $\beta = 0$ represents completely random responding while $\beta = \infty$ represents complete deterministic responding to choose the highest value option.

The parameters that define the standard RL model are represented by:

$$\theta_{Standard} = (\eta, \beta)$$

2.7.2 Utility Model

Most studies that have derived risk-sensitive RL methods have focused on various elaborations to the subjective utility framework[63]. The utility model is an elaboration of the simple RL model that maps objective outcomes to subjective utilities through a nonlinear transformation[6]. The exponential utility framework accounts for the fact that different individuals hold differing degrees of risk aversion; the shape of the utility curve is dependent on the parameter α. A concave curve represents risk-averse tendencies (positive α), a convex curve represents risk-seeking tendencies (negative α), and a linear function represents risk neutrality ($\alpha = 0$).

If $\alpha \neq 0$:

$$U(r_t) = \frac{(1 - e^{-\alpha r_t})}{\alpha}$$

The same rule from the standard RL model is then used to update the expected value:

$$V_{t+1}^k = V_t^k + \eta(U(r_t) - V_t^k)$$

The parameters that define the utility model are represented by:

$$\theta_{Utility} = (\eta, \beta, \alpha)$$

This model does not account for variance, at least explicitly. Instead, it is implicitly encoded by subjective utilities. For example, a risk-averse individual is expected to have a higher subjective utility for a safe stimulus relative to a risky stimulus (even when both have the same mean payoff), arising from a preference for certainty at the expense of a possibly better outcome. Thus, this model's non-linear treatment of uncertain outcomes is what allows it to affect the learning of probabilistic stimuli more than deterministic stimuli.

2.7.3 Asymmetric Learning Model

Unlike the standard RL and utility models, the asymmetric learning model explicitly accounts for variance to capture risk sensitivity. As another variant of the simple RL framework, the only difference is that there are two learning rates (η_{pos} and η_{neg}) – one for positive prediction errors and one for negative prediction errors.

$$V_{t+1}^k = V_t^k + \eta_{pos}(r_t - V_t^k)$$

$$V_{t+1}^k = V_t^k + \eta_{neg}(r_t - V_t^k)$$

The parameters that define the asymmetric learning model are represented by:

$$\theta_{Asymmetric} = (\eta_{pos}, \eta_{neg}, \beta)$$

While the learning rate in the standard RL model is indiscriminate of feedback valence, differential treatment for unexpected good and bad outcomes allows the asymmetric learning model to penalize or favor outcome variance. This allows for nonlinear effects on learning for uncertain outcomes. If the learning rate for negative prediction errors is higher than the learning rate for positive prediction errors, the expected value of a risky stimulus would be driven down, and the model would give rise to risk aversion. Similarly, risk-seeking tendencies could result from a relatively higher learning rate for positive prediction errors compared to negative prediction errors. Through asymmetric learning, the model is able to distinctly affect the learning of probabilistic stimuli[46].

As our final model, we tested a combined model called the super model that involves both mechanisms of the utility and asymmetric models. The risk-sensitive parameters of the super model include η^+ and η^- for asymmetric learning, along with α for the subjective utility transformation.

$$\theta_{Super} = (\eta^+, \eta^-, \alpha, \beta)$$

Prior to data collection, we performed simulations with the first three models to generate expected patterns of trial-by-trial choices for hypothetical subjects categorized as risk-seeking and risk-averse. These simulations allowed us to verify that the models, under various parameter conditions, produced qualitatively and quantitatively distinguishable patterns of behavior that were compatible with our research questions. Afterwards, we conducted a pilot experiment with 20 subjects containing just the choice task. Model fitting and parameter estimation using the pilot data allowed us to verify the robustness of our experimental set-up. All subjects were able to adequately learn the values of the deterministic stimuli, and our choice task allowed us to distinguish between models under optimal parameter conditions. The results also provided potential evidence supporting the asymmetric learning model[40].

2.7.4 Maximum-Likelihood Approach to Model Fitting

In the maximum likelihood approach to model fitting, the goal is to find the parameter values of the model ($\hat{\theta}_m^{MLE}$) that maximize the likelihood of the observed data.

We optimized model parameters by maximizing the trial-by-trial log likelihoods of the data through the entire task (156 trials), given different settings of the model parameters. This is equivalent to minimizing the negative log likelihood (the loss) of the data. We used the Python SciPy[70] function optimize.minimize with method='TNC'. This function performs bound-constrained minimization[72] that uses a truncated Newton algorithm[44] to minimize a function with variables subject to bounds. We also fit separate models to each block (52 trials). The model-fitting procedure was the same, except the model was fit to 52 trials instead of 156.

$$LL = \sum_{t=1}^{T} logp(choice_t|\theta_m)$$

where p is the model's probability of the specific $choice_t$ given the parameters of the model and the information up to that choice.

The following represents the optimal parameters for a specific model:

$$\hat{\theta}_m^{MLE} = \theta_m \arg\max LL$$

Model likelihoods were based on assigning probabilities to the 156 trials (or 52 trials) for each subject, according to the softmax function. Given our focus on intersubject differences, we fit the parameters for each subject individually rather than pooling data across subjects[46].

Given the multiplicative interaction between learning rates and softmax inverse temperatures in the model[35], we imposed parameter constraints to produce realistic parameter fits. We constrained learning rates to the range $0 \leq \eta \leq 1$. The inverse temperature softmax parameter was constrained to the range $0 \leq \beta \leq 40$. The utility parameter was constrained to the range $-5 \leq \alpha \leq 5$. Initial guesses for optimization were determined from the results of an extensive brute optimization procedure. The constraints were determined after successive visualizations of parameter distributions using more or less bounded ranges. All parameters had uniform priors.

2.7.5 Model Performance and Model Comparison

We took the mean of the total log likelihoods produced from the optimized parameters across participants and applied the Bayesian Information Criterion (BIC) to transform these mean log likelihoods into mean BICs. Overfitting is a common problem that results from increasing the number of parameters[75]. BIC attempts to mitigate this problem by penalizing more complex models (model complexity is reflected by its number of parameters). The model with the lowest BIC represents the best-performing model.

$$BIC = -2 * LL + k * log(n)$$

where LL = the maximized log likelihood for the estimated model, n = the number of trials, and k = the number of free parameters to be estimated. For model-fitting on the entire task,

$n = 156$, and for model-fitting on specific blocks, $n = 52$. The super model, which has four parameters (compared to three in the utility and asymmetric learning models), was included in model comparison to assess trade-offs between model performance and complexity.

3

Behavioral Results

Prior to starting our main behavioral analyses, we first performed a manipulation check to make sure that our mood inductions had "worked" by testing for their effects on self-reported mood. We specifically wanted to confirm that self-reported valence had changed significantly after the happy and sad inductions and that all three videos in each video condition were comparable in their mood-inducing effects.

3.1 Effect of Mood Inductions on Self-Reported Mood

Figure 3.1 shows average self-reported mood ratings over time, showing that the mood inductions had driven self-reported mood in the expected directions after each induction.

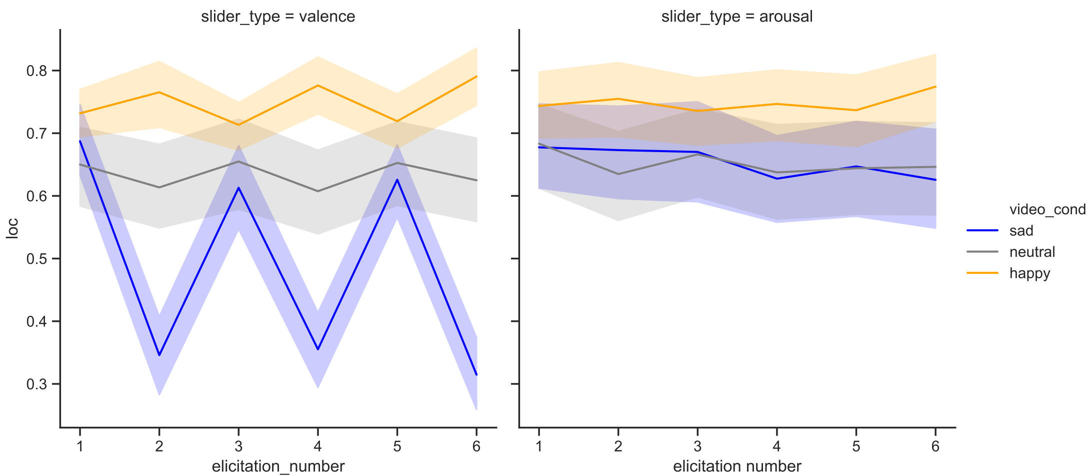

Figure 3.1: Mood Inductions vs. Mood Reports. Elicitation numbers 1, 3, and 5 are pre-mood induction timepoints, while elicitation numbers 2, 4, and 6 are post-mood induction timepoints. Based on self-reported measures of valence and arousal, the mood inductions caused significant differences in these measures as a function of the video condition.

