Essays on Sequential Sampling in Value-Based Choice

Thesis by
Brenden Eum

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2025
Defended August 20, 2024

Brenden Eum
ORCID: 0000-0002-5484-495X

ACKNOWLEDGEMENTS

I am extremely grateful to have had the chance to do my graduate studies at Caltech. The people around me have made these past 5 years unforgettable, and I don't think I've ever been surrounded by so many people who have all had such a strong influence on me.

To Antonio Rangel,

It was one of the greatest privileges of my life to have had you as my advisor. You are (by far) the best advisor I could have ever asked for, and when I think of the biggest role model in my life, whether it be in mentoring, science, or family, it's you. I move on to the next stage in my career with the confidence that, if I could survive your questions during practice talks, there is literally no question out there that I can't handle. Thank you for shaping me into the confident and capable researcher I am today and for inspiring me to be a better person every day.

To Mike Woodford,

Without you, I might have never found my passion for neuroeconomics. Thank you for all the discussions and guidance over the past 7 years. You have altered the course of my life in more ways than one, and I'll always be grateful that you chose to teach your Cognitive Mechanisms of Economic Behavior course in the fall of 2017.

To Colin Camerer,

There's always a spike in my dopamine levels when I run into you. Whatever the topic, you always know an interesting fact or dataset to contribute, which I think correlates with the style of your research (though I can't speak to the direction of causality). When I think of the style in which I'd like to approach research topics in the future, I think of you. Thank you for all of your advice over these past 5 years.

To Charlie Sprenger,

I am deeply grateful for all your encouragement over the years, and I'm always inspired every time I see you give a talk. I've never met anyone better at taking complicated concepts and explaining them so clearly and concisely. When I think of the teacher I'd like to be, I think of you.

I am extremely grateful to have shared my time in the Rangel Neuroeconomics Lab with a special group of people: Zeynep Enkavi, Wenning Deng, Pantelis Vafidis,

Thomas Henning, Liz Schroder, and (for a short time) Doug Lee. You all inspired me to work hard, stay disciplined, and enjoy balance. I'm surprised the third floor could sustain the weight of that much passion for both science and life. Perhaps that's why lab meetings had to be held in the basement.

To the other members of SDN, it was a joy going through grad school together. Sanghyun, Weilun, Aniek, Qianying, Wenning, Thomas, and Sneha, I'll miss these years with you all.

To Kate Huang, thank you for getting me psyched on math and the outdoors. My legs will never forget all the cycling, backpacking, and climbing trips together, and my brain will never forget how you pulled proofs out of thin air along the way.

I would also like to thank a few professors outside of Caltech. To Guillaume Frechette at NYU, you once threw an eraser at me for incorrectly explaining Nash equilibrium. This made me curious how I could get paid to do the same, and with your guidance, I started down the path to academia. To Irasema Alonso at Columbia, you showed me how much a professor can care for her students and gave me the confidence to continue on to a PhD.

To my mom, thank you for always loving and supporting me. Last but not least, to my partner Yujia Wan, who has stuck with me through the highs and lows that this path can bring, thank you for being my foundation.

ABSTRACT

This dissertation comprises three chapters related to the fields of psychology, computational neuroscience, and experimental economics. Chapters 1 and 2 use experimental and computational methods to study the role of attention in simple, value-based choices. Chapter 3 examines risky choices from experience and tests some of the underlying assumptions of sequential sampling models.

A growing body of research has shown that simple choices involve the construction and comparison of values at the time of decision. These processes are modulated by attention in a way that leaves decision makers susceptible to attentional biases. In Chapter 1, co-authored with Stephanie Dolbier and Antonio Rangel, we studied the role of peripheral visual information on the choice process and on attentional choice biases. We used an eye-tracking experiment in which participants (N = 50 adults) made binary choices between food items that were displayed in marked screen "shelves" in two conditions: (a) where both items were displayed, and (b) where items were displayed only when participants fixated within their shelves. We found that removing the nonfixated option approximately doubled the size of the attentional biases. The results show that peripheral visual information is crucial in facilitating good decisions and suggest that individuals might be influenceable by settings in which only one item is shown at a time, such as e-commerce.

In Chapter 2, co-authored with Stephen Gonzalez and Antonio Rangel, we studied the role of attention in aversive risky choices where all outcomes were unpleasant. We used two eye-tracking experiments in which participants made binary choices between two lotteries in two conditions: (a) a gain condition where outcomes for lotteries were weakly positive, and (b) a loss condition where outcomes were weakly negative. Contrary to the predictions of the standard aDDM, we found that attentional choice biases in the loss condition were identical to those found in the gain condition, suggesting that attention nudges choices towards the attended option even in losses. To explain these results, we propose a variation of the Attentional Drift-Diffusion-Model (called the Hybrid aDDM) that incorporates (a) both a value-dependent and a value-independent effect of attention on the choice process and (b) reference-dependent value signals. We show that the observed attentional choice biases and other behavioral signatures in the loss condition can only be explained by the Hybrid aDDM with a reference-point rule that sets the reference-point at or below the minimum possible outcome in a given context.

In Chapter 3, co-authored with Antonio Rangel, we establish that sequential sampling models apply to risky decisions from experience and test some of the underlying assumptions of these models. We ran an online study in which participants chose to Play or Skip a slot machine, based on a stream of samples drawn from its outcome distribution. We found evidence for leakage, collapsing decision boundaries, and a delay in sample integration. We also found evidence of non-linear sample weighting depending on when the sample occurred during the trial. As a bonus, we established a link between the fixed decision boundaries in a Drift-Diffusion-Model and a Modified Probit model, allowing for estimation of decision boundaries in cumulative sample space without the need to fit a computational model.

PUBLISHED CONTENT AND CONTRIBUTIONS

Eum, Brenden, Stephanie Dolbier, and Antonio Rangel (2023). "Peripheral Visual Information Halves Attentional Choice Biases". In: *Psychological Science* 34.9, pp. 984–998. DOI: 10.1177/09567976231184878.

Contributions: analyzed the data, wrote the paper.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

INTRODUCTION

Nearly every aspect of human behavior is the result of some form of decision-making. From the simple choices we make daily, like selecting what to eat for breakfast, to the complex decisions involved in financial investments, understanding how individuals evaluate options and make decisions is crucial. The traditional view in economics assumes that decision-making is a rational process where individuals weigh all available options and make choices that maximize their utility. However, a growing body of research in behavioral economics and cognitive sciences challenges this notion, showing that decision-making is often influenced by factors beyond pure rationality, including attention, cognitive biases, and the context in which choices are presented.

In recent years, the role of attention in decision-making has garnered significant attention, particularly in economic choices where the decision-maker must evaluate and compare the value of different options. Overt attention has been shown to significantly influence the decision-making process, often leading to biases where certain options are favored simply because they receive more attention. These attentional biases can have profound implications, particularly in environments where the presentation of options can be controlled, such as in online shopping platforms or digital advertising.

Chapters 1 and 2 of this dissertation delve into the intricacies of attention and its impact on decision-making. Chapter 1 investigates whether peripheral visual information is processed and incorporated into the choice process, whether consciously or subconsciously. Using a gaze-contingent eye-tracking experiment where participants make binary choices between food items, we explore how the removal of the options in participants' visual periphery influences attentional choice biases. We find that attentional choice biases approximately double in size when peripheral visual information is hidden, showing that peripheral visual information is crucial in facilitating good choices. This study is particularly relevant in understanding how visual environments, such as supermarket shelves or e-commerce websites, can be optimized to guide consumer choices.

Chapter 2 shifts focus to aversive choices, where decision-makers are faced with unpleasant outcomes. We run two experiments involving both appetitive and aversive risky choices and find identical choice biases across the two conditions. These

results suggest that, even in losses, attention is nudging choices towards the attended option. This is surprising given the predictions of a canonical model that links attention to choice. We propose another model that splits the effect of attention on the choice process into two effects and incorporates reference-dependent value signals. We show that the observed attentional choice biases and other behavioral signatures in the loss condition can only be explained by our model with a reference-point rule that sets the reference-point at or below the minimum possible outcome in a given context.

Chapter 3 examines the applicability of sequential sampling models to risky decisions from experience. Using an online study, we test key assumptions of these models, uncovering phenomena such as leakage, collapsing decision boundaries, and non-linear sample weighting depending on when information is accumulated. This investigation not only validates the use of these models in understanding how real-world decisions are made from experience, but also provides new insights into the temporal dynamics of how individuals integrate information. As a bonus, we also establish a link between the decision boundaries of the canonical Drift-Diffusion-Model and a Modified Probit model, allowing for computationally-inexpensive estimation of these decision boundaries in cumulative sample space.

Together, these essays contribute to the growing understanding of the cognitive processes underlying decision-making. By examining the role of attention across different contexts—ranging from everyday food choices to unpleasant decisions under uncertainty—we offer new insights on how attention can both facilitate and bias the decision-making process. As the world increasingly shifts towards digital and visual environments, the insights gained from this research are poised to inform the design of systems and interventions that can improve decision-making quality and outcomes.

Chapter 1

PERIPHERAL VISUAL INFORMATION HALVES ATTENTIONAL CHOICE BIASES

1.1 Introduction

Everyday we face two different types of choice situations. Sometimes we are presented with all of the available options at once, as when we face a supermarket shelf or a buffet table. In other cases, such as many shopping websites, we are presented with one option at a time, which changes sequentially at our own pace. In both cases our overt visual attention is deployed to one option at a time. But the two situations differ on the availability of peripheral visual information about the nonfixated options, which in principle could be used to guide the choice process.

A growing number of experiments have studied the role of visual attention in simple choice and have found that increases in the relative attention received by a desirable option are associated with an increase in the frequency with which it is chosen, all else being equal (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Krajbich, Lu, et al., 2012; S. M. Smith and Krajbich, 2018; S. M. Smith and Krajbich, 2019; J. F. Cavanagh et al., 2014; S. E. Cavanagh et al., 2019; Sepulveda et al., 2020; Thomas et al., 2019; Gluth, Spektor, and Rieskamp, 2018; Gluth, Kern, et al., 2020; Fisher, 2017; Towal, Mormann, and Koch, 2013). Although the exact mechanism behind the attentional bias remains unknown, foveation seems to facilitate the process of value computation and integration in a way that is consistent with overweighting fixated items relative to nonfixated ones. This is formalized in the Attentional Drift-Diffusion-Model (aDDM), which is able to provide a quantitative account of the relationship between fixations, choices, and reaction times (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Krajbich, Lu, et al., 2012; S. M. Smith and Krajbich, 2018; S. M. Smith and Krajbich, 2019). The aDDM predicts that choices can be biased through exogenous manipulations of relative fixation time, consistent with the findings of multiple studies (Armel, Beaumel, and Rangel, 2008; Tavares, Perona, and Rangel, 2017; Hare, Malmaud, and Rangel, 2011; Pärnamets et al., 2015; Ghaffari and Fiedler, 2018; Kunar et al., 2017; Peschel, Orquin, and Mueller Loose, 2019; Shimojo et al., 2003).

Our goal is to study the role of peripheral visual information on the choice process

and on attentional choice biases. In particular, do we use the same choice algorithm when all options are presented simultaneously, as when we shop at the market, and when they are presented sequentially, as when we shop online? If not, does the absence of peripheral information change the fixation process and the magnitude of the attentional biases?

We study these questions using an eye-tracking experiment in which subjects make binary choices between foods that are displayed in marked screen "shelves" in two different conditions: (1) a simultaneous condition in which both items are displayed on the screen at the time of choice, and (2) a gaze contingent condition in which items are only displayed when subjects fixate within their shelves. Most previous studies have used choice tasks in which all options are displayed simultaneously, although a handful have used gaze-contingent presentation of stimuli (Simion and Shimojo, 2006; Folke et al., 2016; Sepulveda et al., 2020; Franco-Watkins and J. G. Johnson, 2011). However, none have compared the two situations directly, which is necessary to understand the effect of peripheral visual information on the choice process.

Based on what is known about choice in the simultaneous case, and the seemingly minor change involved in removing the nonfixated options from peripheral vision, it is natural to hypothesize that similar algorithms are at work in both conditions, albeit with some differences. In particular, the aDDM suggests two non-mutually exclusive mechanisms through which removing the nonfixated options from the visual field might affect choices. First, it might increase the overweighting of fixated relative to nonfixated items, which would result in an increased attentional bias. Second, it might change the fixation process in a way that exacerbates the attentional biases, for example by increasing the asymmetry on fixation time across options.

Understanding the role of peripheral visual information in simple choice is important for multiple reasons. First, despite the robustness of the attentional biases identified in previous work, we do not know what are the channels through which covert and overt visual attention influence decisions, nor their relative contribution to choice. Decades of work in visual attention have shown that a substantial amount of information is processed through peripheral visual attention (Carrasco, 2011; Perkovic et al., 2022; Wästlund, Shams, and Otterbring, 2018), which raises the puzzle of why and how fixations matter so much in economic choices, even when making decisions among familiar items. Second, the transition to e-commerce has increased the frequency with which our decisions are made in sequential presentation

settings. We need to understand the impact that this has on the choice algorithms and their associated biases in order to design interfaces and nudges that enhance choice quality.

To preview the results, we find that removing the nonfixated options has little impact on the average quality of choices. However, we also find that it approximately doubles the magnitude of attentional choice biases, which make decision makers more susceptible to marketing interventions (e.g., packaging) that affect attention independently of the value of products. We also find that the removal of the nonfixated options slows down the fixation and decision process considerably, but that the impact on attentional biases is driven mostly by an increase in the tendency to overweight the value of fixated options.

All data and code are available for download at the Rangel Neuroeconomics Lab website (`www.rnl.caltech.edu`). The design and analysis plans for this study were not preregistered.

1.2 Methods

Task

We investigated the role of peripheral information about nonfixated stimuli using the task depicted in Fig. 1.1. Subjects made decisions in two conditions: (1) a visible condition in which both items were displayed on the screen at the time of choice, and (2) a hidden condition in which items were displayed only when subjects fixated within the location associated with the stimulus. Subjects were asked to refrain from eating for 2 h before the start of the experiment, and to refrain from eating any foods afterwards during a 1 h waiting period, except for the snack that they chose in a randomly selected trial, which was given to them at the end of the experiment.

Subjects participated in two tasks. First, they were asked to provide liking ratings for 60 snack foods available at local stores ("How much would you LIKE to eat this food?", 1 = "don't like" to 5 = "like a lot", 0.25 intervals). Each item was rated twice, in random order, using a slider bar controlled by the arrow keys, and initialized to a random location to reduce anchoring effects. We use the average of the two ratings as a measure of each item's value.

Second, subjects made choices between two food items, shown on the left and right sides of the screen, in two separate conditions: (1) a visible condition where both items were shown simultaneously, and (2) a hidden condition where items were shown only when subjects fixated within their region of interest (ROI). The ROIs

were indicated in both treatments with a white box (Fig. 1.1). Trials started with an enforced 500 ms central fixation. Subjects indicated their choices with the left and right arrow keys, and responded at their own pace. The selected option was highlighted for 1 s and trials were separated by a 1 s blank screen. Subjects made 360 choices in the exploratory sample and 400 in the confirmatory sample, half on each experimental condition. The task was divided into four equal-sized blocks, two with the hidden condition and two with the visible condition, in random order.