We used generalized linear models to test the effect of the video condition on self-reported mood. We also tested for the effect of the video condition in interaction with each of the individual-difference measures collected from the questionnaire. We built separate regression models (instead of a single model) for each individual-difference measure in interaction with the video condition because we found significant correlations among the individual-difference measures.

Specifically, we found positive correlations between positive affect and 7 up ($\rho(135) = 0.33, p < 0.001$) and between negative affect and 7 down ($\rho(135) = 0.46, p < 0.001$), and negative correlations between positive affect and 7 down ($\rho(135) = -0.41, p < 0.001$) and between positive affect and negative affect ($\rho(135) = -0.28, p < 0.001$).

3.1.1 Effect of Video Condition on Valence Change

According to the results of a generalized linear model, the sad induction produced a significant negative effect on valence change relative to the happy induction ($\beta = -2.24, p < 0.001$) and the neutral induction produced a significant negative effect on valence change relative to the happy induction ($\beta = -0.94, p < 0.001$).

To test for possible interactions between the video condition and individual-difference measures on valence change, a numeric encoding of video condition (-1 for sad, 0 for neutral, 1 for happy) was used to derive a single interaction coefficient for each individual-difference measure interaction. No significant interactions were found (all $p > 0.05$). There were also no significant correlations between individual-difference measures and valence change (all $p > 0.05$).

3.1.2 Effect of Video Condition on Arousal Change

According to the results of a generalized linear model, the sad induction produced a significant negative effect on arousal change relative to the happy induction ($\beta = -0.55, p < 0.01$) and the neutral induction produced a significant negative effect on arousal change relative to the happy induction ($\beta = -0.46, p < 0.05$).

A significant negative interaction between the sad induction and positive affect relative to the interaction between the happy induction and positive effect was found ($\beta = -0.045, p < 0.05$). For sad-induced subjects, a higher positive affect measure predicted a larger-magnitude drop in arousal following a mood induction. None of the other interactions with individual-difference measures were significant (all $p > 0.05$). There were no significant correlations between individual-difference measures and arousal change (all $p > 0.05$).

3.1.3 Correlation between Valence and Arousal

We visualized the correlation between valence change and arousal change (Figure 3.2) and found these two measures to be positively correlated overall ($\rho(135) = 0.27, p < 0.05$). However, the only significant correlation existed for the neutral condition ($\rho(46) = 0.43, p < 0.05$). The correlations for the happy and sad video conditions were not significant (both $p > 0.05$). The Fisher Z-transformation confirmed all video condition-specific correlations to be significantly different.

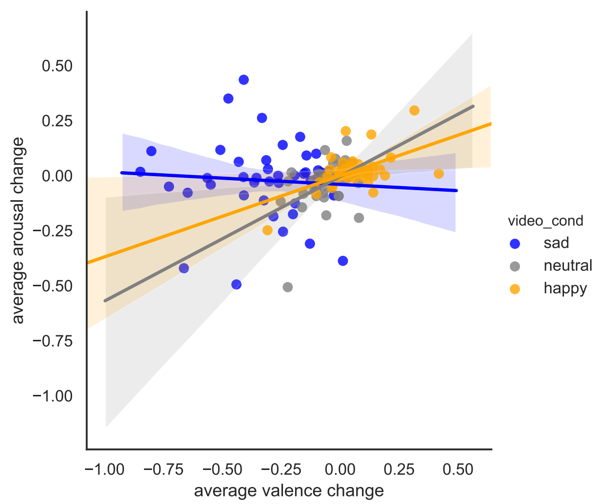

Figure 3.2: Valence Change vs. Arousal Change. This figure depicts the relationship between the average valence change and the average arousal change. We found these two dimensions of mood to be positively correlated overall ($\rho(135) = 0.27, p < 0.05$), the only significant correlation existed for the neutral condition ($\rho(46) = 0.43, p < 0.05$). The correlations for the happy and sad video conditions were not significant (both $p > 0.05$). The Fisher Z-transformation confirmed all video condition-specific correlations to be significantly different.

3.1.4 Comparison of Videos within Video Conditions

We tested for the comparability of the videos for each video condition in their mood-inducing effects according to self-reported measures (Figure 3.3). No statistically significant differences between group means of valence change or arousal change were found by one-way ANOVA for any of the video conditions (all $p > 0.05$).

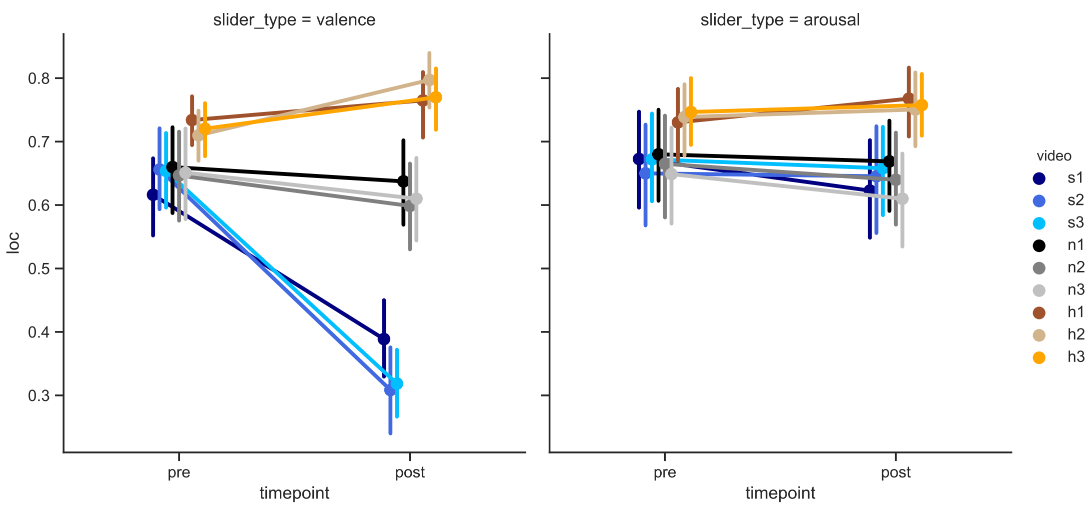

Figure 3.3: Video Comparability within Video Conditions. No statistically significant differences between group means of valence change or arousal change were found by one-way ANOVA for any of the video conditions (all $p > 0.05$).

3.1.5 Interim Summary

A significant effect of the video condition on valence change confirmed the effectiveness of our mood inductions in inducing the desired valence (at least according to the self-reports). However, a significant effect of the video condition on arousal change presents a possible limitation. Since the happy induction (relative to the sad induction) produced increases in both valence and arousal, we cannot be certain that possible effects on risk preference were independent of changes in arousal. The finding that valence and arousal reports were only significantly correlated for subjects in the neutral condition makes it more likely that the sad and happy inductions had affected both the valence and arousal dimensions of mood. It is still notable, however, that the video condition had a more significant effect on valence change. Additionally, for sad-induced subjects, a higher positive affect measure predicted a larger-magnitude drop in arousal change, suggesting this particular state-level measure to be an important mediator in our results. No significant differences were found for the videos in each condition, confirming the comparability of the videos. This was a necessary check, as it mitigates the possibility that certain videos had elicited conflicting moods.

3.2 Effect of Mood on Overall Learning

Through trial and error, subjects were able to learn the relative values of the deterministic stimuli ($S1$, $S2$, $S3$). Overall learning was quantified by the proportion of "test" trials (84 trials) to which subjects had correctly responded. Forced trials and "risk" trials were excluded from this analysis. The learning curve in Figure 3.4 depicts averaged performance (coded either 0 or 1 for objective correctness of choice) of all subjects per trial. The red dashed lines separate the three blocks. The curve shows that most learning had taken place in block 1.

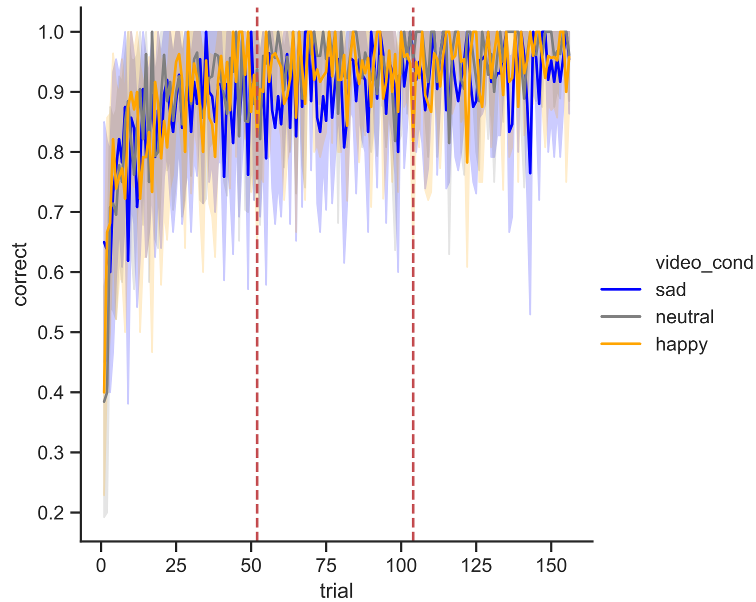

Figure 3.4: Overall Learning Curve. On the y-axis, *correct* is a binary variable that indicates whether or not the subject chose the stimulus of higher expected value. It does not reflect the received reward. Over time, average performance improved as subjects learned the values of stimuli through sequential decision-making. The vertical dotted lines separate the experiment into the three blocks. The curve shows that most learning had taken place in block 1.

To more specifically assess the effect of mood on overall learning, we fit a generalized linear mixed effects model (Table 3.1) on a binary variable that encodes the correctness of the choice. We used a categorical encoding of the video condition as a predictor. Generalized linear mixed effects models incorporate both fixed-effects and random-effects parameters in a linear predictor via maximum likelihood. Our model was fit with random intercepts by subject ID. Block was included as a fixed-effects variable to test how the effect of the video condition may have varied by block.

Results showed block to have a significant positive effect on overall learning ($p < 0.01$). The mood inductions did not have a significant effect (all $p > 0.05$).

	Dependent variable:
	correct
block	1.183***
	(0.153)
video_condneutral	0.012
	(0.300)
video_condsad	0.152
	(0.300)
block:video_condneutral	−0.002
	(0.198)
block:video_condsad	−0.200
	(0.196)
Constant	0.839***
	(0.227)
Observations	11,358
Log Likelihood	−2,861.551
Akaike Inf. Crit.	5,741.102
Bayesian Inf. Crit.	5,807.141
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 3.1: Effect of Video Condition on Overall Learning. The results of a generalized linear mixed effects model showed block to have a significant positive effect on overall learning ($p < 0.01$). The mood inductions did not have a significant effect (all $p > 0.05$)

.

3.2.1 Interim Summary

The only significant predictor was block, which is in line with the expected upward trajectory of performance through sequential trial-and-error learning. Additionally, the learning curve shows performance to improve most substantially in block 1. This finding aligns with expected patterns of model-free learning, where the learning rate decays over iterations[46]. Our main finding that the mood inductions had not affected overall learning supports our hypothesis for Aim 1 that mood would not bias the learning of deterministic stimuli.

3.3 Effect of Mood on Risk-Sensitive Learning

"Mood" in our study could be represented in several ways. As predictors, we used the video condition (categorical encoding of mood) and self-reported mood (continuous encoding of mood) in separate generalized linear mixed effects models to test for their effects on behavioral risk attitudes. Given the significant positive correlation we found between valence change and arousal change (at least for the neutral condition), we tested for the effects of self-reported valence change and arousal change as predictors in separate models. Risk preference was quantified by the proportion of "risk" trials (36 trials) in which subjects chose the probabilistic stimulus ($S4$) over the deterministic stimulus ($S2$).

Prior to starting our main analyses, we found no significant correlations between risk preference and any of the individual-difference measures (all $p > 0.05$).

3.3.1 Video Condition as a Predictor

According to the results (Table 3.2, the sad induction predicted a lower preference for risk ($p < 0.05$). The block also had a significant negative effect on risk preference ($p < 0.01$), predicting subjects to become increasingly risk-averse with each block (Figure 3.5). The slope of this decline did not differ based on the video condition, as the interaction effect (derived using a numerical encoding of video condition) was not significant ($p > 0.05$).

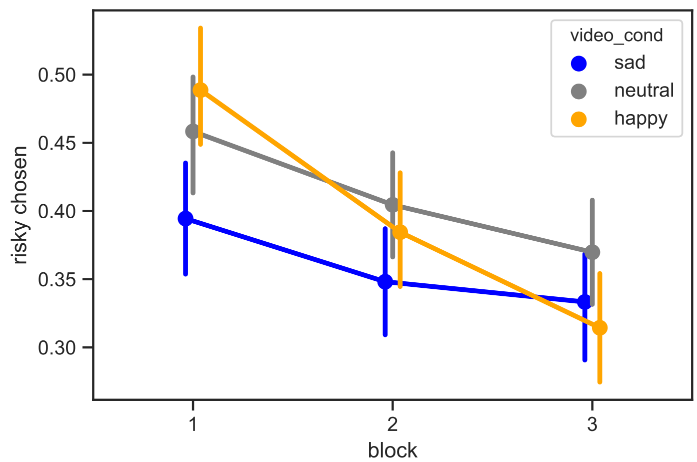

Figure 3.5: Block vs. Risk Preference. On the y-axis, risky chosen is a binary variable that indicates whether or not the subject chose the risky stimulus ($S4$) over the safe stimulus ($S2$). On average, subjects were generally risk-averse regardless of video condition (they were risky less than 50% of the time). According to the results of a generalized linear mixed effects model, going from one block to the next decreased the probability of choosing the risky stimulus by an estimated 59% ($p < 0.01$). The error bars in the figure should be interpreted with caution, since they only represent summary statistics of the data. The mixed-effects regression showed that this downward trend did not vary by video condition, as the interaction term was found to have no significant effect ($\beta = -0.16, p > 0.05$). Thus, the three slopes depicted in this figure were not significantly different.

	Dependent variable:
	risky_chosen
block	-0.591^{***}
	(0.166)
video_condneutral	-0.336
	(0.380)
video_condsad	-0.847^{**}
	(0.389)
block:video_condneutral	0.204
	(0.230)
block:video_condsad	0.325
	(0.234)
Constant	0.545^{**}
	(0.275)
Observations	4,932
Log Likelihood	$-2,714.015$
Akaike Inf. Crit.	5,446.030
Bayesian Inf. Crit.	5,504.562
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 3.2: Effect of Video Condition on Risk Preference. The results of a generalized linear mixed effects model showed a sad mood induction to predict a lower preference for risk relative to a happy mood induction ($p < 0.05$). The block also had a significant negative effect on risk preference ($p < 0.01$), predicting subjects to become increasingly risk-averse with each block. The slope of this decline did not differ based on the video condition, as the interaction effect (derived using a numerical encoding of video condition) was not significant ($p > 0.05$).

3.3.2 Self-Reported Mood as a Predictor

Self-reported valence change and self-reported arousal change (Models 1 and 2 in Table 3.3) were not predictive of risk preference (both $p > 0.05$). The block was again shown to have a significant negative effect ($p < 0.01$). The interaction between self-reported valence change and block was not significant ($p > 0.05$), but the interaction between self-reported arousal change and block was significantly positive ($\beta = 0.22, p < 0.05$). This interaction is depicted in Figure 3.6. Higher drops in arousal were associated with faster declines in risk preference through blocks.

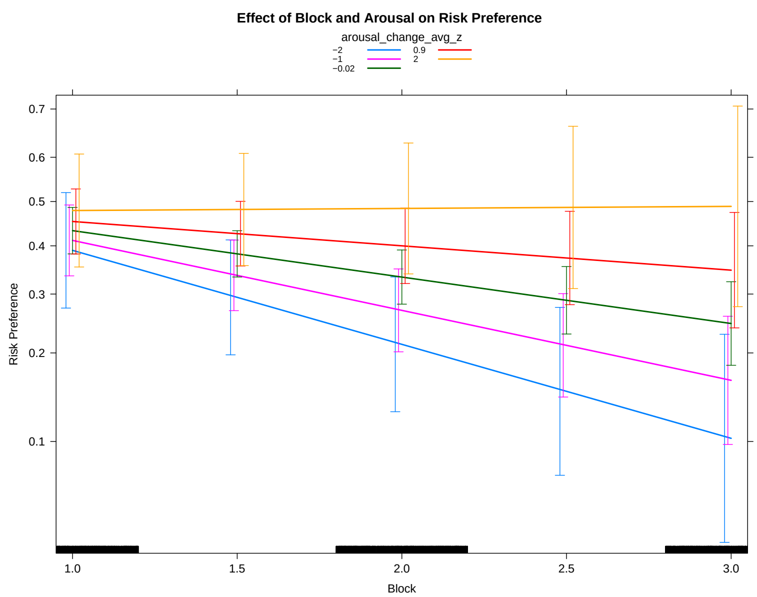

Figure 3.6: Interaction between Block and Arousal Change on Risk Preference. The results of a generalized linear mixed effects model showed a significantly positive interaction effect between the block and self-reported arousal change ($\beta = 0.22, p < 0.05$). More negative drops in arousal were associated with faster declines in risk preference through blocks.