The choice pairs in the exploratory data set were randomly selected from all 60 food items. In the confirmatory sample, they were constructed as follows. We used each subject's ratings to prune the stimulus set down to the 40 food items that resulted in the most uniform distribution of ratings, in order to maximize the spread of rating differences across choice pairs. Stimuli for each trial were then randomly selected, subject to the constraint that they be used four times per block in the exploratory data set, and five times per block in the confirmatory data set. All 60 foods were shown once every 30 trials in the exploratory data set, and all 40 foods were shown once every 20 trials in the confirmatory data set.

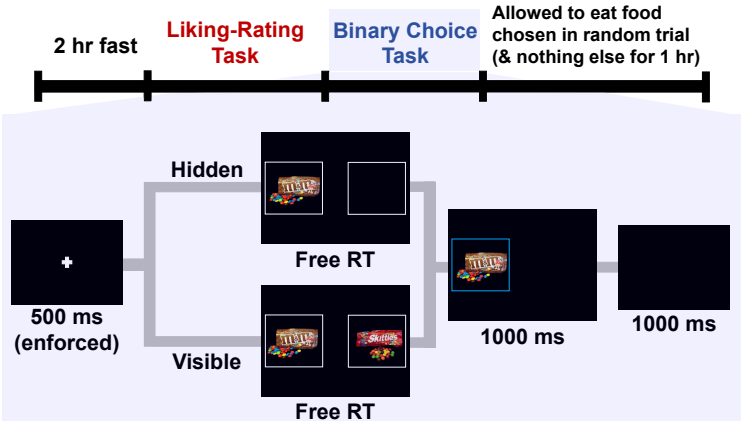

Figure 1.1: Task. Subjects had to fixate on a center fixation cross for 500 ms for the trial to start. In the visible condition subjects were presented with two snack food items simultaneously, each located within a white box on the left and right sides of the screen. In the hidden condition subjects had to fixate within the white boxes in order to reveal the snack food item inside. Subjects indicated their response at their own pace with a keyboard press. Once a choice was made, a blue box highlighted the selection for 1 s, followed by a 1 s inter-trial interval.

Participants

50 subjects (mean age = 30.8, 34 female) were recruited from Caltech and the surrounding community using flyers. We pre-screened subjects for a self-reported

liking for snack foods (e.g. candy and potato chips) and against requiring glasses for vision correction that might interfere with eye-tracking. Subjects were paid a $35 participation fee. The experiment was approved by Caltech's IRB.

In order to obtain high quality data, we implemented a subject filter at the data collection stage. Immediately after data collection we deleted subjects who failed any of the following criteria: (1) correlation between the two liking ratings of at least 70%, (2) mean RT in choice trials between 0.7 and 6 s, (3) probability of choosing the best item significantly different from chance (based on a binomial test), and (4) at most 10% missing fixation data. Data collection continued until 50 subjects passed the data quality criteria. The first 25 subjects were allocated to the exploratory sample, the other 25 to the confirmatory sample. The number of trials per subject, and the number of subjects, were chosen based on related studies which have shown that this sample size provides reliable estimates of the parameters and effects of interest.

Eye-Tracking

Subjects' fixation patterns were recorded using an EyeLink 100 desk eye-tracker at 500 Hz. Subjects sat approximately 60 cm from a 1920×1080 pixel monitor. Food image sizes were 403×302 pixels. Fixations within the ROI for the left food were classified as "left", those within the right food's ROI were classified as "right", and those outside the two ROIs were classified as "blank". If a sequence of blank fixations was recorded between two fixations of the same type (e.g. left-blank-blank-left), they were re-coded as a fixation of the same type (e.g., left-left-left-left), since blank fixations of this type are typically due to eye-tracking noise and tend to be quite short. Blank fixations recorded between two fixations of different types (e.g. left-blank-right) were coded as a saccade period between fixations. Trials in which any eye-tracking information is missing are dropped from further analysis (with a mean of 6 and 4 trials per subject in the exploratory and confirmatory datasets, respectively).

Data Analysis Strategy

In order to be able to explore the data in detail, while avoiding the type of statistical problems that have raised questions about the validity of some published research, we collected two separate datasets with 25 subjects each. We used the first one to carry out exploratory analyses until we understood the data generating process in sufficient detail. Based on this, we pinned down a set of analyses and tests

that were carried out in a second confirmatory dataset of equal size. Thus, the confirmatory dataset serves as a replication of our findings, and provides unbiased statistics for hypothesis testing. Given the similarity of the estimates and findings in both samples, and in the spirit of meta-analysis, we also provide results on the pooled sample and describe summary statistics in terms of the pooled estimates.

Computational Model

As illustrated in Fig. 1.2, the Attentional Drift-Diffusion-Model (aDDM) is a version of the Drift-Diffusion-Model of binary choice (Ratcliff and McKoon, 2008; Gold and Shadlen, 2007; Ratcliff, P. L. Smith, et al., 2016) in which value sampling is affected by fixation location. Subjects integrate noisy values signals into an evolving evidence process. Evidence starts every trial at an initial location b, which may include some bias towards one of the options if $b \neq 0$. A choice is made the first time evidence crosses one of two pre-specified barriers, which are fixed at 1 for the left item, and -1 for the right item. The identity of the barrier crossed determines which option is chosen. Critically, evidence evolves as the following diffusion process:

$$Evidence_t = Evidence_{t-1} + \mu_t + e_t \tag{1.1}$$

where e_t is i.i.d. white Gaussian noise with variance σ^2, and the slope of the process depends on the fixation location. In particular, when the left item is fixated, the slope of integration is $\mu_t = d(V_{left} - \theta V_{right})$, and when the right item is fixated is $\mu_t = d(\theta V_{left} - V_{right})$, where d is a parameter controlling the speed of integration and θ is a parameter controlling the attentional bias. When $\theta = 1$, the fixations do not affect choices, there are no attentional biases, and the process reduces to a standard DDM. In contrast, when $\theta < 1$, the value of the fixated item is overweighted relative to the nonfixated value, which results in an attentional bias that increases as θ gets smaller.

Importantly, the aDDM assumes that the fixation process is orthogonal to the state of evidence in any given trial. Thus, when simulating the model, we sample fixations from the observed fixation distributions, separately for first and middle fixations.

aDDM Fitting

We fit the aDDM using a hierarchical Bayesian model, separately for the visible and hidden conditions, using the methods and associated toolbox developed by Lombardi and Hare (Lombardi and Hare, 2021). We estimate the model separately for the

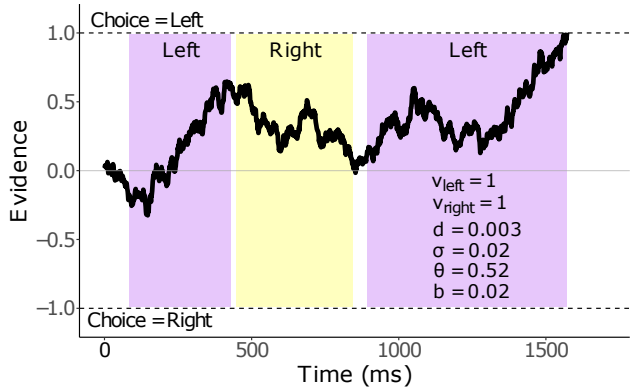

Figure 1.2: aDDM Example. An illustration of how the aDDM makes decisions in a sample trial. Colored vertical bands denote fixation locations.

exploratory, confirmatory, and pooled datasets. In every case, the model is fit using only the odd trials, as the even trials are reserved for out-of-sample predictions. As described in Fig. A.1, the model has the following free parameters, both at the group and individual levels: the evidence accumulation drift rate (d), the standard deviation of the Gaussian noise for the drift process (σ), the attentional bias parameter (θ), and the initial bias of the drift process (b). Posterior distributions were estimated using Markov Chain Monte Carlo methods with 3 chains for a total of 55,000 burn-in samples and 30,000 samples from each of the posteriors. Gelman-Rubin statistics for all estimates are at or below 1.1, indicating convergence.

The model that we estimate and report in the paper specifies priors without any correlation of parameters across the hidden and visible conditions. We do this to maximize the extent to which our posterior estimates are driven by the data. However, in order to investigate the role of the uncorrelated priors on our model fits, we also estimated a version of the model in which the priors for the same parameter in the visible and hidden conditions are correlated. In particular, each parameter (x) consists of two parts: a baseline ($\bar{x}$) and a hidden-condition deviation (Δx). See Fig. A.2 for details. As discussed further below, both models generated very similar parameter estimates.

Out-of-Sample Simulations

Even-numbered trials were set aside as out-of-sample data, in order to compare them to the predictions of aDDM model fitted on the odd trials. We simulate 10 datasets for each subject and condition, using the same rating pairs encountered in the experiment. For each simulated data set we sample a set of parameters from

the joint posterior distribution for that subject and condition. Then we simulate each trial as follows. We sample all fixation duration statistics from their observed empirical distributions in the even trials, conditional on the hidden or visible condition. For example, when simulating a hidden condition trial, evidence for the trial is initialized at the bias parameter and a trial-specific drift rate parameter is sampled. Evidence evolves based only on the noise up to the duration of the sampled latency to first fixation. Afterwards, a maximum first fixation duration is sampled from the distribution of first fixations in the hidden condition, and evidence evolves according to the drift rate, noise, and attentional bias parameters depending on the fixation location, as described in the Computational Model section. If a barrier is crossed before the maximum fixation duration is reached, the process is terminated and the choice and RT are recorded. Otherwise, a new saccadic duration and maximum fixation duration are sampled from the distributions of saccades and middle fixations in the hidden condition, respectively. The process is repeated until a choice is made. Note that this assumes that the value of the nonfixated item is known during the first fixation, which is unrealistic and interferes with the quality of our fits.

Hierarchical regressions

All the logistic and linear regressions reported in the paper are based on standard hierarchical models with random coefficients for all parameters. The regressions are implemented using the brms R-package (Bürkner, 2017; Bürkner, 2018) and used the default weakly informative priors, occasionally scaled depending on the units of the independent variable. Posterior distributions were estimated using 3 chains for a total of 9,000 burn-in samples and 9,000 samples from each of the posteriors. See the companion data and code package for details (`https://www.rnl.caltech.edu/publications/`).

1.3 Results

Basic Psychometrics

The top row of Fig. 1.3 depicts the psychometric choice curve, with each experimental condition and dataset separated. See Table A.1 for the associated regression estimates and test statistics. We find a small but significant increase in the responsivity of choices to value differences in the hidden condition. The middle row of Fig. 1.3 depicts reaction times (RTs) as a function of choice difficulty. We find that RT increases with choice difficulty, that average RTs are about 32% (520 ms) slower in the hidden condition, and that this slowdown does not vary significantly

with choice difficulty. The bottom row of Fig. 1.3 depicts the number of fixations as a function of choice difficulty. We find that the number of fixations increases with choice difficulty, and are approximately similar in both conditions, except for a small flattening in the slope of the fixation curve in the hidden condition.

Together, these results show that removing the nonfixated items slows down the choice process, but has a negligible effect on the quality of average choices (probability best chosen visible = 0.865 ± 0.007 (±SEs), probability best chosen hidden = 0.876 ± 0.007, d = 0.19, t(49) = 1.81, p = 0.08).

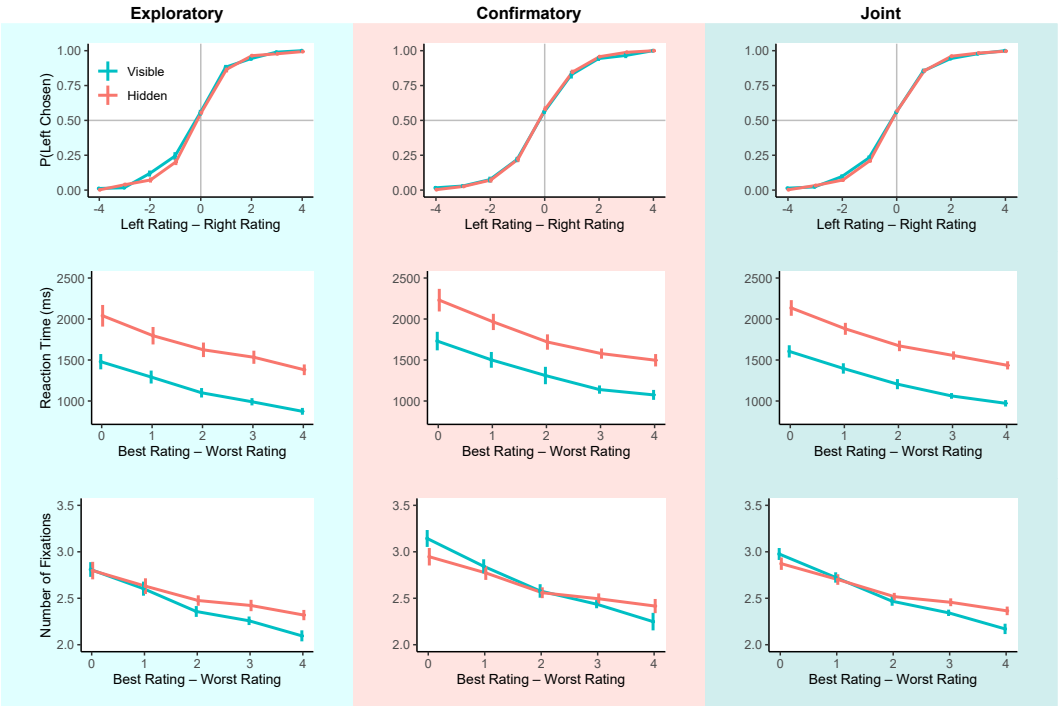

Figure 1.3: Basic Psychometrics. (Top) The probability of choosing the left item as a function of its relative value. (Middle) Response time as a function of trial difficulty, as measured by the rating difference between the best and worse items. (Bottom) The number of fixations as a function of trial difficulty. Columns indicate which dataset generated the figures. Error bars show standard errors of the mean across participants.

Fixation Process

Fig. 1.4 and Table A.2 explore the fixation process in more detail. The goal here is to understand the impact that removing nonfixated items has on the fixation process, which is essential to understand how it affects attentional biases.

The first row of Fig. 1.4 depicts the probability that the first fixation is to the best

item, as a function of choice difficulty. The first fixation location is at chance in both conditions. Fig. A.3 (top row) shows that there is no difference between conditions on the probability of first fixating left (probability first fix. left visible = 0.759 ± 0.042, probability first fix. left hidden = 0.803 ± 0.048, d = 0.14, t(49) = 1.64, p = 0.11). Fig. A.3 (bottom row) and Table A.3 show that there is also no difference between conditions on the latency to the start of the first fixation.