	Dependent variable:	
	risky_chosen	
	(1)	(2)
block	-0.430^{***}	-0.421^{***}
	(0.100)	(0.094)
valence_change_avg_z	0.193	
	(0.129)	
block:valence_change_avg_z	-0.044	
	(0.078)	
arousal_change_avg_z		-0.129
		(0.178)
block:arousal_change_avg_z		0.220^{**}
		(0.105)
Constant	0.223	0.156
	(0.164)	(0.159)
Observations	4,932	4,932
Log Likelihood	$-2,715.185$	$-2,713.727$
Akaike Inf. Crit.	5,444.371	5,441.453
Bayesian Inf. Crit.	5,489.895	5,486.978
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 3.3: Effect of Self-Reported Mood Change on Risk Preference. The results of a generalized linear mixed effects model showed that self-reported valence change and self-reported arousal change (Models 1 and 2) were not predictive of risk preference (both $p > 0.05$). The block had a significant negative effect ($p < 0.05$). The interaction between self-reported valence change and block was not significant ($p < 0.05$), but the interaction between self-reported arousal change and block was significantly positive ($\beta = 0.22, p < 0.05$). More negative drops in arousal were associated with faster declines in risk preference through blocks.

Correlations within each Video Condition

We calculated correlations between self-reported mood change (both valence and arousal) on behavioral risk preference. For the sad video condition, the valence change was found to be negatively correlated with risk preference ($\rho(43) = -0.31, p < 0.05$). All other video condition-specific correlations were not significant ($p > 0.05$), and Fisher Z-transformations showed no significant differences between video condition-specific correlations. No significant correlations were found between arousal change and risk preference (all $p > 0.05$).

3.3.3 Interim Summary

As hypothesized for Aim 1, the video condition had a significant effect on risk preference, with a sad induction making subjects more risk-averse, relative to a happy induction.

Interestingly, subjects generally become more risk-averse across blocks, regardless of the video condition. Over time, subjects gained a better idea of the probabilistic stimulus through more exposures. Most learning took place in block 1, where subjects were expected to have chosen the risky stimulus more frequently in order to figure out its variance in outcome. Naturally, after accumulating more knowledge, subjects overall could have chosen the risky stimulus less frequently, resulting in what looks like greater risk aversion over time. At the present moment, however, the specific effect of block is unclear; for example, perhaps watching more videos (thus, their cumulative effects regardless of video condition), modulated another aspect of cognition that left subjects generally less likely to choose the risky stimulus (discussion continued in Chapter 4 and Chapter 5).

Another interesting finding was that self-reported mood changes were not predictive of risk preferences, while the video condition was. This result suggests a possible disconnect between conscious reflections of mood and truly experienced mood. More broadly, it suggests that the simple act of watching a mood-inducing video has predictable effects on behavior (at least in our experimental context). At the same time, our results are less interpretable. Are the mood-reports unreliable? How can we be sure that our mood inductions had truly affected mood? Could they have affected another aspect of cognition to influence distinct risk attitudes? The differential effects by block on risk preference based on self-reported arousal further point to the need to clarify these questions (discussion continued in Chapter 5).

3.4 Self-Reported Risk Preference vs. Behavioral Risk Preference

The "goodness" difference is the normalized difference between self-reported goodness ratings for the risky ($S4$) and safe ($S2$) stimuli. It is essentially a measure of hedonic preference. There was a significant positive correlation between the goodness difference and behavioral risk preference ($\rho(135) = 0.56, p < 0.001$), shown in Figure 3.7. Fisher Z-transformations showed no significant differences between video condition-specific correlations.

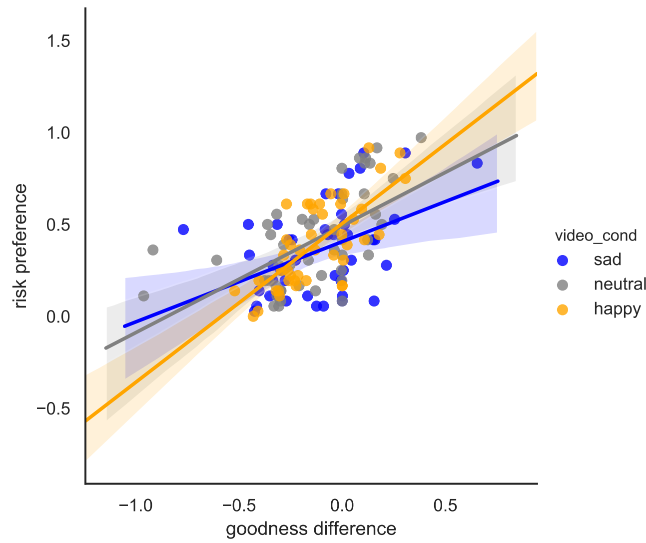

Figure 3.7: Goodness Rating Difference vs. Risk Preference. The "goodness" difference is the normalized difference between self-reported goodness ratings for the risky and safe stimuli. Risk preference is a behavioral measure that is quantified by the proportion of times the risky stimulus was chosen when presented with the safe stimulus. A significant positive correlation between these two variables was found ($\rho(135) = 0.56, p < 0.001$). Fisher Z-transformations showed no significant differences between video condition-specific correlations.

3.4.1 Interim Summary

A significant positive correlation suggests the goodness difference to be a reliable indicator of behavioral risk preference. Since the self-reported goodness ratings were collected at the end of the task, they represent reflections – a result of looking back on the task and assessing how subjectively "good" the stimuli were. These ratings provide a glimpse into subjective valuations and potentially inform us of the extent to which subjects may have been consciously aware of their preferences for risk.

4

Computational Modeling Results

The study by Niv (2012) proposed the asymmetric learning model, among the candidate models, as the best-fitting model of risk-sensitive learning. Since we utilized a variant of the task used in this study, we expected to find similar results. However, while we did utilize a very similar choice task, there was one major difference: we experimentally manipulated mood, which then drove distinct risk preferences. Thus, given our novel experimental framework, it was still necessary to compare other candidate models through model-fitting on our behavioral results. Specifically, we tested to see (1) if the asymmetric learning model would best capture risk-sensitive behavior (relative to other candidate models) and (2) if the asymmetric learning model would also account for the observed mood-driven differences in risk attitudes.

In total, we performed model-fitting and model comparison across four models: standard RL model, utility model, asymmetric learing model, and super model. The standard RL model was not expected to capture risk attitudes because it has no "risk-sensitive" parameter that would allow for any sort of non-linearity (in outcome valuation or in learning). The utility model, which does not explicitly track variance, was expected to suffer from drawbacks in our model-free framework (where subjects were given minimal information and had to learn through trial and error). The asymmetric learning model, by virtue of explicit tracking of variance, was predicted to perform the best. The super model was included in model comparison to assess trade-offs between model performance and complexity.

4.1 Results of Model-Fitting

4.1.1 Capturing the General Behavioral Pattern

After model-fitting on the entire task, the asymmetric learning model had the lowest mean BIC (86.02), followed by the utility model (89.91), the super model (90.14), and the standard RL model (92.19). A lower BIC indicates a better fit of the data, so the asymmetric learning model performed the best (Table 4.1).

The difference in mean BICs between the asymmetric and utility models was -3.88 and the standard deviation of the difference was 9.06. The difference in mean BICs between the asymmetric and standard Models was -6.17 and the standard deviation of the difference was

13.95. The difference in mean BICs between the asymmetric and super models was -4.12 and the standard deviation of the difference was 5.47.

Given that we saw a significant effect of block in driving risk preference (with subjects on average becoming more risk-averse through time), we fit the models separately to each block. In block 1, the standard model performed the best, followed by the asymmetric learning, utility, and super models. In block 2, the asymmetric learning model performed the best, followed by the standard, utility, and super models. In block 3, the asymmetric learning model performed the best, followed by the utility, standard, and super models.

Model	Mean BIC (All)	Block 1	Block 2	Block 3
Standard	92.19	37.26	36.43	36.57
Asymmetric	86.02	37.64	35.40	33.97
Utility	89.91	38.88	38.09	36.32
Super	90.14	40.67	38.17	36.14

Table 4.1: BIC Approximations. Overall, the asymmetric learning model had the lowest mean BIC, signaling the best fit on the data. In block 1, the standard model performed the best, and in blocks 2 and 3, the asymmetric learning model performed the best.

The asymmetric learning model emerged as the best-fitting model in capturing the general behavioral pattern, and the next step was to see if this mechanism was invariant of the video condition. For example, a positive mood could have more substantially modulated the learning of uncertain stimuli (suggesting the asymmetric learning framework), while a negative mood could have more substantially modulated the utility of uncertain stimuli (suggesting the utility framework). Essentially, the mechanism driving risky decision-making may have been dependent on the video condition. To address this possibility, we compared model fits to each video condition group 4.1.

Across all video conditions, the asymmetric learning model was found perform the best.

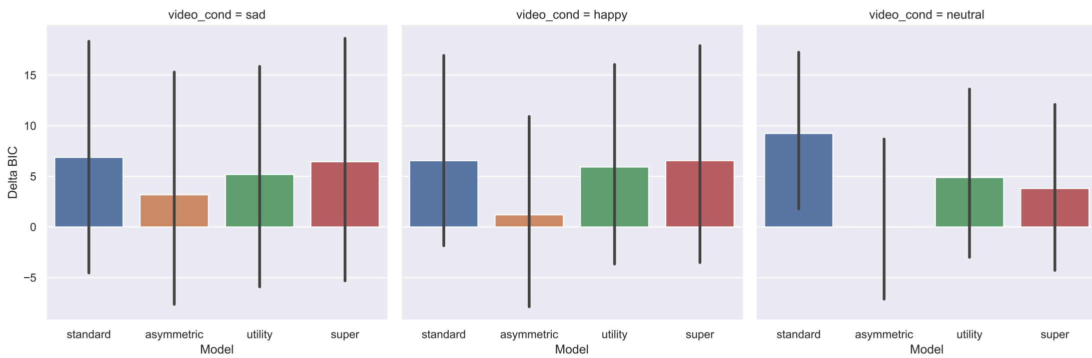

Figure 4.1: Model Comparisons Across Video Conditions. The Bayesian Information Criterion (BIC) was used as a metric for model comparison. The model with the lowest BIC is preferred. Previous results showed the asymmetric learning model for subjects in the neutral condition to produce the best fit of the data. Thus, we subtracted this mean BIC value from the others to display the differences in mean BICs relative to this best-fitting model.

The relative model rankings were also consistent, except for in the neutral video condition. Here, while the asymmetric learning model still performed the best, the utility model performed worse than the super model (for the other video conditions, the utility model performed better than the super model). When comparing the asymmetric learning model across video conditions, its performance was best in the neutral condition.

4.1.2 Interim Summary

Our computational modeling results support our hypothesis (Aim 2) for the asymmetric learning model in best capturing risk-sensitive learning in the task. As expected, the utility model came in second place. The super model performed less well than the asymmetric learning and utility models (while it performed the best when comparing mean likelihoods, it came in third place when comparing mean BICs). Despite the penalty for model complexity, it is notable that the super model still performed better than the standard RL model (which only had two parameters compared to four). This result confirmed that a simple RL framework cannot adequately capture risk attitudes.

The results of model-fitting to each block showed interesting patterns. In block 1, the standard RL model performed the best. Since this was when subjects were doing most of the learning, it may have been too early for distinct risk preferences to emerge. Thus, risk-sensitive models may not have been able to gauge differences in risk attitudes this early on across subjects during a time when they were trying to learn not only the values of the deterministic stimuli, but also the variance of the probabilistic stimulus. Across blocks starting from the first to the third, all models besides the standard RL model produced progressively better fits of the data. Risk-sensitive models are expected to perform well under conditions in which subjects have fully learned the outcomes of the deterministic stimuli and have developed more discernible risk preferences. As expected, the asymmetric learning model did indeed perform the best in block 2, but interestingly, the utility model still performed worse than the standard RL model. It was only in block 3 that the utility model moved up to second place. These results potentially provide further support for the asymmetric learning model, which was able to more sensitively capture risk preferences at an earlier stage of trial-and-error learning. Our results showing the utility model performing worse than the standard RL model in blocks 1 and 2 support findings that while the subjective utility framework has seen useful applications in modeling risk attitudes, it performs poorly as a model-free RL algorithm[43,7]. In consideration of the significant effect of block that we found on risk preferences, our results provide potential support for the idea that this effect could have been a function of learning. The fact that our risk-sensitive learning models had the best fits on block 3 trials suggests that distinct risk attitudes are more discernible after a certain amount of learning had taken place. At first glance, subjects seem to have become more risk-averse through time, but this trend may have resulted because true risk attitudes simply emerged at a later time upon more exposure to the deterministic and probabilistic stimuli.

The video condition-specific results of model-fitting showed the asymmetric learning model to consistently perform the best across all video conditions. This provides support for the asymmetric learning framework as a generally robust model of risk-sensitive learning. Whether or not subjects' risk preferences had been modulated by experimentally induced mood, the asymmetric learning model produced the best fits. Another interesting finding was that this model performed the best when fitted to subjects in the neutral condition. This possibly suggests the need for a more complex model in which mood is a modulating factor, as the results show poorer performance in cases where risk-preference was driven by mood inductions (sad and happy). Nevertheless, the synthesis of our results above provide strong evidence in support of the asymmetric learning model. The next step was to then assess if this model could also explain how distinct mood-driven risk preferences emerged.

4.2 Capturing Mood-Driven Effects on Risk Preference

Aim 3 of the study was to investigate the role of mood in modulating risk preference in the context of our best-fitting model. Given the general superiority of the asymmetric learning model in capturing behavioral patterns across video conditions and blocks, we focused our attention on the fitted parameters of the asymmetric learning model.

This model has three parameters: η^+, η^-, and β. The two learning rates are our risk-encoding parameters. A higher positive difference between η^+ and η^- would lead to higher risk-seeking tendencies. Prior to analyses, we normalized the difference in learning rates (*eta_diff* $= \frac{\eta^+ - \eta^-}{\eta^+ + \eta^-}$) to produce a comparable metric across participants.

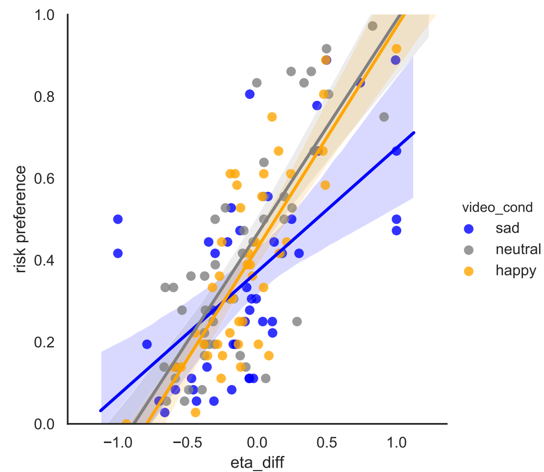

Figure 4.2: Difference in Asymmetric Etas vs Risk Preference. eta_diff represents the normalized difference between η^+ and η^-. The correlation between *eta_diff* and risk preference was found to be significant ($\rho(135) = 0.70, p < 0.001$), with the highest correlation for the neutral condition ($\rho(46) = 0.77, p < 0.001$). Fisher Z-transformations showed no significant differences between video condition-specific correlations.

As a sanity check, we first visualized the correlation between the normalized difference in learning rates (*eta_diff*) and behavioral risk preference to make sure the model was working as expected, with a more positive difference in learning rates correlated with a higher risk preference. The correlation between *eta_diff* and risk preference was found to be significant ($\rho(135) = 0.70, p < 0.001$), with the highest correlation for the neutral condition ($\rho(46) = 0.77, p < 0.001$). Fisher Z-transformations showed no significant differences between video condition-specific correlations, but it is notable that this result aligns with the model's best fit on data for the neutral condition.

Before performing statistical tests to determine possible video condition-specific differences in parameters (η^+, η^-, and β), we visualized their distributions and found the three parameters to be skewed. We used the Kruskal-Wallis test, a non-parametric method for testing whether samples originate from the same distribution, instead of ANOVA to test

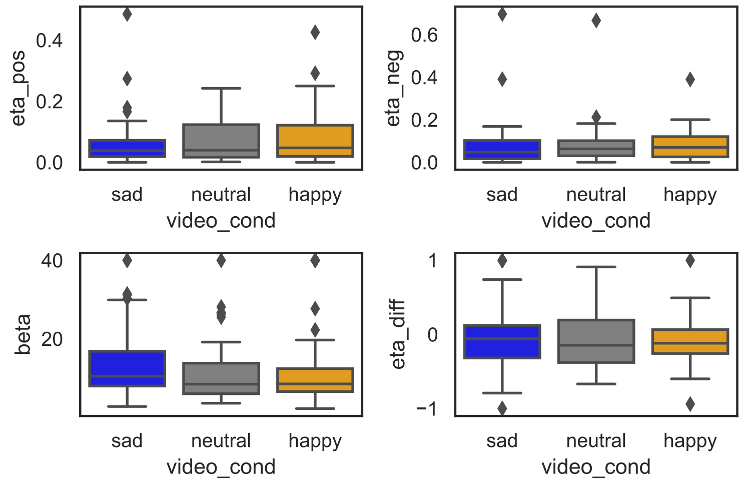

Figure 4.3: Parameter Distributions for Asymmetric Learning Model. η^+, η^-, and β are the free parameters of our best-fitting model, the asymmetric learning model. They represent the learning rate from positive prediction errors, the learning rate from negative prediction errors, and the softmax parameter, respectively. *eta_diff* is the normalized difference between η^+ and η^-. We found no significant differences between group means for any of the three free parameters as determined by the Kruskal-Wallis test ($p > 0.05$). Most importantly, we found no significant differences between group means for *eta_diff* as determined by one-way ANOVA ($F(2, 133) = 0.22, p > 0.05$).

for video condition-related differences in these parameters. No significant differences were found for any of these parameters (all $p > 0.05$). Most importantly, there were no significant differences between group means for *eta_diff* as determined by one-way ANOVA ($F(2, 133) = 0.22, p > 0.05$). When *eta_diff* was correlated with individual-difference measures, we found no significant correlations across any of these measures (all $p > 0.0.5$).