The second row of Fig. 1.4 depicts the mean duration of first, middle, and last fixations, separately for the two conditions. We find that the three types of fixations are longer in the hidden condition by about 40% on average (Δfirst=160ms, d=1.43, t(49)=11.3, p=3e$-$15; Δmiddle=145ms, d=0.82, t(49)=8.51, p=3e$-$11; Δlast=191ms, d=1.86, t(49)=17.02, p=0). Note that this is consistent with the RT results above: an average trial has 3 fixations, and each fixation is on average 165 ms longer in the hidden condition, which implies that decisions should take 495 ms longer, just shy of the observed RT difference.

The third row of Fig. 1.4 depicts middle fixation durations as a function of choice difficulty. We find that middle fixation durations increase with choice difficulty. Fig. A.4 and Table A.4 show that this difference is driven by the value of the fixated item: in the hidden condition middle fixation durations increase with the value of the fixated item, whereas the opposite occurs in the visible condition. Interestingly, Fig. A.4 also shows that middle fixation durations decrease with the value of the nonfixated item even in the hidden condition.

The fourth row of Fig. 1.4 depicts the first fixation duration as a function of choice difficulty. We find that duration is independent of value in both conditions, and about 46% (160 ms) longer in the hidden condition. See Fig. A.4 and Table A.4 for additional results.

The bottom row of Fig. 1.4 shows the relationship between relative value and relative fixation time. In both conditions, the relationship exhibits an S-shape. Both items are fixated the same amount when they have equal value, but otherwise the better item is fixated longer, with the asymmetry on fixation time increasing in the value advantage. In addition, the effect is stronger in the hidden condition, and as a result the distribution of net fixation times is more asymmetric in favor of the better item in this condition. Note that, since the fixated item is overweighted in the aDDM, this asymmetry in relative fixation time facilitates choosing the better option.

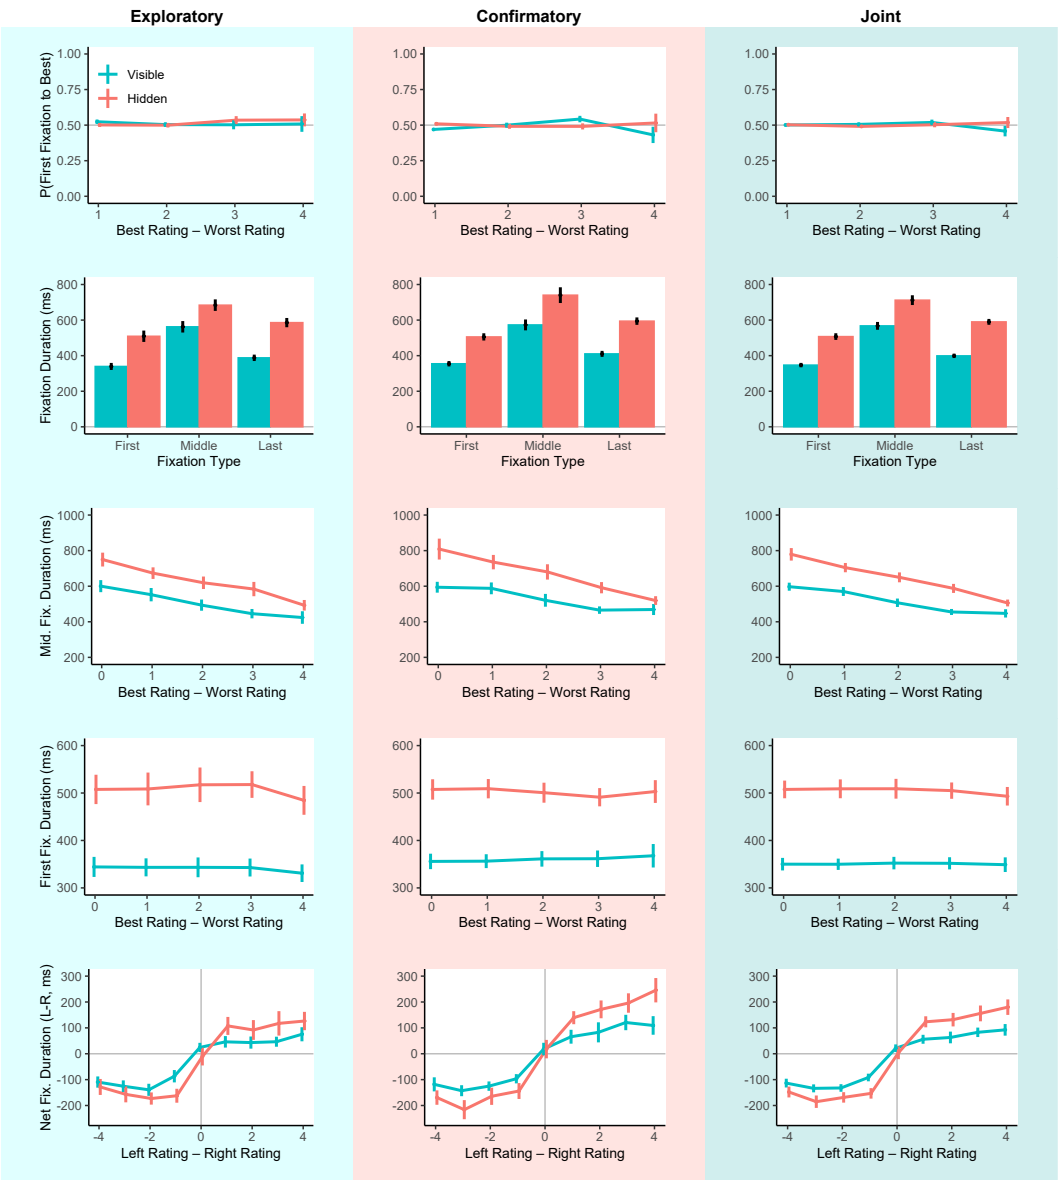

Figure 1.4: Fixation Properties. (Row 1) The probability that the first fixation is to the best item as a function of choice difficulty. (Row 2) Fixation durations by fixation type. (Row 3) Middle fixation duration as a function of choice difficulty. (Row 4) First fixation duration as a function of choice difficulty. (Row 5) Net fixation duration to the left item as a function of its relative value. Columns indicate which dataset generated the figures. Error bars show standard errors of the mean across participants.

Choice Biases

Fig. 1.5 and Table A.5 depict the attentional bias in both conditions. The goal here is to provide a model-free test of the extent to which removing nonfixated items affects attentional biases.

The top row depicts the probability of choosing the left item as a function of its relative rating and the location of the last fixation. In the absence of an attentional bias, the location of the last fixation should not matter and the choice curves should lie on top of each other. In contrast, and consistent with previous studies (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Krajbich, Lu, et al., 2012; S. M. Smith and Krajbich, 2018; S. M. Smith and Krajbich, 2019; Fisher, 2017; Tavares, Perona, and Rangel, 2017), we find a substantial attentional bias in the visible condition: on average, when the left and right items are equally valued, the left item is 2.5 times more likely to be chosen when the last fixation is to the left than when it is to the right. The bias is substantially larger in the hidden condition, where the left item is 5 times more likely to be chosen when the last fixation is to the left than when it is to the right.

The middle row depicts the relationship between net fixation time and the corrected probability of choice. The choice measure is corrected by subtracting from each choice observation (coded as 1 if left chosen, and 0 otherwise) the proportion with which left is chosen at each relative value. As a result, in the absence of an attentional bias, the corrected probability of choice should be 0, independent of net fixation time. In contrast, we find that shifting net fixation time towards the left item by 1 second increases its choice probability by 24% in both conditions.

The bottom row depicts the relationship between excess first fixation durations and the corrected choice probability of the first seen item, using the same correction described above. Excess first fixation duration is defined as first fixation duration minus mean first fixation duration (computed for each subject). In the absence of an attentional bias, the corrected probability should be 0 regardless of excess first fixation duration. In contrast, we find that an increase in the excess first fixation duration by 1 second increases the choice probability by about 22% in the visible condition, but that there is no such effect in the hidden condition.

aDDM

Given that the aDDM has been shown to provide good quantitative accounts of the relationship between fixations, choices, and RTs, we fit a hierarchical approximation of this model to our data, separately for the visible and hidden conditions. The goal is to investigate the impact of removing the nonfixated items on the parameters of the aDDM and the attentional biases that they predict.

Table 1.1 summarizes the maximum a posteriori (MAP) estimates for group-level

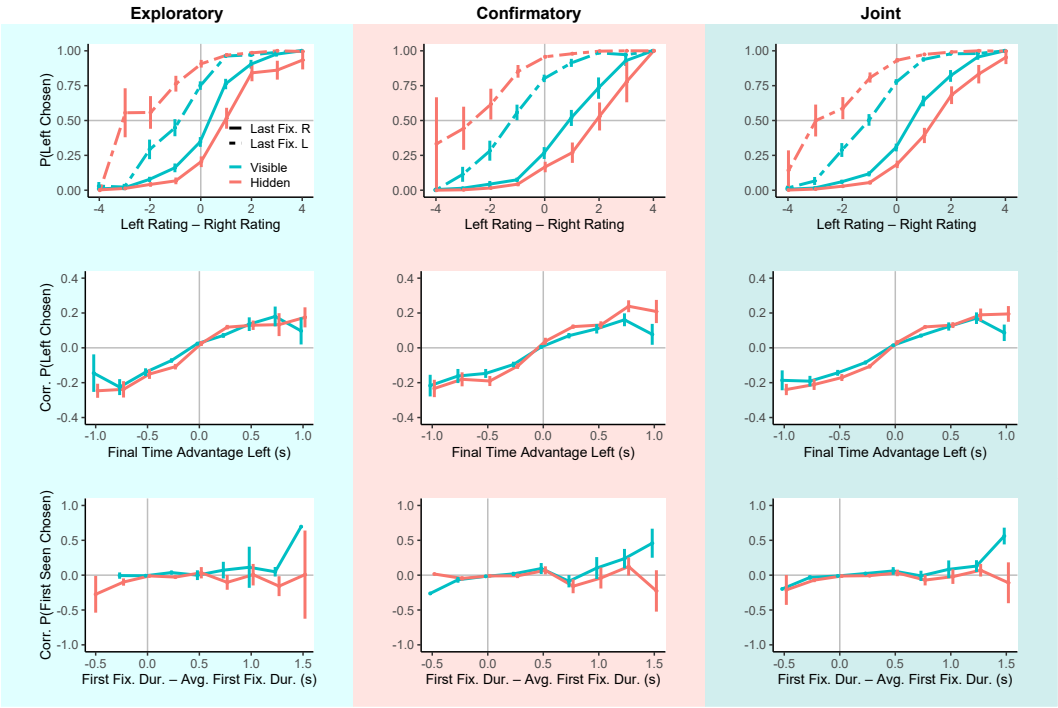

Figure 1.5: Choice Biases. (Top) Probability of choosing the left item as a function of its relative value, conditional on last fixation location. (Middle) Corrected probability of choosing the left item as a function of the net fixation time to the left item. The corrected probability is computed by subtracting from each choice observation (coded as 1 if left chosen, and 0 otherwise) the proportion with which left is chosen at each relative value. (Bottom) Corrected probability that the first seen item is chosen as a function of the excess first fixation duration, defined as first fixation duration minus mean first fixation duration (computed for each subject). Columns indicate which dataset generated the figures. Error bars show standard errors of the mean across participants.

mean parameters. We find that $\theta^V_{group} = 0.52$ and $\theta^H_{group} = 0.29$ ($\theta^V_{group} - \theta^H_{group}$ 95% $CI = [0.12, 0.35]$) which means that the attentional bias parameter in the hidden condition worsens by a factor of two, consistent with the results described above. We also find differences in the estimated parameters for the slope ($d^V_{group} = 0.003$ vs. $d^H_{group} = 0.002$, $d^V_{group} - d^H_{group}$ 95% $CI = [0.0004, 0.0009]$), and noise ($\sigma^V_{group} = 0.022$ vs. $\sigma^H_{group} = 0.017$, $\sigma^V_{group} - \sigma^H_{group}$ 95% $CI = [0.003, 0.006]$). As shown in Table A.6, the estimates using the model with correlated priors led to very similar conclusions.

	Exploratory		Confirmatory		Joint	
	H	V	H	V	H	V
d	0.002	0.003	0.002	0.003	0.002	0.003
	[0.002, 0.002]	[0.002, 0.003]	[0.002, 0.002]	[0.002, 0.003]	[0.002, 0.002]	[0.002, 0.003]
σ	0.018	0.023	0.016	0.021	0.017	0.022
	[0.016, 0.020]	[0.021, 0.025]	[0.015, 0.018]	[0.019, 0.023]	[0.016, 0.018]	[0.021, 0.023]
θ	0.38	0.54	0.20	0.51	0.29	0.52
	[0.19, 0.54]	[0.41, 0.67]	[0.02, 0.35]	[0.37, 0.63]	[0.17, 0.39]	[0.44, 0.61]
b	0.02	0.03	0.02	0.00	0.02	0.02
	[-0.07, 0.10]	[-0.05, 0.12]	[-0.07, 0.10]	[-0.08, 0.09]	[-0.03, 0.07]	[-0.03, 0.07]

MAP estimate and 95% HDI of group-level mean.

Table 1.1: Group-level MAP Parameter Estimates for Model with Uncorrelated Priors across Datasets and Conditions.

The hierarchical model also provides individual parameter estimates for each subject, which are shown in Fig. 1.6. Except for bias, the parameters in the visible condition are larger for most subjects. We estimate θ without the typical bounds at 0 and 1. In the visible condition, 0 out of 50 subject-level MAP estimates for θ fall below zero and 1 falls above one. In the hidden condition, 7 out of 50 fall below zero and 0 fall above one. However, in each of these cases, the 95% highest density intervals (HDIs) include the traditional boundaries. See Figs. A.5 and A.6 for a comparison of the out-of-sample predictions of the fitted model and the data in the even trials. See Figs. A.7 and A.8 for a comparison of the fixations and choice biases for subjects with estimated θ^H below and above zero. As shown in Fig. A.9, the model with correlated priors leads to very similar individual parameter estimates.

Mechanisms of Choice Bias

Our results show that attentional biases are approximately twice as large in the hidden condition, and that this is accompanied by a change in fixation durations, a change in the key attentional bias parameter θ, and changes in other aDDM parameters. In this section, we use out-of-sample simulations to investigate the extent to which the attentional biases are driven by changes in the fixation process, changes in θ, or changes in non-attentional model parameters.

The simulations are shown in Fig. 1.7. We start the analysis by comparing the observed and simulated attentional bias in the visible condition. To do this, we simulate 10 datasets for every subject in the out-of-sample even trials, using the empirical fixation patterns from the even trials and the aDDM parameters fitted in the odd trials of the visible condition (see Methods for details). As shown in the top panel, we find a good quantitative match between the observed and simulated data.