4.2.1 Interim Summary

We found no significant video condition-specific differences in the asymmetric learning model's risk-sensitive parameters. Thus, while the model was the best at capturing general behavioral patterns for risk-sensitive learning, we could not pinpoint the mechanism by which the mood inductions had led to distinct risk preferences. Had we found significant group-level differences in *eta_diff* based on the video condition, we may have been more certain in attributing mood's effect on risk preference to the asymmetric learning process. At the present moment, however, we can neither confirm nor deny this hypothesis.

Our finding that the asymmetric learning model best captures risky decision-making for neutral-induced subjects possibly offers a hint as to why we could not derive significant mood-related differences directly from the model's risk-sensitive parameters. Our results possibly suggest that a model with mood directly as a parameter in the asymmetric learning framework may offer a more nuanced and accurate way of capturing mood-driven risk preferences.

5

General Discussion

How does mood influence one's preference for risk through experiential learning?

To explore this question, we tested subjects using a risk-sensitive reinforcement learning task containing experimental mood manipulations. A happy mood induction caused higher risk-taking tendencies, while a sad mood induction caused lower risk-taking tendencies. The overall learning of deterministic stimuli was left untouched by mood. The results of computational modeling suggested the asymmetric learning model to best capture trial-by-trial choices through its nonlinear effects on the learning of probabilistic stimuli. Whether this model can explain the mood-driven effects on risk preference is unclear; our results point to future directions aimed at clarifying how exactly mood modulates risk-sensitive computations.

5.1 Evidence for Mood-Driven Risk-Sensitive Learning

Our behavioral and modeling results provide support for our hypothesis that mood would distinctly bias the learning of risky stimuli (with no effects on overall learning). Additionally, our hypothesis that the asymmetric learning model would be the best-fitting model of behavioral patterns was supported through model selection. Our final hypothesis that mood would modulate the risk-sensitive parameters of the asymmetric learning model to bias risk preference remains an open possibility. At the present moment, it is simply unclear; whether due to low statistical power or other possibilities, we found no significant differences in the asymmetric learning process as a function of video condition. Nevertheless, our results provide strong evidence for a mood-biased risk-sensitive RL process and suggest insights for a more nuanced computational model.

5.2 Implications for Theories of Mood and Decision-Making

Our findings are consistent with theoretical and experimental findings implicating positive mood with higher risk-taking tendencies and negative mood with lower risk-taking tendencies[27]. This includes the Affect Infusion Model[18] and a body of findings, particularly for pathological mood states, showing mania to result in more risk-seeking behavior[50,24] and depression to result in more risk-averse behavior[77]. Thus, our results are inconsistent with the Mood Maintenance Hypothesis[31], which would have predicted happy-induced subjects to be less willing than sad-induced subjects to choose the risky stimulus, as well as experimental findings suggesting associations between negative affect and risk-seeking attitudes[55], and between positive affect and risk-averse attitudes[48,33].

At the same time, it is difficult to draw direct comparisons because risk-sensitivity is highly context-dependent across time and space. There is the fundamental question of whether experimental risk preferences reflect real-world behavior[69]. Additionally, people can be risk-seeking in one domain and risk-averse in others (like health and financial decisions)[61]. Our own results have also shown that interactions between state-level and trait-level effects further complicate the relationship between mood and risk attitudes. Perhaps most importantly in light of our study, there are distinct cognitive differences when information about risk is explicitly provided compared to when this information must be learned experientially[46].

Thus, we frame our discussion of risky decision-making specifically around the context of trial-and-error learning, where mood inductions were found to have distinct effects on risk attitudes. Now in a more comparable context, we can more robustly claim consistency in results with those found by Niv (2012)[46]. Our results of model comparison selected the asymmetric learning model as the most likely cognitive representation driving risk-sensitive learning; this validates the modeling results by Niv (2012).

The study by Eldar & Niv (2015) was the second motivating work behind our study, as it provided evidence for a mood-biased learning paradigm. While Eldar & Niv (2015) did not focus on risky decision-making specifically, their study established mood's distinct biasing role in the reinforcement learning process (specifically, on the perception of outcomes). Our results showing mood's predictable influence on model-free decision-making are consistent with their study's finding of mood's biasing effect on learning. Concurrently, our models propose distinct mechanistic roles of mood. In our proposed model, we predicted mood to exacerbate asymmetric learning by pushing the relative difference between η^{+} and η^{-} in one direction or the other, leading to nonlinear effects on mood. In the model outlined in Eldar & Niv (2015), mood was proposed to directly bias the perception of outcomes through a mechanism involving symmetric learning with a mood-biased nonlinear transformation of outcome. While in general conceptual agreement, these two models propose distinct cognitive effects of mood. Future work could be dedicated to investigating whether the former, the latter, or a combination of the mood effects is occurring to explain mood-biased risky decision-making.

5.3 Limitations of our Design

The biggest limitation of our study concerns our ability to make claims about the role of "mood" in modulating risk preference, given the possible unreliability of our self-reported mood measures. We trusted that our mood inductions had "worked" by seeing how self-reported mood had changed as a result of these inductions. However, we found conflicting results, as the video condition emerged as the only reliable predictor of risk preference. To add to this limitation, happy-induced subjects relative to sad-induced subjects reported not only different valence changes, but also different arousal changes. Arousal's implication in other interactions involving state-level positive affect and block further suggest that valence may not explain the full story. Our inability to pinpoint the particular causal effect of the videos suggests a need to explore new metrics to evaluate the efficacy of our mood inductions. Self-reports are not the gold standard; previous studies have utilized more objective procedures like pupil diameter analysis[74] and voice analysis[39,12] to determine features like the valence and arousal of the induced mood, the extent to which the induced mood lasted throughout the experiment, and the specific emotions they had experienced (like anxiety and excitement).

Another limitation of our study is that we did not probe for risk attitude as an individual-difference measure. While we collected state-level and trait-level features related to our independent variable (mood inductions), we did not have a measure related to our dependent variable (risk attitudes). Thus, we could not test baseline risk-taking tendencies in interaction with the video condition in our mixed-effects models or visualize potentially informative correlations.

A related limitation comes from our between-subjects study design in which each subject only experienced one of the three mood inductions. As a result, we do not have a good indication as to how the mood inductions changed risk preferences within each subject relative to baseline risk-taking tendencies. A design in which the same subject receives different kinds of mood inductions would allow for better assessment of how risk preferences could change within-subject as a function of mood.

In our computational models, we had a fixed learning rate instead of a decaying learning rate, which would have been more representative of the sequential learning process in a model-free environment[35]. When we observed model fits to specific blocks, the asymmetric and utility models were shown to perform poorly, particularly in block 1. Utilizing a more dynamic learning rate that reflects the slowing of learning through trials may have produced better model fits and controlled for important interactions with learning in a risk-sensitive framework.

5.4 Future Directions

To address the most pressing limitation of our study, a future study should implement a design that captures an objectively reliable measure of mood. For example, speech is a rich, natural marker whose semantic and acoustic features have been shown to be highly predictive markers of various emotions and cognitive states[8,3]. Happy, sad, angry, scared, curious, confused, embarrassed, and stubborn are just a few examples of emotions that can be detected using deep learning[12]. Based on findings suggesting that emotions of the same valence can have differential effects on risk-taking tendencies[41,23], it is all the more important to gain more clarity about the specific cognitive states that are affected by different mood manipulations to give rise to distinct risk preferences.

In our study, we found the difference between subjects' "goodness" ratings of the risky and safe stimuli to be strongly correlated with behavioral risk preference. Given the observed effect of block on risk preference, we could collect "goodness" ratings after every block (instead of after just the third block). Visualizing separate correlations by block may clarify whether behavioral risk preference in that block was a function of inadequate learning (represented by a weaker correlation) or a function of true risk attitudes (represented by a stronger correlation). A subject's "goodness" of a stimulus is likely to encapsulate the subject's thoughts and feelings about both its mean and variance, both of which are expected to contribute to the overall "goodness" of the stimulus. To gain a clearer idea of subjective valuations, we can also ask specifically for subjects' explicit beliefs about the probabilities associated with stimuli outcomes. These subjective probabilities could help test (more explicitly) if different risk attitudes had emerged due to different subjective probabilities (representing mood-modulated perceptions of probability). For example, a positive mood may promote higher risk-taking tendencies by biasing the perceived probability of future positive outcomes. Studies have found that happy people, relative to sad people, report being more optimistic about the probability of a future positive outcome[73]. If we consider our two models of interest, the utility and asymmetric learning models, such a framework may align more closely with the utility model, since it involves a change of outcome valuation. On the other hand, if happy-induced and sad-induced subjects report the same probability measures for the risky stimulus, this result may align more closely with the asymmetric learning model, since objective probabilities are preserved in this model; the learning process itself would be more likely to have been affected by a mood-biased effect to produce distinct risk attitudes.

Niv (2012) found the nucleus accumbens (implicated in reinforcement learning[19]) to be sensitive to risk in a model-free context. Other studies focused on model-based decision-making (involving explicit knowledge of risk) showed several cortical areas to be risk-sensitive[29]. These findings suggest a high contextual dependence in how risk is tracked and represented in the brain. Given the unique design of our task – a model-free environment that was interspersed with deliberate mood manipulations – it would be interesting to analyze the patterns of activation[65] that are associated with mood-driven risk-sensitive learning. While we would expect to find risk sensitivity in the nucleus accumbens, mood's implication in a variety of brain regions (most notably, the limbic system[14]) presents interesting possibilities for future functional magnetic resonance imaging (fMRI) studies of this mood-driven effect.

In our experiment, we targeted a very specific type of risk preference – a risk preference for gains where stimuli of equal means but different variance were presented. We could imagine another case in which choosing the risky option would have resulted in a loss, or one in which the risky stimulus differed in the mean instead of the variance. These other possibilities represent assessments of risk in different neuroeconomic frameworks; future studies could test how mood's modulating effects differ across various types of risk measures. Loss aversion is a well-known cognitive bias in which individuals prefer avoiding losses to obtaining equivalent gains [34]. Given the heightened sensitivity for losses, mood may be more likely to affect risk-taking that elicits loss aversion, particularly because such an aversion inherently involves more emotions in the first place and would thus affect choices to a greater degree [2].

In the future, it would be interesting to replicate an improved version of our experiment on a clinical population characterized by high mood instability (like patients with bipolar disorder). A higher susceptibility to changes in mood would predict a higher effect of experimental mood inductions. Greater mood changes as a result of trait-level tendencies may produce more exacerbated differences in risk-taking tendencies as a function of mood, which may provide more clarity on mood-biased interactions on risk-sensitive reinforcement learning processes. Modeling data from a clinical population could not only provide insights into how destructive cognitive processes may result, but research could also inform more general computational models of cognition as a function of individual differences.

The computational modeling work presented here provides the groundwork for more rigorous approaches to modeling. Given that neurocomputational processes are incredibly complex and that our candidate models represent massive oversimplifications of the inner workings of the mind, future directions for research involve building more nuanced models of cognition, given the insights we have gained from this study. There is always room for a better model. Thus, building upon our current models is an immediate next step to better capture how mood might be driving asymmetric learning (or possibly another risk-sensitive learning framework) to produce our observed behavioral effects.

Additionally, future work should focus on more accurate ways for model comparison. Although our computational modeling results led to the selection of the asymmetric learning model, the difference in means between this model and the utility model was not drastic. While we can say that the asymmetric learning framework offered a better explanation, we cannot outrule the subjective utility framework. Previous work has proposed that hierarchical procedures yield better estimates of model parameters than do nonhierarchical, independent approaches[60]. Thus, to more accurately assess model performance, we may want to estimate parameters using Bayesian hierarchical parameter estimation.

5.5 Conclusion

Navigating through the stochasticity of life, we constantly make decisions under uncertainty, all under the filter of mood. After watching a feel-good movie, would you be more likely to order that fish taco? Our results suggest so. The simple act of watching an emotionally-charged video predictably influences our preferences for risk. Whether under the influence of mood or not, risk attitudes have been shown to be driven by nonlinear effects on learning for uncertain outcomes.

Research in mood-modulated risky behavior can have significant clinical implications. Computational psychiatry is a burgeoning field that draws from computational neuroscience and machine learning to more objectively and quantitatively extract insights from psychiatric conditions[4,5]. It offers a promising avenue for theory-driven and data-driven research to join hands in deriving better models of cognitive disturbance such that an informed mechanistic understanding of pathology could lead to more accurate diagnostic methods and better treatment outcomes. Given that risk-taking behavior is widely implicated in a range of mood-related disorders, a long-term goal of this research is to help elucidate how the computational workings of the mind could go askew to result in disturbances in mood, cognition, and behavior.

Our results suggest interesting mood-driven effects in risk-sensitive learning. Future studies should be dedicated to elucidating mood's specific computational role in mediating preferences for risk.

A

Honor Code

This paper represents my own work in accordance with University regulations.

Claire Lee

References

[1] American Psychiatric Association (2013). *Diagnostic and statistical manual of mental disorders: DSM-5*. Autor, Washington, DC, 5th ed. edition.

[2] Barkley-Levenson, E. E., Leijenhorst, L. V., and Galván, A. (2013). Behavioral and neural correlates of loss aversion and risk avoidance in adolescents and adults. *Developmental Cognitive Neuroscience*, 3:72–83.

[3] Bedi, G., Carrillo, F., Cecchi, G. A., Slezak, D. F., Sigman, M., Mota, N. B., Ribeiro, S., Javitt, D. C., Copelli, M., Corcoran, C. M., and et al. (2015). Automated analysis of free speech predicts psychosis onset in high-risk youths. *npj Schizophrenia*, 1(1).

[4] Bennett, D. and Niv, Y. (2018). Opening burtons clock: Psychiatric insights from computational cognitive models.

[5] Bennett, D., Silverstein, S. M., and Niv, Y. (2019). The two cultures of computational psychiatry. *JAMA Psychiatry*, 76(6):563.

[6] Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23.

[7] Calisir, S. and Pehlivanoglu, M. K. (2019). Model-free reinforcement learning algorithms: A survey. *2019 27th Signal Processing and Communications Applications Conference (SIU)*.

[8] Cecchi, G., Gutierrez, E., and Corcoran, C. (2019). 24.2 natural language processing studies of psychosis and its risk states. *Schizophrenia Bulletin*, 45.

[9] Chou, K.-L., Lee, T. M. C., and Ho, A. H. Y. (2007). Does mood state change risk taking tendency in older adults? *Psychology and Aging*, 22(2):310–318.

[10] Corrado, G. and Doya, K. (2007). Understanding neural coding through the model-based analysis of decision making. *Journal of Neuroscience*, 27(31):8178–8180.

[11] Crump, M. J. C., Mcdonnell, J. V., and Gureckis, T. M. (2013). Evaluating amazons mechanical turk as a tool for experimental behavioral research. *PLoS ONE*, 8(3).

[12] Davletcharova, A., Sugathan, S., Abraham, B., and James, A. (2015). Detection and analysis of emotion from speech signals. *Procedia Computer Science*.

[13] Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., and Dolan, R. J. (2011). Model-based influences on humans choices and striatal prediction errors. *Neuron*, 69(6):1204–1215.

[14] Dell'Acqua, F. and Thiebaut de Schotten, M. (2013). A revised limbic system model for memory, emotion and behavior. *Neuroscience and biobehavioral reviews*, 37.

[15] Eldar, E. and Niv, Y. (2015). Interaction between emotional state and learning underlies mood instability. *Nature Communications*, 6(1).

[16] Eldar, E., Rutledge, R. B., Dolan, R. J., and Niv, Y. (2016). Mood as representation of momentum. *Trends in Cognitive Sciences*, 20(1):15–24.

[17] Elliott, R., Rees, G., and Dolan, R. (1999). Ventromedial prefrontal cortex mediates guessing. *Neuropsychologia*, 37(4):403–411.

[18] Forgas, J. (1995). Mood and judgment: The affect infusion model (aim). *Psychological bulletin*, 117:39–66.

[19] Gale, J. T., Shields, D. C., Ishizawa, Y., and Eskandar, E. N. (2014). Reward and reinforcement activity in the nucleus accumbens during learning. *Frontiers in Behavioral Neuroscience*, 8.

[20] Gradinaru, A. (2014). The contribution of behavioral economics in explaining the decisional process. *Procedia Economics and Finance*, 16:417–426.

[21] Greifeneder, R., Bless, H., and Keller, J. (2005). When feelings matter: Determinants of the use of cognitive subjective feelings. *PsycEXTRA Dataset*.

[22] Gruber, J., Johnson, S., Oveis, C., and Keltner, D. (2008). Risk for mania and positive emotional responding: Too much of a good thing? *Emotion (Washington, D.C.)*, 8:23–33.

[23] Habib, M., Cassotti, M., Moutier, S., Houdé, O., and Borst, G. (2015). Fear and anger have opposite effects on risk seeking in the gain frame. *Frontiers in Psychology*, 6.

[24] Herman, A., Critchley, H., and Duka, T. (2018). Risk-taking and impulsivity: The role of mood states and interoception. *Frontiers in Psychology*, 9:1625.