In row 2, we repeat the exercise by changing one component of the simulations at a

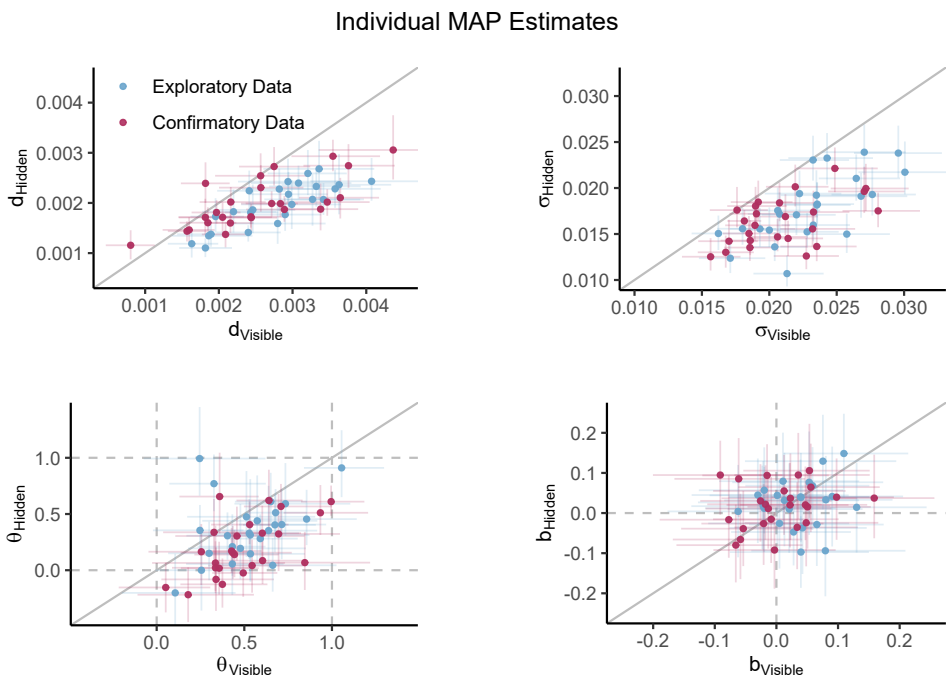

Figure 1.6: Comparison of aDDM Parameter Estimates from Model with Uncorrelated Priors. Subject-level aDDM MAP parameter estimates in the visible and hidden conditions, for both the exploratory and confirmatory datasets. Colored lines denote 95% HDIs.

time. Panel b depicts data simulated using the fixations from the visible condition, but using the values of θ fitted in the hidden trials. The panel shows that this change by itself generates a good qualitative account of the increased attentional bias in hidden trials.

Panel c depicts data simulated using the parameters fitted in the visible condition but using the fixation process from the hidden condition (ΔFix.). To clarify, when we use the fixation process from the hidden condition, we mean that we are sampling properties of the fixation process (probability of first fixation to the left, latency to first fixation, first fixation duration, middle fixation duration, saccadic duration) from their empirical distributions across the hidden trials, separately for each subject. All properties of the fixation process are independently sampled once per trial, except for middle fixation durations and saccade durations, which are independently sampled until the drift diffusion process terminates. We find that this change, by itself, has a negligible impact on the attentional bias, and thus cannot account for observed data in the hidden condition.

Panel d depicts data simulated using the fixations from the visible condition, but using the values of (d, σ, b) fitted in the hidden trials. Again, this change, by itself, is unable to provide a good qualitative account of the increased attentional bias in hidden trials.

Row 3 depicts simulations in which two of the components are changed at a time. In panel e we use the θ parameters and the fixation process from the hidden condition, but the value of the other parameters are taken from the visible condition. In panel f we use the value of the other parameters (d, σ, b) and the fixation process from the hidden condition, but the value of θ is taken from the visible condition.

Finally, the bottom row depicts a simulation in which all parameters as well as fixations are taken from the hidden condition.

A comparison of these plots shows that the model can account for the large differences in attentional choice biases as long as the change in the θ parameter is taken into account, but not otherwise. This shows that the impact on attentional biases is driven mostly by an increase in the tendency to overweight the value of fixated options.

One natural concern with this simulation analysis is that changes in the θ parameter might be correlated with changes in fixation durations, across subjects. Fig. A.10 shows that this is not the case.

For completeness, Fig. A.11 and Table A.7 show that the estimated model parameters are able to qualitatively account for the observed choice biases associated with net fixation time and excess first fixation duration (in the visible condition) in the bottom rows of Fig. 1.5. However, they are unable to account for the observed disappearance of excess first fixation bias in the hidden condition.

1.4 Discussion

Our experiment was designed to study the impact of peripheral visual information on the decision algorithm and its performance. Removing the nonfixated option has little impact on the quality of average choices, although it slows down the choice process by about 32% (or 520 ms). More importantly, we find that attentional choice biases are approximately twice as large when the nonfixated option is not shown.

The conclusion about the relative magnitude of the attentional biases in the two conditions is based on two different sets of analyses. A model free way of measuring the size of the attentional bias, that does not depend on the assumption that the aDDM

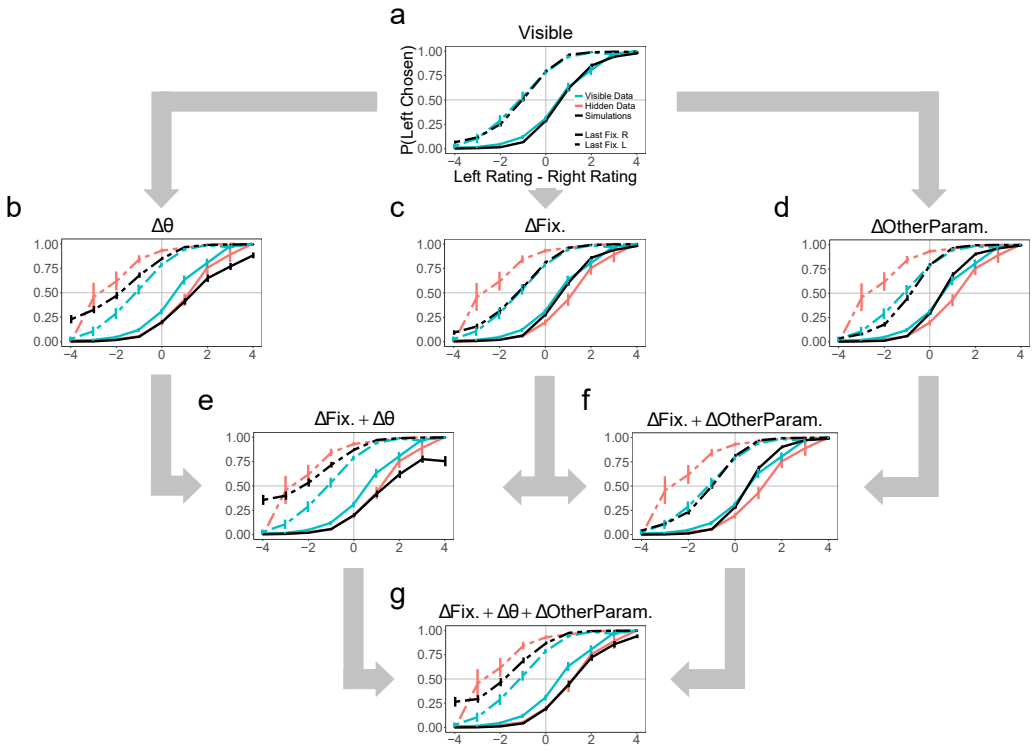

Figure 1.7: Mechanisms of Choice Bias. Probability of choosing the left item, as a function of its relative value and last fixation location, in observed (blue, red) and simulated (black) data. Each panel differs on the assumptions that were used to simulate the data. (a) Simulated data for out-of-sample even trials use the empirical fixation patterns and MAP parameters fitted in the visible condition. (b) Simulated data now uses the θ MAP parameter fitted out-of-sample in the hidden condition. (c) Simulated data uses the empirical fixation patterns from the hidden condition (ΔFix.). (d) Simulated data uses the (d, σ, b) MAP parameters fitted from the hidden condition. (e) Simulated data now uses the empirical fixation patterns and the θ MAP parameter from the hidden condition. (f) Simulated data uses the empirical fixation patterns and the (d, σ, b) MAP parameters fitted from the hidden condition. (g) Simulated data now uses fixations and all parameters from the hidden condition. The figures show that the simulations provide a good qualitative match for the difference between the visible and hidden conditions when the attentional bias parameter is modified, but not otherwise. The simulations include 10 observations per trial, per subject.

is a good description of the data generating process (Mormann and Russo, 2021), is to ask what is the probability of choosing the last fixated item when decisions have equal value (Fig. 1.5, top row). In the absence of an attentional bias, both items should be chosen with equal probability. In contrast, the last seen item is 2.5 times more likely than the other item to be chosen when all items are shown

simultaneously, and 5 times more likely when nonfixated items are hidden. Another way of measuring the attentional bias is based on the aDDM. In this model, the value of nonfixated options at any given time is downweighted by a parameter θ. When $\theta = 1$, there is no attentional bias. When $\theta < 1$, there is an attentional bias in favor of the fixated item, which is stronger for lower values of θ. Our mean estimates are $\theta = 0.52$ when all items are shown and $\theta = 0.29$ when nonfixated options are hidden. In both cases, the results show that removing peripheral visual information doubles the size of the attentional biases.

We find that middle fixations slow down by about 23% and first fixations slow down about 46% in the hidden condition, independently of the stimuli's value. There are two natural hypotheses for this change. One hypothesis, based on bottom-up control of the fixation process, is that the removal of peripheral stimuli changes the priority map that controls fixation durations and locations (Itti and Koch, 2000; Towal, Mormann, and Koch, 2013). This is consistent with findings from the visual search literature, which have found that a decrease in the saliency of peripheral stimuli, of which removal is an extreme case, increases fixation durations (Machner, Lencer, et al., 2020), as well as with the finding that fixation durations increase in patients with hemispatial neglect (Machner, Dorr, et al., 2012). An alternative hypothesis, based on top-down control of the fixation process, is that fixations slow down to accommodate the increased difficulty of generating value samples for the nonfixated stimuli in the absence of peripheral visual information.

Beyond showing that attentional choice biases increase substantially when only one item is shown at a time, our findings also provide some novel clues about the mechanisms at work in simple choice.

First, we find that in the absence of peripheral stimuli the attentional bias parameter (θ) is greater than zero on average, which means that the values of nonfixated items are still being processed by some, even if they are underweighted. This suggests that foveation facilitates the extraction of value samples, but that it is not necessary, at least after the second fixation when the identity of both stimuli becomes known. This also implies that covert attention is paid to the nonfixated item, at least after the second fixation. In fact, one interpretation of our results is that removing the nonfixated item reduces the amount of covert attention that it receives (see Carrasco (2011) for an outstanding review of the role of covert visual attention).

Second, the estimated parameters in the visible condition, and specifically the attentional bias parameter, are consistent with related literature. When fitting the aDDM

to individuals in their dataset, Krajbich, Armel, and Rangel (2010) found that the average value of attentional bias among subjects was $\theta = 0.52 \pm 0.3$, although their best-fitting model had $\theta = 0.3$. Our estimates of attentional bias in both conditions are similar to the estimated influences of gaze on choice in Weilbächer and co-authors' study (Weilbächer et al., 2021) where all options were hidden and had to be recalled from memory at the time of decision-making (attentional discounting parameter: hidden mean = 0.12, visible mean = 0.42). Interestingly, our estimate of θ in the hidden condition is larger than the one in the Weilbächer study, which suggests that the attentional bias is stronger when all information about the choice stimuli has to be recalled from memory (in their study) than when it is available conditional on foveation (as in the hidden condition in this study).

Third, Bayesian models of information sampling in simple choice have proposed that fixations matter because they control which value samples are obtained, and that samples matter because they shift the value estimates from a common initial prior to posteriors that are closer to the true value of each stimulus. As a result, the value estimates of better-than-average items tend to increase with additional fixation time, and the opposite is true for worse-than-average items (Armel, Beaumel, and Rangel, 2008; Callaway, Rangel, and Griffiths, 2021; Jang, Sharma, and Drugowitsch, 2021; Li and Ma, 2021). This Bayesian perspective could account for the increased attentional bias when nonfixated items are hidden. Value samples must be taken in parallel from both choice options, and either the rate of sampling must be slower, or the sampled information must be noisier for the nonfixated item. These variations should be even more exaggerated when nonfixated items are not present in peripheral vision. Existing Bayesian models do not account for the former, though they do account for the latter (Jang, Sharma, and Drugowitsch, 2021).

Finally, our results also have implications for the growing field of choice architecture, which seeks to understand how seemingly minor changes in the choice environment affect decisions, and how to apply this information to help individuals make better decisions (E. J. Johnson et al., 2012). We find substantially larger attentional biases in settings where only one option is shown at a time—as is done on many shopping websites—than in settings where all options are presented simultaneously, such as supermarket shelves. This suggests that individuals might be more susceptible to marketing influences that attract attention (e.g., salient packaging or point-of-sale ads) in the growing domain of e-commerce than in traditional retail settings. Although our experiments only measure the effect of removing peripheral

stimuli, similar issues could arise in contexts where choice options are described sequentially using other sensory modalities (e.g., when a waiter describes the menu specials). Extrapolating from our results, we also hypothesize that similar increases in attentional biases could be induced simply by increasing the spatial separation between stimuli, so that it becomes difficult to process nonfixated options using peripheral vision. Consistent with this hypothesis, others have found that subjects with a narrower spatial attention tend to exhibit larger attentional choice biases than those with broader spatial attention (S. M. Smith and Krajbich, 2018).

Several aspects of our study might limit the generalizability of the findings. First, our results are limited to the context of binary choice, whereas in many decision contexts more than two options are available for selection. The impact of peripheral visual information on the choice process might depend on the complexity of the environment. Second, based on previous work, we use food stimuli as a basis for understanding attentional effects on value-based choices (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011). However, it is possible that the quantitative influence of peripheral information might depend on the nature of the stimuli (e.g. lotteries, toys, concert tickets), especially if it differs on how easily it can be processed in peripheral vision. Third, in the real-world, it may be more costly and slower for consumers to switch between different options than it is in our simple gaze-contingent paradigm. For instance, consumers may have to walk between two different shelves at a super market or click through a list online, whereas in our paradigm they simply need to fixate between two regions of interest. The impact of such variables in choices needs further study.

References

Armel, Carrie, Aurelie Beaumel, and Antonio Rangel (2008). "Biasing simple choices by manipulating relative visual attention". In: *Judgment and Decision Making* 3.5.

Bürkner, Paul-Christian (2017). "brms: An R Package for Bayesian Multilevel Models Using Stan". In: *Journal of Statistical Software* 80, pp. 1–28. DOI: 10.18637/jss.v080.i01.

– (2018). "Advanced Bayesian Multilevel Modeling with the R Package brms". In: *The R Journal* 10.1, pp. 395–411.

Callaway, Frederick, Antonio Rangel, and Thomas L. Griffiths (2021). "Fixation patterns in simple choice reflect optimal information sampling". In: *PLOS Computational Biology* 17.3, e1008863. DOI: 10.1371/journal.pcbi.1008863.

Carrasco, Marisa (2011). "Visual attention: The past 25 years". In: *Vision Research* 51.13, pp. 1484–1525. DOI: 10.1016/j.visres.2011.04.012.

Cavanagh, James F. et al. (2014). "Eye Tracking and Pupillometry are Indicators of Dissociable Latent Decision Processes". In: *Journal of experimental psychology. General* 143.4, pp. 1476–1488. DOI: 10.1037/a0035813.

Cavanagh, Sean E. et al. (2019). "Visual fixation patterns during economic choice reflect covert valuation processes that emerge with learning". In: *Proceedings of the National Academy of Sciences* 116.45, pp. 22795–22801. DOI: 10.1073/pnas.1906662116.

Fisher, Geoffrey (2017). "An attentional drift diffusion model over binary-attribute choice". In: *Cognition* 168, pp. 34–45. DOI: 10.1016/j.cognition.2017.06.007.

Folke, Tomas et al. (2016). "Explicit representation of confidence informs future value-based decisions". In: *Nature Human Behaviour* 1.1, pp. 1–8. DOI: 10.1038/s41562-016-0002.

Franco-Watkins, Ana M. and Joseph G. Johnson (2011). "Applying the decision moving window to risky choice: Comparison of eye-tracking and mouse-tracing methods". In: *Judgment and Decision Making* 6.8, pp. 740–749. DOI: 10.1017/S1930297500004174.

Ghaffari, Minou and Susann Fiedler (2018). "The Power of Attention: Using Eye Gaze to Predict Other-Regarding and Moral Choices". In: *Psychological Science* 29.11, pp. 1878–1889. DOI: 10.1177/0956797618799301.

Gluth, Sebastian, Nadja Kern, et al. (2020). "Value-based attention but not divisive normalization influences decisions with multiple alternatives". In: *Nature Human Behaviour* 4.6, pp. 634–645. DOI: 10.1038/s41562-020-0822-0.

Gluth, Sebastian, Mikhail S Spektor, and Jörg Rieskamp (2018). "Value-based attentional capture affects multi-alternative decision making". In: *eLife* 7. Ed. by Michael J Frank, e39659. DOI: 10.7554/eLife.39659.

Gold, Joshua I. and Michael N. Shadlen (2007). "The Neural Basis of Decision Making". In: *Annual Review of Neuroscience* 30.1, pp. 535–574. DOI: 10.1146/annurev.neuro.29.051605.113038.

Hare, Todd, J. Malmaud, and Antonio Rangel (2011). "Focusing Attention on the Health Aspects of Foods Changes Value Signals in vmPFC and Improves Dietary Choice". In: *Journal of Neuroscience* 31.30, pp. 11077–11087. DOI: 10.1523/JNEUROSCI.6383-10.2011.

Itti, Laurent and Christof Koch (2000). "A saliency-based search mechanism for overt and covert shifts of visual attention". In: *Vision Research* 40.10, pp. 1489–1506. DOI: 10.1016/S0042-6989(99)00163-7.

Jang, Anthony I, Ravi Sharma, and Jan Drugowitsch (2021). "Optimal policy for attention-modulated decisions explains human fixation behavior". In: *eLife* 10, e63436. DOI: 10.7554/eLife.63436.

Johnson, Eric J. et al. (2012). "Beyond nudges: Tools of a choice architecture". In: *Marketing Letters* 23.2, pp. 487–504. DOI: 10.1007/s11002-012-9186-1.

Krajbich, Ian, Carrie Armel, and Antonio Rangel (2010). "Visual fixations and the computation and comparison of value in simple choice". In: *Nature Neuroscience* 13.10, pp. 1292–1298. DOI: 10.1038/nn.2635.

Krajbich, Ian, Dingchao Lu, et al. (2012). "The Attentional Drift-Diffusion Model Extends to Simple Purchasing Decisions". In: *Frontiers in Psychology* 3. DOI: 10.3389/fpsyg.2012.00193.

Krajbich, Ian and Antonio Rangel (2011). "Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions". In: *Proceedings of the National Academy of Sciences* 108.33, pp. 13852–13857. DOI: 10.1073/pnas.1101328108.

Kunar, Melina A. et al. (2017). "The influence of attention on value integration". In: *Attention, Perception, & Psychophysics* 79.6, pp. 1615–1627. DOI: 10.3758/s13414-017-1340-7.

Li, Zhi-Wei and Wei Ji Ma (2021). "An uncertainty-based model of the effects of fixation on choice". In: *PLOS Computational Biology* 17.8, e1009190. DOI: 10.1371/journal.pcbi.1009190.

Lombardi, Gaia and Todd Hare (2021). *Piecewise constant averaging methods allow for fast and accurate hierarchical Bayesian estimation of drift diffusion models with time-varying evidence accumulation rates*. preprint. DOI: 10.31234/osf.io/5azyx.

Machner, Björn, Michael Dorr, et al. (2012). "Impact of dynamic bottom-up features and top-down control on the visual exploration of moving real-world scenes in hemispatial neglect". In: *Neuropsychologia* 50.10, pp. 2415–2425. DOI: 10.1016/j.neuropsychologia.2012.06.012.

Machner, Björn, Marie C. Lencer, et al. (2020). "Unbalancing the Attentional Priority Map via Gaze-Contingent Displays Induces Neglect-Like Visual Exploration". In: *Frontiers in Human Neuroscience* 14, p. 41. DOI: 10.3389/fnhum.2020.00041.

Mormann, Milica and J. Edward Russo (2021). "Does Attention Increase the Value of Choice Alternatives?" In: *Trends in Cognitive Sciences* 25.4, pp. 305–315. DOI: 10.1016/j.tics.2021.01.004.

Pärnamets, Philip et al. (2015). "Biasing moral decisions by exploiting the dynamics of eye gaze". In: *Proceedings of the National Academy of Sciences* 112.13, pp. 4170–4175. DOI: 10.1073/pnas.1415250112.

Perkovic, Sonja et al. (2022). "Covert attention leads to fast and accurate decision-making." In: *Journal of Experimental Psychology: Applied*. DOI: 10.1037/xap0000425.

Peschel, Anne O., Jacob L. Orquin, and Simone Mueller Loose (2019). "Increasing consumers' attention capture and food choice through bottom-up effects". In: *Appetite* 132, pp. 1–7. DOI: 10.1016/j.appet.2018.09.015.

Ratcliff, Roger and Gail McKoon (2008). "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks". In: *Neural Computation* 20.4, pp. 873–922. DOI: 10.1162/neco.2008.12-06-420.

Ratcliff, Roger, Philip L. Smith, et al. (2016). "Diffusion Decision Model: Current Issues and History". In: *Trends in Cognitive Sciences* 20.4, pp. 260–281. DOI: 10.1016/j.tics.2016.01.007.

Sepulveda, Pradyumna et al. (2020). "Visual attention modulates the integration of goal-relevant evidence and not value". In: *eLife* 9, e60705. DOI: 10.7554/eLife.60705.

Shimojo, Shinsuke et al. (2003). "Gaze bias both reflects and influences preference". In: *Nature Neuroscience* 6.12, pp. 1317–1322. DOI: 10.1038/nn1150.

Simion, Claudiu and Shinsuke Shimojo (2006). "Early interactions between orienting, visual sampling and decision making in facial preference". In: *Vision Research* 46.20, pp. 3331–3335. DOI: 10.1016/j.visres.2006.04.019.

Smith, Stephanie M. and Ian Krajbich (2018). "Attention and choice across domains." In: *Journal of Experimental Psychology: General* 147.12, pp. 1810–1826. DOI: 10.1037/xge0000482.

– (2019). "Gaze Amplifies Value in Decision Making". In: *Psychological Science* 30.1, pp. 116–128. DOI: 10.1177/0956797618810521.

Tavares, Gabriela, Pietro Perona, and Antonio Rangel (2017). "The Attentional Drift Diffusion Model of Simple Perceptual Decision-Making". In: *Frontiers in Neuroscience* 11, p. 468. DOI: 10.3389/fnins.2017.00468.

Thomas, Armin W. et al. (2019). "Gaze bias differences capture individual choice behaviour". In: *Nature Human Behaviour* 3.6, pp. 625–635. DOI: 10.1038/s41562-019-0584-8.

Towal, R. Blythe, Milica Mormann, and Christof Koch (2013). "Simultaneous modeling of visual saliency and value computation improves predictions of economic choice". In: *Proceedings of the National Academy of Sciences* 110.40. DOI: 10.1073/pnas.1304429110.

Wästlund, Erik, Poja Shams, and Tobias Otterbring (2018). "Unsold is unseen ... or is it? Examining the role of peripheral vision in the consumer choice process using eye-tracking methodology". In: *Appetite* 120, pp. 49–56. DOI: 10.1016/j.appet.2017.08.024.

Weilbächer, Regina Agnes et al. (2021). “The influence of visual attention on memory-based preferential choice”. In: *Cognition* 215, p. 104804. DOI: 10.1016/j.cognition.2021.104804.

Chapter 2

ATTENTION IN AVERSIVE CHOICE: EVIDENCE FOR REFERENCE-DEPENDENT SEQUENTIAL SAMPLING

2.1 Introduction

Literature

Everyday, we make decisions about things that we want to consume or experience. What do I want to eat for dinner? What product do I want to buy? What show do I want to watch? Previous studies have established the existence of robust attentional choice biases that nudge choices towards the attended option during these types of decisions (S. M. Smith and Krajbich, 2018; Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Tavares, Perona, and Rangel, 2017).

Among these daily decisions, we also frequently encounter choices that we have to make, but where the outcomes are not pleasant. Which bill should I pay off first? Do I really need to repair my shaky car? Must I vote for one of these candidates? In these aversive scenarios, it's not clear whether attention still nudges choices in the direction of the attended option. A goal of this paper is to establish a clear relationship between choices, response times, and fixation patterns in aversive choices.

Previous work has demonstrated that manipulating attention in favor of an option biases choices towards that option (Tavares, Perona, and Rangel, 2017), and that this effect reverses in aversive choices with forced fixations and fixed response times (Armel, Beaumel, and Rangel, 2008). However, by forcing decision-makers into a forced sequential presentation format, it's possible that the authors altered the choice process in a manner that encourages these attentional bias reversals. For instance, it encourages separate evaluations of the options and discourages their comparison, which is more likely to induce attentional choice bias reversals than an environment with free response time (Eum, Dolbier, and Rangel, 2023; Basu and Savani, 2017; Basu and Savani, 2019).

In this paper, we develop a variation of the Attentional Drift-Diffusion-Model (Krajbich, Armel, and Rangel, 2010) called the Hybrid Attentional Drift-Diffusion-Model to quantitatively explain the role of attention in aversive choice. The aDDM and its variations provide a framework to investigate the role of attention during the

decision-making process, while taking into consideration other factors like noise and response caution (Krajbich, Armel, and Rangel, 2010; S. M. Smith and Krajbich, 2018). The attentional parameter in the aDDM yields a quantitative measurement of the effect of attention on the decision process.

Previous estimates of the attentional parameter in the aDDM have shown that individuals discount the value of the nonfixated option during the value comparison process (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Krajbich, Lu, et al., 2012; S. M. Smith and Krajbich, 2018; S. M. Smith and Krajbich, 2019; Tavares, Perona, and Rangel, 2017; Eum, Dolbier, and Rangel, 2023). Based on these estimates (and assuming that participants are sampling the expected values of the lotteries), we hypothesized that in simple aversive choices, attention to an option would push choices towards the *unattended* option. This is because a discounted negative value is actually *larger* than its original value, and therefore discounting the nonfixated option would make it seem better than it actually is. See Fig. 2.1 for an example. Behaviorally, this would imply that decision-makers should be (a) less likely to choose the option they attended to longer and (b) less likely to choose the option that they last fixated on.

We test these hypotheses in two studies involving simple, risky choices. In both studies, participants made a series of binary choices between two lotteries, separated into blocks based on 2 conditions: (a) a gain condition in which all lotteries had weakly positive outcomes and (b) a loss condition in which all lotteries had weakly negative outcomes. Study 1 represented lotteries with 100 colored dots in a circle to represent probabilities over fixed outcomes, and study 2 presented information about the lotteries numerically.

Previous studies have investigated the role of attention in risky choices. Most of them only involved positive outcomes and have found that the lottery that received more attention is more often selected (Fiedler and Glöckner, 2012; S. M. Smith and Krajbich, 2018; Stewart, Hermens, and Matthews, 2016; Alsharawy et al., 2021), echoing the predictions of the aDDM in choices between gains. However, contrary to the predictions of the aDDM in losses, Brandstätter and Körner (2014) found a bias towards the last fixated lottery, even if the last fixated outcome was negative. Pachur et al. (2018) document a causal relationship between attention to loss outcomes and loss aversion estimates from cumulative prospect theory (Tversky and Kahneman, 1992). Lejarraga et al. (2019) document increased attention to loss outcomes, relative to gain outcomes, even in the absence of loss aversion. Taken together

with the results from Pachur et al. (2018), this shows that increased attention to loss outcomes is a necessary, but not sufficient, condition for exhibiting loss aversion in risky choice.

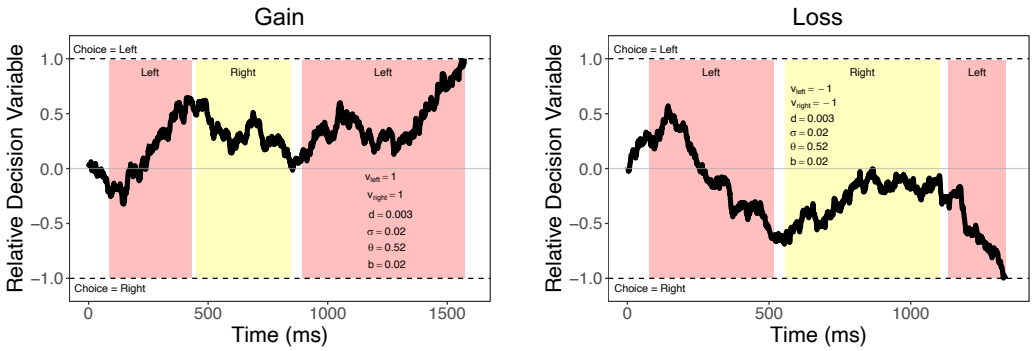

Figure 2.1: aDDM Examples. (Gain) With positive value signals and an attentional parameter between $(0, 1)$, the accumulator should be biased towards the fixated option. (Loss) With negative value signals, the accumulator should be biased towards the nonfixated option. Colored vertical bands illustrate fixation locations.

2.2 Methods

Participants

As a pre-commitment to high quality data, we filtered out participants at the data collection stage if they were missing more than 10% of fixation data. Eye-tracking data may fail to record if participants are squinting or excessively moving or blinking. In study 2, we also filtered out participants who failed more than 25% of our sanity check trials (see *Tasks* below).