[25] Hertwig, R., Barron, G., Weber, E. U., and Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8):534–539.

[26] Hewig, J., Hagemann, D., Seifert, J., Gollwitzer, M., Naumann, E., and Bartussek, D. (2005). A revised film set for the induction of basic emotions. *Cognition and Emotion*, 19.

[27] Hockey, G. R. J., Maule, A. J., Clough, P. J., and Bdzola, L. (2000). Effects of negative mood states on risk in everyday decision making. *Cognition and Emotion*, 14(6):823–855.

[28] Hsu, M. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science*, 310(5754):1680–1683.

[29] Huettel, S. A. (2005). Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices. *Journal of Neuroscience*, 25(13):3304–3311.

[30] Huettel, S. A., Stowe, C. J., Gordon, E. M., Warner, B. T., and Platt, M. L. (2006). Neural signatures of economic preferences for risk and ambiguity. *Neuron*, 49(5):765–775.

[31] Isen, A. M. and Patrick, R. (1983). The effect of positive feelings on risk taking: When the chips are down. *Organizational Behavior and Human Performance*, 31(2):194–202.

[32] Jessup, R. K., Bishara, A. J., and Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science*, 19(10):1015–1022.

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

[34] Kahneman, D. and Tversky, A. (2000). Choices, values, and frames. *Choices, Values, and Frames*, page 1–16.

[35] Katahira, K. (2015). The relation between reinforcement learning parameters and the influence of reinforcement history on choice behavior. *Journal of Mathematical Psychology*, 66:59–69.

[36] Kuhnen, C. M. and Knutson, B. (2005). The neural basis of financial risk taking. *Neuron*, 47(5):763–770.

[37] Kuppens, P., Oravecz, Z., and Tuerlinckx, F. (2010). Feelings change: Accounting for individual differences in the temporal dynamics of affect. *Journal of personality and social psychology*, 99:1042–60.

[38] Kusev, P., Purser, H., Heilman, R., Cooke, A. J., Schaik, P. V., Baranova, V., Martin, R., and Ayton, P. (2017). Understanding risky behavior: The influence of cognitive, emotional and hormonal factors on decision-making under risk. *Frontiers in Psychology*, 8.

[39] Lee, C. M. and Narayanan, S. (2005). Toward Detecting Emotions in Spoken Dialogs. In *IEEE Transactions on Speech and Audio Processing*, volume 12, pages 293–303.

[40] Lee, C. S. (2020). Psy 360 computational models of cognition final project.

[41] Leith, K. P. and Baumeister, R. F. (1996). Why do bad moods increase self-defeating behavior? emotion, risk tasking, and self-regulation. *Journal of Personality and Social Psychology*, 71(6):1250–1267.

[42] Mellers, B. A., Schwartz, A., Ho, K., and Ritov, I. (1997). Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychological Science*, 8(6):423–429.

[43] Mihatsch, O. and Neuneier, R. (2002). *Machine Learning*, 49(2/3):267–290.

[44] Nash, S. G. (1984). Newton-type minimization via the lanczos method. *SIAM Journal on Numerical Analysis*, 21(4):770–788.

[45] Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3):139–154.

[46] Niv, Y., Edlund, J. A., Dayan, P., and Odoherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, 32(2):551–562.

[47] Niv, Y., Joel, D., Meilijson, I., and Ruppin, E. (2002). Evolution of reinforcement learning in uncertain environments: A simple explanation for complex foraging behaviors. *Adaptive Behavior*, 10(1):5–24.

[48] Nygren, T. E., Isen, A. M., Taylor, P. J., and Dulin, J. (1996). The influence of positive affect on the decision rule in risk situations: Focus on outcome (and especially avoidance of loss) rather than probability. *Organizational Behavior and Human Decision Processes*, 66(1):59–72.

[49] Patel, R., Lloyd, T., Jackson, R., Ball, M., Shetty, H., Broadbent, M., Geddes, J. R., Stewart, R., Mcguire, P., Taylor, M., and et al. (2015). Mood instability is a common feature of mental health disorders and is associated with poor clinical outcomes. *BMJ Open*, 5(5).

[50] Pavlickova, H., Varese, F., Smith, A., Myin-Germeys, I., Turnbull, O. H., Emsley, R., and Bentall, R. P. (2013). The dynamics of mood and coping in bipolar disorder: Longitudinal investigations of the inter-relationship between affect, self-esteem and response styles. *PLoS ONE*, 8(4).

[51] Payne, J. W. (1985). Psychology of risky decisions. *Behavioral Decision Making*, page 3–23.

[52] Pighin, S., Bonini, N., Savadori, L., Hadjichristidis, C., and Schena, F. (2014). Subjective feelings questionnaire. *PsycTESTS Dataset*.

[53] Posner, J., Russell, J. A., and Peterson, B. S. (2005). The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and Psychopathology*, 17(03).

[54] Preuschoff, K., Bossaerts, P., and Quartz, S. R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron*, 51(3):381–390.

[55] Raghunathan, R. and Pham, M. T. (1999). All negative moods are not equal: Motivational influences of anxiety and sadness on decision making. *Organizational Behavior and Human Decision Processes*, 79(1):56–77.

[56] Reddy, L. F., Lee, J., Davis, M. C., Altshuler, L., Glahn, D. C., Miklowitz, D. J., and Green, M. F. (2013). Impulsivity and risk taking in bipolar disorder and schizophrenia. *Neuropsychopharmacology*, 39(2):456–463.

[57] Reyna, V. F. and Zayas, V. (2015). *Neuroscience of Risky Decision Making*. American Psychological Association.

[58] Ricciardi, V. (2008). *The Psychology of Risk: The Behavioral Finance Perspective*, pages 85–111.

[59] Schaefer, A., Nils, F., Sanchez, X., and Philippot, P. (2010). Assessing the effectiveness of a large database of emotion-eliciting films: A new tool for emotion researchers. *Cognition and Emotion*, 24(7):1153–1172.

[60] Scheibehenne, B. and Pachur, T. (2014). Using bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychonomic Bulletin Review*, 22(2):391–407.

[61] Scholer, A., Zou, X., Fujita, K., Stroessner, S., and Higgins, E. (2010). When risk seeking becomes a motivational necessity. *Journal of Personality and Social Psychology*, 99:215–231.

[62] Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599.

[63] Shen, Y., Tobia, M. J., Sommer, T., and Obermayer, K. (2014). Risk-sensitive reinforcement learning. *Neural Computation*, 26(7):1298–1328.

[64] Shepperd, J. A. and Mcnulty, J. K. (2002). The affective consequences of expected and unexpected outcomes. *Psychological Science*, 13(1):85–88.

[65] Sherman, L., Steinberg, L., and Chein, J. (2018). Connecting brain responsivity and real-world risk taking: Strengths and limitations of current methodological approaches. *Developmental Cognitive Neuroscience*, 33:27–41.

[66] Sutton, R. and Barto, A. (1998). Reinforcement learning: An introduction. *IEEE Transactions on Neural Networks*, 9(5):1054–1054.

[67] Tamir, M. and Robinson, M. (2007). The happy spotlight: Positive mood and selective attention to rewarding information. *Personality social psychology bulletin*, 33:1124–36.

[68] Tymula, A., Belmaker, L. A. R., Roy, A. K., Ruderman, L., Manson, K., Glimcher, P. W., and Levy, I. (2012). Adolescents risk-taking behavior is driven by tolerance to ambiguity. *Proceedings of the National Academy of Sciences*, 109(42):17135–17140.

[69] Verschoor, A., D'Exelle, B., and Perez-Viana, B. (2016). Lab and life: Does risky choice behaviour observed in experiments reflect that in the real world? *Journal of Economic Behavior Organization*, 128:134–148.

[70] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Jarrod Millman, K., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C., Polat, İ., Feng, Y., Moore, E. W., Vand erPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and Contributors, S. . . (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*.

[71] Watson, D., Clark, L. A., and Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The panas scales. *Journal of Personality and Social Psychology*, 54(6):1063–1070.

[72] Wright, S. (1999). Numerical optimization. *Springer Series in Operations Research and Financial Engineering*.

[73] Wright, W. and Bower, G. (1992). Mood effects on subjective probability assessment. *Organizational Behavior and Human Decision Processes*, 52:276–291.

[74] Yamanaka, K. and Kawakami, M. (2009). Convenient evaluation of mental stress with pupil diameter. *International Journal of Occupational Safety and Ergonomics*, 15(4):447–450.

[75] Yang, X.-S. (2019). *Introduction to algorithms for data mining and machine learning*. Academic Press.

[76] Youngstrom, E., Murray, G., Johnson, S., and Findling, R. (2013). The 7 up 7 down inventory: A 14-item measure of manic and depressive tendencies carved from the general behavior inventory. *Psychological assessment*, 25.

[77] Yuen, K. and Lee, T. (2003). Could mood state affect risk-taking decisions? *Journal of affective disorders*, 75:11–8.