A total of 160 participants were recruited from Caltech and the surrounding community using flyers. We pre-screened participants against requiring glasses for vision correction that might interfere with eye tracking. Ninety-one participants were recruited to study 1, of which 19 were excluded for failing our participant filter, leaving 72 participants (age: mean = 25 years, range = 18-41; gender: 26 male, 45 female, 1 non-binary; ethnicity: 27 Asian, 3 Black, 10 Hispanic, 3 Middle Eastern, 0 Native American, 28 White, 1 Abstain). The first 36 participants were allocated to the exploratory sample, and the other 36 to the confirmatory sample (see *Inference Strategy* below).

Sixty-nine participants were recruited to study 2, but 6 were excluded due to a change in the instructions and 13 were excluded for failing the participant filter,

leaving 50 participants (age: mean = 27.68 years, range = 18-55; gender: 13 male, 34 female, 3 non-binary; ethnicity: 17 Asian, 5 Black, 10 Hispanic, 1 Middle Eastern, 1 Native American, 16 White, 0 Abstain). Participants were equally split between the exploratory and confirmatory sample.

The number of participants and trials per participant were chosen based on related studies that have shown that this sample size provides reliable estimates of the parameters and effects of interest. All participants gave informed consent, and all experiments were approved by Caltech's Institutional Review Board.

Tasks

To investigate the role of attention in the aversive choice process, we ran two studies involving binary risky choices (see Fig. 2.2). Both studies were split equally into gain or loss conditions where participants either chose from positive- or negative-outcome lotteries.

In study 1, participants began the trial with a forced fixation cross in the center of the screen for 1 s. In the gain condition, participants were presented with two grey circles on the left and right sides of the screen. Inside each circle were 100 green or white dots. The number of white dots represented the probability of receiving nothing; the number of green dots represented the probability of gaining $10. Participants were given free response time to select the lottery they prefer with the arrow keys. After selecting, they were given feedback about their choice for 1 s. In the loss condition, trials were similar, except the green dots were replaced with red dots representing the probability of losing $10. In all trials, the probability of winning or losing was drawn from a uniform distribution over integers between 45 and 55.

Participants completed 400 trials, split into 2 blocks of 200 trials each. 1 block was in the gain condition, the other in the loss condition, order counterbalanced. After the experiment was over, participants drew a number between 1 and 200 out of an urn. The lottery associated with this number in blocks 1 and 2 was played out. Participants were paid a $40 show-up fee, and gains (losses) earned from the experiment were added (deducted) to this amount.

One potential issue with presenting lotteries as perceptual stimuli is that participants can easily switch between accumulating different types of evidence. For instance, participants might compare the ratio of green to white dots in the gain condition, but then switch to comparing the count of white dots in the loss condition. Both strategies can arrive at the optimal solution, but would have different implications

for the role of attention in the choice process. To check that this evidence switching was not happening, we ran a second study with numerical representations of the lotteries.

In study 2, participants began the trial by fixating for 1 s on a fixation cross that can appear on the left, middle, or right side of the screen. In the gain condition, participants were presented with two lotteries in green. Lotteries comprised a positive outcome and probability p of earning that outcome, with an implicit \$0 outcome with probability $1 - p$. Absolute outcomes were bounded between [\$6, \$12], and absolute differences in expected value of the lotteries were bounded between \$0 and \$4. Lottery outcomes and probabilities were designed in a way to prevent stochastic dominance of a choice. For instance, if lottery A had a better outcome than lottery B, then its probability was strictly bounded above by the probability in lottery B. Thus, choices were always between a higher amount with lower probability and a lower amount with higher probability. Participants indicated their choices using the arrow keys and were free to respond at their own pace. Upon selecting an option, feedback about the selected option was presented for 1 s.

Participants completed 340 trials, split into 4 blocks of 85 trials each. Two blocks were in the gain condition, and the others in the loss condition, with order counterbalanced. After the experiment was over, participants drew two numbers between 1 and 170 out of an urn, the first for a lottery in block 1 and the second for a lottery in block 2. The lotteries were played out, and participants were paid a \$40 show-up fee, plus (or minus) the gains (losses) earned from the experiment.

Four random trials per block were sanity check trials where there was clearly an option with higher expected value. In the gain condition, the options were (a) \$10 with $p = 1$ or (b) \$0 with $p = 1$. In the loss condition, (a) \$0 with $p = 1$ or (b) -\$10 with $p = 1$. If participants chose (b) on one of these trials, we counted it as one failed sanity check.

The 81 remaining trials that were not sanity checks were equally split into 3 attentional manipulation conditions, where the fixation cross was positioned on the left, center, or right side of the screen to manipulate where participants first fixated.

Eye-Tracking

Eye-tracking data was collected using an Eyelink 1000 sampling at 500 Hz. participants were instructed to sit approximately 60 cm from a 1920×1080 pixel monitor. Study 1 was coded in Psychtoolbox (Kleiner et al., 2007) and used lottery stimuli

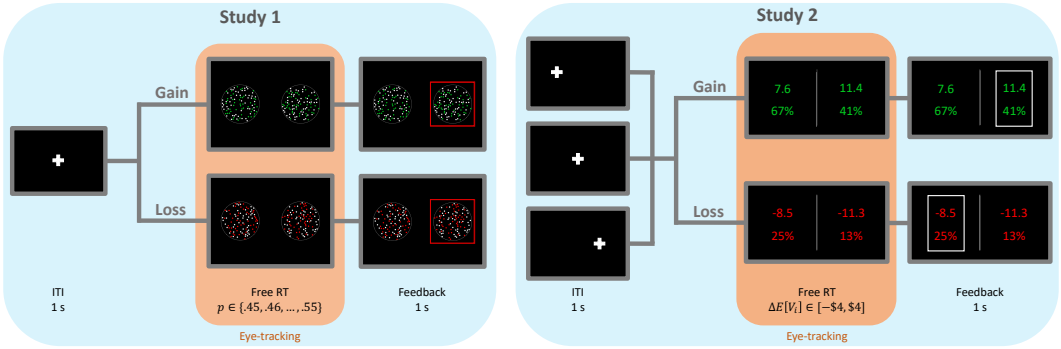

Figure 2.2: Tasks. Trials begin with a forced fixation cross for 1 s. In study 1, the fixation cross is fixed in the center; but in study 2, it varies evenly between the left, center, and right side of the screen. Participants then make a binary choice between two lotteries on the left and right. In the gain condition, the lotteries have weakly positive outcomes. In the loss condition, they have weakly negative outcomes. During this phase, participants have free response time as we record the location of their gaze. After the choice, feedback about the chosen option is presented for 1 s. In study 1, each lottery is presented as 100 dots in a grey circle. Green dots represent the probability of gaining \$10, white dots represent the probability of \$0, and red dots represent the probability of -\$10. The number of green or red dots is drawn uniformly from $p \in \{45, ..., 55\}$, and the number of white dots is $100 - p$. In study 2, information about each lottery is presented in numerical format, with differences in expected value ($\Delta E[V_i]$) bounded between $[-\$4, \$4]$.

and ROIs that were 300×300 pixels. Study 2 was coded in PsychoPy (Peirce et al., 2019) with text height set to 108 pixels and ROIs set to 384×432 pixels. For study 2, we had to use recorded response times to make a slight correction to last fixations (typically around 5 ms) due to a short lag in "stop recording" communication from PsychoPy to Eyelink. Fixations to the left ROI (lottery stimulus) were classified as "left", those to the right as "right", and outside the two ROIs as "blank". Sometimes, a sequence of blank fixations would occur between two fixations of the same type (e.g. right-blank-blank-right). These were re-coded as fixations of the same type (e.g. right-right-right-right) since these types of fixation sequences are typically due to eye-tracker noise or blinking and tend to be short. If a sequence of blank fixations was recorded between different types of fixations (e.g. left-blank-blank-right), then they were coded as a saccade period between fixations.

Inference Strategy

As an alternative to pre-registration, we instead collect two separate datasets. The first (exploratory) dataset is used to explore the data and pin down our analyses. We

then attempt to replicate the results from our exploratory dataset in the second (confirmatory) dataset. In the spirit of meta-analysis, if the results from the exploratory and confirmatory datasets are similar, we also provide results on the pooled sample and describe summary statistics in terms of the pooled estimates.

Computational Model

We develop a variation of the Attentional Drift-Diffusion-Model ("aDDM") by Krajbich, Armel, and Rangel (2010), which we call the Hybrid aDDM. Like the aDDM, value sampling is still affected by the location of one's gaze. However, the role of attention in the hybrid model is split into two parts.

In the Hybrid aDDM, decision-makers integrate noisy value signals into an accumulator that evolves over time, RDV_t. The accumulator initializes at 0, $RDV_0 = 0$, ignoring any biases towards choosing options presented on the left or right sides of the screen. Once RDV_t crosses one of two pre-specified boundaries (± 1), a choice is made based on the identity of the boundary. Hitting the upper boundary indicates a choice for the left option. Hitting the lower boundary indicates a choice for right. The evolution of this accumulator looks like the following diffusion process:

$$RDV_t = RDV_{t-1} + \mu_t + \epsilon_t \tag{2.1}$$

where ϵ_t is i.i.d. white Gaussian noise with variance σ^2. At every point in time, the evidence in the diffusion process depends on the location of the decision-maker's gaze and is given by:

$$\mu_t := \begin{cases} d[f(V_L) - \theta f(V_R)] + \eta & \text{if looking left at time } t \\ d[\theta f(V_L) - f(V_R)] - \eta & \text{if looking right at time } t \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

where d is the drift rate parameter controlling the speed of integration, and $f(V_i)$ is a function of the value of option $i \in \{L, R\}$. $\theta \in [0, 1]$ is an attentional discounting parameter that discounts the value of the nonfixated option. It captures the value-dependent role of attention in the choice process, and its effect on the choice process scales with the magnitude of the values. $\eta \geq 0$ is an attentional parameter that introduces a value-independent effect of attention. It nudges the accumulator at a constant rate towards the attended option, regardless of that option's value.

Importantly, the Hybrid aDDM takes fixations as exogenous to the model and is therefore agnostic to modeling the fixation process, though variations of the aDDM with endogenous fixations do exist (Gluth et al., 2020; Towal, Mormann, and Koch, 2013).

In addition to combining an additive and multiplicative effect of attention, the Hybrid aDDM also focuses on the functional form of the value signals. The value of a lottery option is the expected value over its *reference-dependent* outcomes:

$$f(V_i) = E[V_i|r] = p_i(u_i - r) \tag{2.3}$$

where $i \in \{L, R\}$, p_i is the probability of non-zero outcome u_i, and r is a reference-point selected according to a reference-point rule.

Due to its flexibility with value encoding and the role of attention, the Hybrid aDDM nests multiple variations of the Drift-Diffusion-Model ("DDM"). First, we consider the case when the reference-point rule is Status Quo, $r = 0$. If $\theta = 1$ and $\eta = 0$, the model reduces to the standard DDM (Gold and Shadlen, 2007; Ratcliff and McKoon, 2008). If $\theta < 1$ and $\eta = 0$, the model reduces to the standard aDDM (Krajbich, Armel, and Rangel, 2010; Krajbich, Lu, et al., 2012). And if $\theta = 1$ and $\eta > 0$, the model reduces to the additive model of attention (Cavanagh et al., 2014; S. M. Smith and Krajbich, 2019).

Next, we consider cases when the reference-point rule differs from Status Quo. If $\theta \leq 1$, $\eta = 0$, and the reference-point rule is $r = 2u_i - \max(u)$, then the model reduces to the goal-relevant aDDM when selecting the most disliked item from a set (Sepulveda et al., 2020). The model is also capable of incorporating other reference-point rules (see Baillon, Bleichrodt, and Spinu (2020) and O'Donoghue and Sprenger (2018)).

For the purposes of this paper, we would like to highlight two nested models in the Hybrid aDDM. The first, which we've already discussed above, is the additive model of attention ("AddDDM"). Note that if $\theta = 1$, any translation of the value function will result in the same value difference. Therefore the AddDDM is agnostic to any reference-point rule that applies the same translation to the outcomes of the lotteries. The second nested model we'd like to highlight is one we call the Reference-Dependent aDDM ("RaDDM"). This is the case when $\theta \leq 1$, $\eta = 0$, and a non-Status-Quo reference-point rule is applied. To hint at our results, the reference-point rule we select sets the reference-point equal to the minimum possible outcome

in a block, $r = \min(u|\text{block})$. Note that participants do not know this value at the start of a block, which affects the quality of our model fitting. See *Results* for details.

Hybrid aDDM Fitting

We split our data into training and out-of-sample data. If the trial number was a multiple of 10, then it was allocated to the out-of-sample data. All remaining trails were allocated to the training data. We fit the aDDM to the training sample using the ADDM Toolbox in Julia (Enkavi et al., 2024). The toolbox utilizes the same forward Euler estimation process as the aDDM Toolbox in Python (Tavares, Perona, and Rangel, 2017), but introduces very useful features like custom likelihood functions, discrete marginal posterior distributions for parameters, and Bayesian model comparison. It also significantly reduces computation time compared to the toolbox in Python. More details can be found at `https://addm-toolbox.github.io/dev/`.

Because evidence in the model varies moment-to-moment, the likelihood function does not have a closed-form solution, and instead needs to be numerically approximated. The algorithm used in this paper discretizes both the time space ($dt = 10$ ms) and the accumulator state space ($dx = 0.01$), allowing it to efficiently calculate the probability of crossing a decision boundary at every time-step. While the algorithm allows for gradient approaches to estimation, it is best suited for grid search in order to properly use the model comparison features. In this paper, we utilize grid search for parameter estimation. A set of estimates is pulled from a space of all parameter combinations and a likelihood is numerically approximated for that set using the algorithm described above. The grid covers the space of all reasonable parameter values with fine step sizes (see Fig. B.1). Boundaries were fixed to ± 1 and non-decision time was fixed to 100 ms. The algorithm yields point-estimates for parameters of the aDDM in a fraction of the time previously required, without relying on approximations of the accumulation process (Thomas et al., 2019; Lombardi and Hare, 2021; S. M. Smith, Krajbich, and Webb, 2019).

Out-of-Sample Simulations

We compare out-of-sample data with simulated behavior from the aDDM, fit to the training data. We simulated 10 datasets for every participant and condition, using the same lotteries encountered during the experiment. Each trial was simulated as follows. Fixation durations statistics were sampled from their empirical distributions in the out-of-sample data, conditional on the gain or loss condition. For example,

in a loss condition trial, RDV was initialized at 0, then evolved only according to the noise up to the duration of the sampled latency to first fixation. Afterwards, a maximum first fixation duration and location was sampled from the loss condition, and RDV evolved according to the drift, noise, and attentional bias estimates. If a decision boundary was crossed before the maximum duration was reached, then the process was terminated and the choice and RT were recorded. Otherwise, a saccade and middle fixation duration were sampled from their empirical distributions, and the process would repeat until a decision boundary was crossed. We assume the value of the nonfixated option is known during the first fixation, which is unrealistic and reduces the quality of our model fits.

Mixed Effects Regressions

All regressions reported in the paper are based on mixed effects regressions. Regressions were implemented in R version 4.3.1 (R Core Team, 2023) using the lme4 package (Bates et al., 2015) with a quadratic approximation optimizer ("BOBYQA") and maximum iterations set to $2e5$ where necessary.

The code package and data for this paper can be found at `rnl.caltech.edu`.

2.3 Results

Basic Psychometrics

The top row of Fig. 2.3 depicts a psychometric curve, separately for each experimental condition (color), study (line type), and dataset (column). Table B.1 contains associated statistical analysis. Choices were more sensitive to relative expected value ($\Delta E[V]$) in study 2 than in study 1. This is likely due to the narrow range of $E[V]$ differences in study 1 ($\Delta E[V] \in [-1, 1]$) compared to study 2 ($\Delta E[V] \in [-4, 4]$). We found a small but significant increase in choice sensitivity to $\Delta E[V]$ in the loss condition in study 1, but this disappears in study 2.

The middle row of Fig. 2.3 depicts response time ("RT") as a function of normalized choice difficulty. In both studies, as in previous literature, we replicated increasing RTs in choice difficulty. In study 1, RTs were not significantly different in the loss condition, regardless of choice difficulty. In study 2, RTs were on average 0.7 s slower in the loss condition. This difference shrinks at a small but significant rate as choices become easier. We believe the slower RTs in the study 2 loss condition were partially driven by the difficulty of considering an implicit \$0 outcome with a loss lottery, which many participants self-reported after the experiment concluded.

The bottom row of Fig. 2.3 depicts the number of fixations as a function of choice difficulty. As with RTs, the number of fixations did not significantly differ by condition in study 1, while they increased by roughly 0.6 on average in the loss condition in study 2.

Overall, there was a negligible difference in the quality of choices between conditions in both studies. RTs did not significantly differ between conditions in study 1, but they were significantly slower in the loss condition in study 2.

For results regarding the attentional manipulation in study 2, please refer to Appendix II.

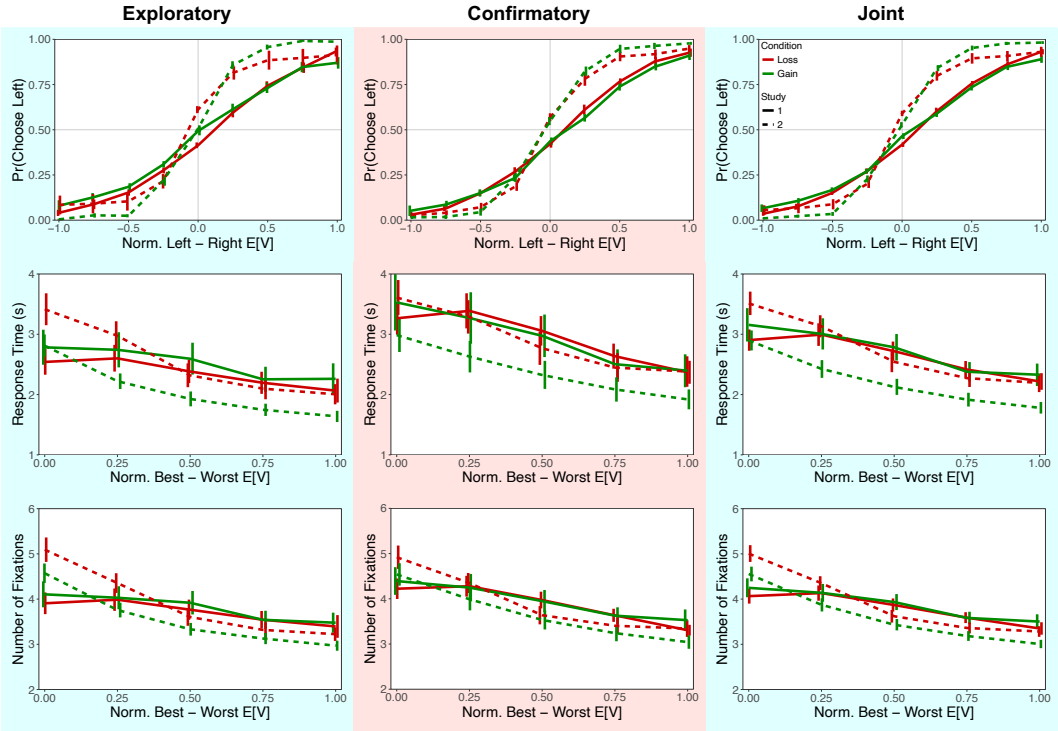

Figure 2.3: Basic Psychometrics. The top row depicts the probability of choosing the left item as a function of the normalized expected value difference between the left and right lotteries ("relative value"). Differences in expected values were normalized by dividing by the maximum magnitude of the difference. The middle row depicts response time as a function of normalized choice difficulty, measured by the expected value difference between the best and worst lottery then divided by the maximum magnitude of the difference. The bottom row depicts the number of fixations as a function of normalized choice difficulty. Columns indicate which data set generated the figures. Error bars show the standard error of the mean across participants.

Fixation Process

Fig. 2.4 and Table B.2 explore how the fixation process differs between the gain and loss condition in studies 1 and 2.

The first row of Fig. 2.4 depicts the probability of first fixating on the best lottery as a function of choice difficulty. In both conditions in both studies, this probability remained at chance regardless of choice difficulty, suggesting proper counterbalancing of the location of stimuli.

The second row of Fig. 2.4 depicts mean fixation durations in the gain and loss conditions in both studies, separately for first, middle, and last fixations. In study 1, there were no significant differences across conditions in mean fixation duration for any fixation type (First, Gain-Loss: $p = 0.29$; Middle, Gain-Loss: $p = 0.75$; Last, Gain-Loss: $p = 0.66$). But across all fixation types in study 2, fixations in the loss condition were slightly longer (First, Gain-Loss: $p = 0.02$; Middle, Gain-Loss: $p = .01$; Last, Gain-Loss: $p < 0.001$), though only by 30 to 50 ms on average (i.e. less than half a typical blink).

The third row of Fig. 2.4 depicts mean middle fixation duration as a function of choice difficulty. In both studies, middle fixation durations were not significantly sensitive to condition and slightly increased in choice difficulty.

The fourth row of Fig. 2.4 depicts mean first fixation duration as a function of the normalized expected value of the first option. In both studies, first fixation durations were not sensitive to the expected value of the first fixated lottery.

The fifth row of Fig. 2.4 depicts net fixation duration to the left lottery as a function of $\Delta E[V]$. In both studies, the relationship exhibits a significant positive relationship. In study 1, participants tended to look right a bit more in the loss condition compared to the gain condition, but this pattern did not persist in study 2. When lotteries had equal expected value, they were fixated about the same amount in study 1, but the right-side lottery received significantly more attention in study 2.

This right-side attentional favoritism occurred because participants most frequently utilized 2 fixations when the fixation cross started left, whereas they most frequently utilized 3 fixations when the fixation cross started right (see Fig. B.2 and its associated Table B.3). This phenomenon is not explained by Bayesian models of attention allocation.

Overall, there are small (but still significant) differences in fixation patterns between the gain and loss conditions in both studies. While there is significant attentional

favoritism towards the right side option in study 2, we do not see this reflected as a choice bias in the first row of Fig. 2.3, hinting that attention might be playing a smaller role in our risky choices compared to previously studied choices (S. M. Smith and Krajbich, 2018; Tavares, Perona, and Rangel, 2017).

Attentional Choice Biases

Fig. 2.5 and Table B.4 present attentional choice biases across conditions and studies. The goal here is to characterize how attention affects our choices in the loss domain by comparing these attentional biases across gains and losses.

The top row of Fig. 2.5 and Table B.4 depicts the corrected probability of choosing the left lottery as a function of the net fixation to the left lottery ("net fixation bias"). We corrected the probability by subtracting from each choice observation (1=left, 0=right) the proportion with which left was chosen at each $\Delta E[V]$. Note that in the absence of an attentional bias, this relationship should be flat. Instead, we find evidence of net fixation bias in the gain condition, consistent with previous literature. However, we predicted a reversal of this bias in the loss condition based on predictions of the aDDM (see Fig. B.3 row 1). Surprisingly, in both studies, we found no evidence of a reversal! In fact, the magnitude and direction of net fixation bias was identical across conditions.

In the bottom row of Fig. 2.5, we plot the probability of choosing the last fixated lottery as a function of the relative expected value of the last fixated lottery ("last fixation bias"). Note that in the absence of an attentional bias, the probability of choosing the last fixated lottery should cross 50% when the relative value of the last fixated lottery is 0. Replicating previous studies, we find evidence of last fixation bias in the gain condition. But once again, we do not observe a reversal of this bias in the loss condition as hypothesized.

Overall, we found identical patterns of net fixation and last fixation bias in the gain and loss conditions. This was surprising for us, given the results of previous studies and the fact that the aDDM predicts attentional choice bias reversals in the loss condition.

Model Selection

We have two goals in this section. The first is to determine what variations of the aDDM are capable of predicting the lack of attentional choice bias reversals seen in Fig. 2.5. The second is to determine if there is a reference-point rule that allows the

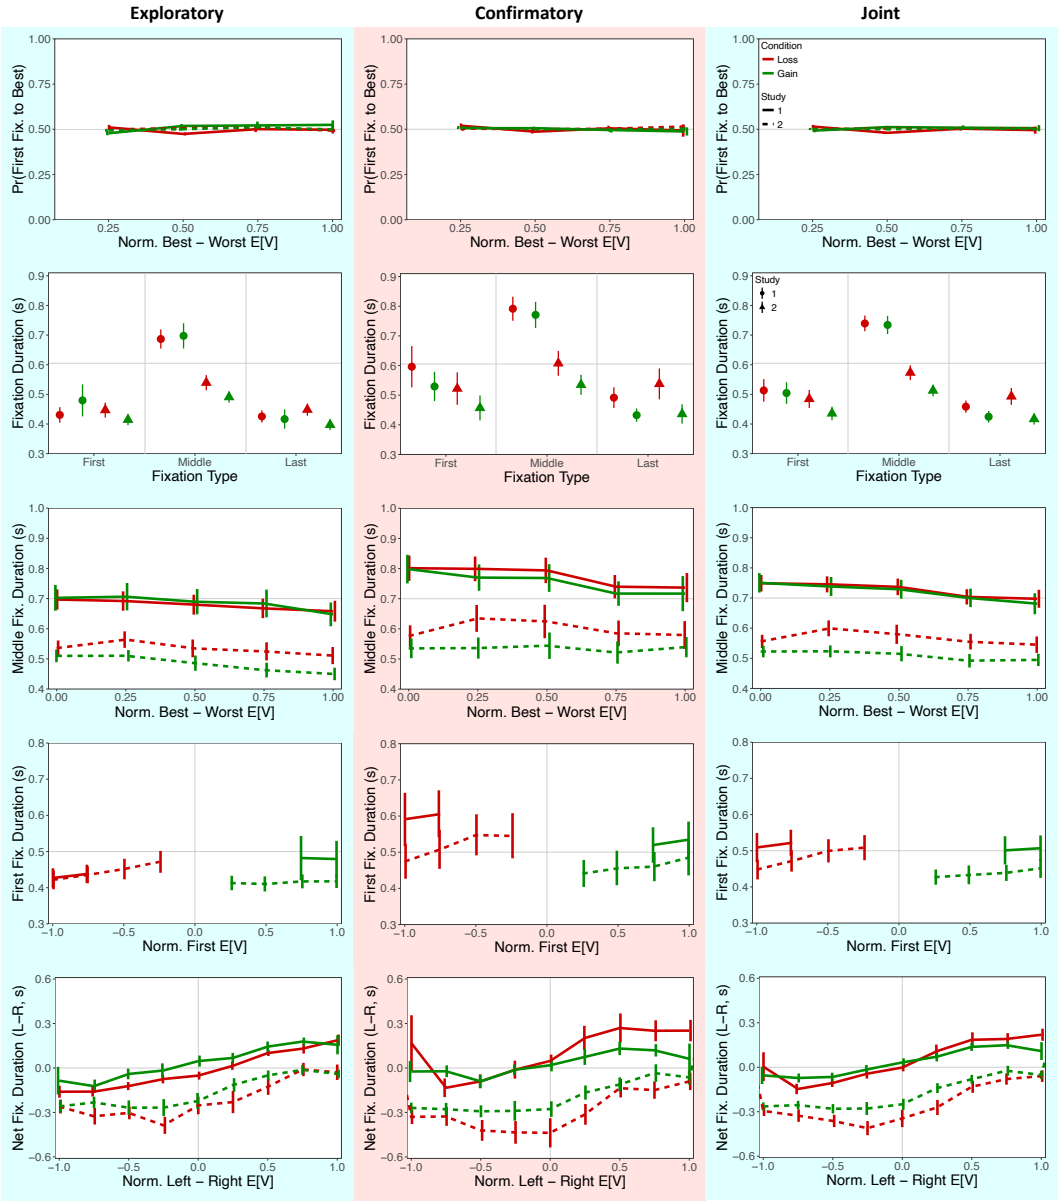

Figure 2.4: Fixation Process. The first row depicts the probability of first fixating on the best lottery as a function of normalized choice difficulty. The best lottery is determined by expected value, ignoring risk preferences. The second row depicts fixation durations as a function of fixation type. The third row depicts middle fixation durations as a function of normalized choice difficulty. The fourth row depicts first fixation durations as a function of normalized expected value of the first fixated option. The fifth row depicts net fixation duration to the left lottery as a function of its normalized relative expected value. Columns indicate which data set generated the figures. Error bars show the standard error of the mean across participants.

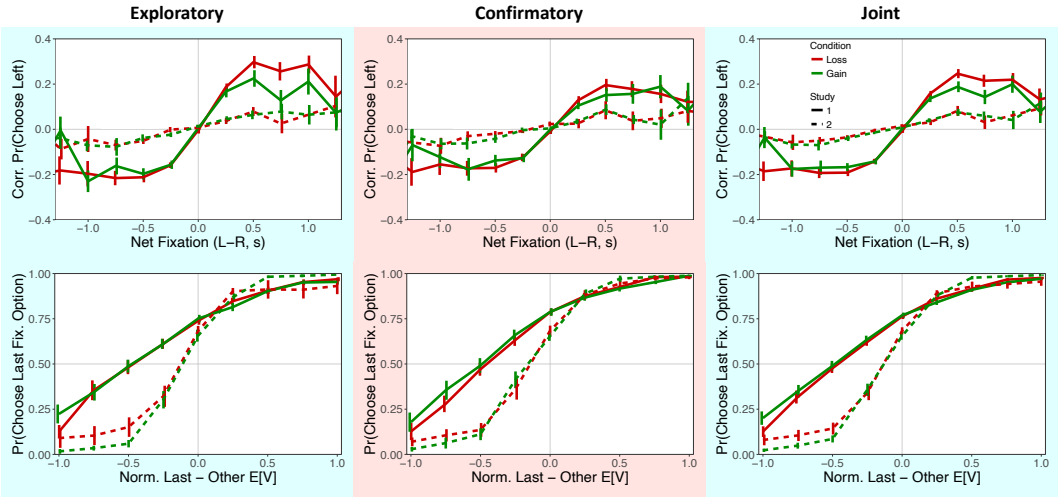

Figure 2.5: Attentional Choice Biases. The top row depicts the corrected probability of choosing the left lottery as a function of the net fixation to the left lottery. The corrected probability is computed by subtracting from each choice observation (1=left, 0=right) the proportion of which left is chosen at each relative value. The bottom row depicts the probability of choosing the last fixated lottery as a function of the normalized relative expected value of the last fixated lottery. Columns indicate which data set generated the figures. Error bars show the standard error of the mean across participants.

Hybrid aDDM or RaDDM to also make these predictions.

Having identical attentional choice biases in the gain and loss condition suggests that attention is pulling choices towards the attended option, regardless of the sign of its value. This immediately hints at a value-independent role of attention in the choice process that favors the attended option, like in the AddDDM. In row 4 of Fig. B.3, we simulate behavior using the AddDDM and find that it is capable of predicting identical attentional choice biases across our conditions. Other explanations, like (a) selecting the opposite of the predicted choice or (b) accumulating over absolute values and selecting the opposite of the predicted choice, are incapable of this (see Fig. B.3 rows 2 and 3).

An alternative explanation is that participants are range normalizing values before integrating them to make a decision (Rangel and Clithero, 2012; Frydman and Jin, 2021). This would transform negative value signals to positive signals. However, the DDM and its variations are not identifiable with range normalized value signals. Since the range in the denominators of the value signals can be factored into the drift rate, a DDM with range normalized value signals will generate the same likelihood

as a DDM with reference-dependent value signals, where the reference-point is equal to the minimum value. For our purposes, we limit our estimation to the latter case, where the reference-point rule sets the reference-point to the minimum possible outcome in a given block of trials ("Minimum-Outcome Rule"). With this reference-point rule, both the Hybrid aDDM and its nested RaDDM can predict our observed attentional choice biases (see Fig. B.3 rows 5 and 6).

This leaves us with three potential explanations of our data: (1. AddDDM) attention nudges choices towards the attended option, regardless of value; (2. RaDDM) attention reweights the reference-dependent value signals during the comparison process; (3. Hybrid) or both. Before fitting the models and doing model comparison, there is one more test we can run to distinguish explanation (1. AddDDM) from (2. RaDDM) and (3. Hybrid).

Previous studies have recorded a negative relationship between overall value ($V_L + V_R$) and RTs (S. M. Smith and Krajbich, 2019; Shevlin et al., 2022). In choices between gains, this relationship is consistent with the aDDM because of the value-dependent role of attentional discounting. For example, without loss of generality, suppose $d = 1$. Then the AddDDM would predict the same RT distribution for $V_L = V_R = 1$ as it would for $V_L = V_R = 10$, since RT is only dependent on the size of η and $(V_L - V_R)$. However, with attentional discounting, evidence will accumulate faster when $V_L = V_R = 10$ compared to when $V_L = V_R = 1$, since the value of the non-fixated option is multiplied by $\theta \in (0, 1)$ during the value comparison process.

We can apply this same logic to choices between losses. If attentional effects are value-independent like in the AddDDM, then RTs and overall value should not show any correlation. If attentional effects are value-dependent and participants are using the Status Quo reference-point rule like in the aDDM, then RTs will be decreasing as values move further away from 0 (i.e. RTs should be increasing in overall value). However, if attentional effects are value-dependent and participants are using a reference-point rule that sets the reference-point to less than or equal to the minimum value in a condition (like in the Hybrid aDDM and RaDDM), then RTs in aversive choices will be *decreasing* in overall value. Therefore, in aversive choices, the sign of the relationship between RTs and overall value serves as a distinguishing test between (a) the AddDDM and (b) the Hybrid aDDM and RaDDM. In order to look for a negative relationship between RT and overall value, we run the following regression with mixed effects, separately for gains and losses:

$$log(RT) \sim \beta_0 + \beta_1 \left| \frac{V_L - V_R}{\max(|V_L - V_R|)} \right| + \beta_2 \left(\frac{V_L + V_R}{\max(|V_L + V_R|)} \right) \tag{2.4}$$

and test if $\beta_2 < 0$. This is the same test used by S. M. Smith and Krajbich (2019), except with normalized value difference and overall value.

In Table B.5, we simulate data using the AddDDM, RaDDM, and Hybrid aDDM, then run the regression test from Eq. 2.4. Just like in gains, the AddDDM does not predict a relationship between RT and overall value in losses. However, both the Hybrid aDDM and RaDDM with the Minimum-Outcome Rule predict a negative relationship in both gains and losses.

Table 2.1 repeats this regression test in our data, separately for each study and condition. For both studies, we find a significant, negative relationship between RT and overall value in the loss condition. This evidence suggests that participants are integrating over reference-dependent value signals during the value comparison process. We also replicate the negative relationship between RTs and overall value in the study 1 gain condition, but this relationship is not significant in the study 2 gain condition. This suggests that participants may be using a reference point set to \$0 for outcomes in the gain condition, which is also the minimum possible value in those blocks. For a visual representation of the RT-overall value relationship in our data, see Fig. B.4. For replications of additional tests from S. M. Smith and Krajbich (2019), see Fig. B.5.

Loss Condition

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Study 1	β_0 Intercept	0.36	0.16	*	0.62	0.15	*	0.50	0.11	*
$log(RT)$	β_1 N. Abs. Net Value	-0.21	0.03	*	-0.29	0.04	*	-0.24	0.02	*
(Linear)	β_2 N. Overall Value	-0.41	0.14	*	-0.37	0.17	*	-0.39	0.11	*
Study 2	β_0 Intercept	0.78	0.10	*	0.94	0.10	*	0.86	0.07	*
$log(RT)$	β_1 N. Abs. Net Value	-0.55	0.04	*	-0.56	0.05	*	-0.59	0.04	*
(Linear)	β_2 N. Overall Value	-0.37	0.06	*	-0.25	0.08	*	-0.32	0.05	*

Gain Condition

Study 1	β_0 Intercept	1.24	0.17	*	1.09	0.19	*	1.17	0.12	*
$log(RT)$	β_1 N. Abs. Net Value	-0.24	0.03	*	-0.34	0.05	*	-0.29	0.03	*
(Linear)	β_2 N. Overall Value	-0.53	0.16	*	-0.18	0.15		-0.36	0.11	*
Study 2	β_0 Intercept	0.85	0.06	*	0.89	0.10	*	0.87	0.06	*
$log(RT)$	β_1 N. Abs. Net Value	-0.46	0.05	*	-0.48	0.05	*	-0.49	0.04	*
(Linear)	β_2 N. Overall Value	-0.06	0.05		-0.02	0.05		-0.04	0.04	

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
"N.": Normalized. "Abs.": Absolute.

Table 2.1: Model Comparison Regression Test

We then fit the Hybrid aDDM separately for each participant in a study-condition, using the parameter grid specified in Fig. B.1. This allows us to calculate the posterior model probabilities ("PMP") of the Hybrid aDDM, RaDDM, and AddDDM, where the PMPs sum to 1 for each participant in a study-condition. In Fig. 2.6, we compare the PMPs of all three models across individuals using a probability simplex. We plot the PMP of the RaDDM as a function of the PMP of the Hybrid aDDM. Note that the PMP of the AddDDM is equivalent to 1 minus the sum of the PMPs of the other models. This means that points close to the upper left corner of the triangle favor the RaDDM, points close to the lower right corner favor the Hybrid aDDM, and points close to the lower left corner favor the AddDDM. Consistent with our regression test results, we find that most participants are best-fit by the Hybrid aDDM in study 1, and that most participants are best-fit by either the Hybrid aDDM or RaDDM in study 2. The Hybrid aDDM is the most favored in the study 2 gain condition. Fig. B.6 shows these plots for the exploratory and confirmatory datasets.

Based on the results from our regression test and posterior model probability comparison, we posit that the risky choice process in gains and losses is best explained by the Hybrid aDDM. In the next section, we discuss model-based results using this model.

Hybrid aDDM

Fig. 2.7 presents the maximum a posteriori (MAP) estimates from the Hybrid aDDM. For each participant, we plot their estimate of $(d, \sigma, \theta, \eta)$ in the loss condition as a function of their respective estimate in the gain condition. On average, in study 1, $\eta_{1,G} = \eta_{1,L} = 0.003$, and in study 2, $\eta_{2,G} = \eta_{2,L} = 0.002$. This suggests that attention has a value-independent role in the choice process. We also find that in study 1, $\theta_{1,G} = \theta_{1,L} = 0.90$, and in study 2, $\theta_{2,G} = 0.86$ and $\theta_{2,L} = 0.91$. Taken all together, these results suggest that attention plays two roles in the choice process, split between a value-dependent effect and a value-independent effect. Because of the Minimum-Outcome reference-point rule, both of these effects nudge choices towards the attended option. These effects did not significantly differ across conditions in either study ($\theta_{1,G} - \theta_{1,L} = -0.001$, Cohen's d $= -0.006$, $t(71) = -0.04$, $p = 0.97$; $\theta_{2,G} - \theta_{2,L} = -0.053$, Cohen's d $= -0.38$, $t(49) = -1.84$, $p = 0.07$; $\eta_{1,G} - \eta_{1,L} = -8e-5$, Cohen's d $= -0.03$, $t(71) = -0.26$, $p = 0.79$; $\eta_{2,G} - \eta_{2,L} = -0.0001$, Cohen's d $= -0.07$, $t(49) = -0.39$, $p = 0.70$).

Figs. B.7 and B.8 present the MAP estimates across participants for the RaDDM

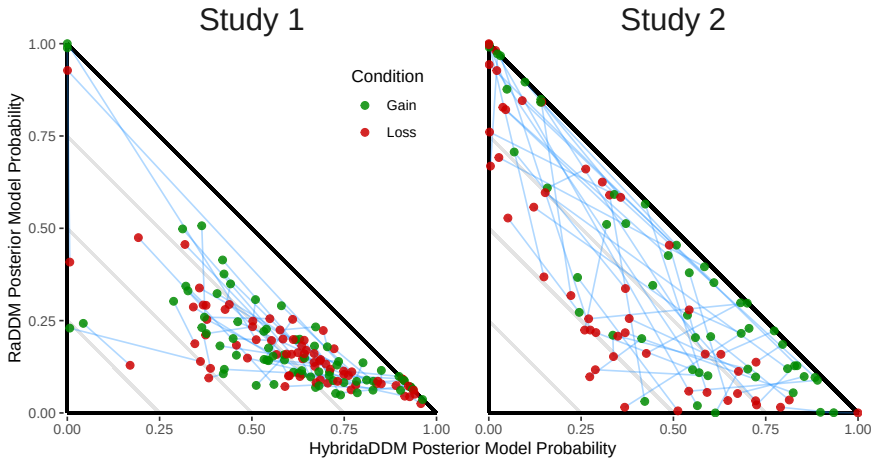

Figure 2.6: Model Comparison Between RaDDM, HybridaDDM, and AddDDM, Plotted as a Probability Simplex. For each participant, the posterior model probability of the RaDDM is plotted as a function of the posterior model probability of the HybridaDDM. The posterior model probability of the AddDDM for each individual is equal to 1 minus the sum of their posterior model probabilities for the RaDDM and HybridaDDM. Points closer to the origin indicate participants that were best-fit by the AddDDM, points closer to the top-left indicate participants that were best-fit by the RaDDM, and points closer to the bottom-right indicate participants that were best-fit by the HybridaDDM. Blue lines connect participants across conditions to visualize consistency of each participant's best-fitting model.

and AddDDM. Compared to estimates from the RaDDM and AddDDM, the Hybrid aDDM estimates θ closer to 1 and η closer to 0, suggesting lower levels of attentional bias from each of these sources ($\theta^{Hybrid}_{1,G} - \theta^{RaDDM}_{1,G} = 0.09$, Cohen's d = 0.72, $t(71) = 9.29$, $p = 7e-14$; $\theta^{Hybrid}_{1,L} - \theta^{RaDDM}_{1,L} = 0.09$, Cohen's d = 0.73, $t(71) = 9.04$, $p = 2e-13$; $\theta^{Hybrid}_{2,G} - \theta^{RaDDM}_{2,G} = 0.14$, Cohen's d = 0.94, $t(49) = 7.54$, $p = 1e-9$; $\theta^{Hybrid}_{2,L} - \theta^{RaDDM}_{2,L} = 0.08$, Cohen's d = 0.46, $t(49) = 5.86$, $p = 4e-7$; $\eta^{Hybrid}_{1,G} - \eta^{AddDDM}_{1,G} = -0.003$, Cohen's d = -0.85, $t(71) = -9.07$, $p = 2e-13$; $\eta^{Hybrid}_{1,L} - \eta^{AddDDM}_{1,L} = -0.002$, Cohen's d = -0.85, $t(71) = -11.08$, $p < 2e-16$; $\eta^{Hybrid}_{2,G} - \eta^{AddDDM}_{2,G} = -0.001$, Cohen's d = -0.67, $t(71) = -5.32$, $p = 3e-6$; $\eta^{Hybrid}_{2,L} - \eta^{AddDDM}_{2,L} = -0.0009$, Cohen's d = -0.40, $t(71) = -5.03$, $p = 7e-6$). This makes sense because the role of attention in the Hybrid aDDM is split, and both effects are nudging choices towards the attended option. In the RaDDM and AddDDM, these split effects are captured by a single attentional parameter, either θ or η, and therefore must be overestimated to account for both effects. In Fig. B.9, we show that estimates of θ from the RaDDM and η from the AddDDM are

significantly correlated across participants.

We did not observe differences in drift rates or noise across conditions, except for a small difference in drift rate in study 2 ($d_{1,G} = 0.006$ vs. $d_{1,L} = 0.006$, $d_{1,G} - d_{1,L} = 0.0001$, Cohen's d = 0.04, $t(71) = 0.34$, $p = 0.74$; $d_{2,G} = 0.003$ vs. $d_{2,L} = 0.002$, $d_{2,G} - d_{2,L} = 0.0006$, Cohen's d = 0.52, $t(49) = 3.78$, $p = 4e-4$; $\sigma_{1,G} = 0.06$ vs. $\sigma_{1,L} = 0.06$, $\sigma_{1,G} - \sigma_{1,L} = 0.002$, Cohen's d = 0.13, $t(71) = 1.73$, $p = 0.09$; $\sigma_{2,G} = 0.05$ vs. $\sigma_{2,L} = 0.05$, $\sigma_{2,G} - \sigma_{2,L} = 0.002$, Cohen's d = 0.16, $t(49) = 1.50$, $p = 0.14$). This is consistent with slower RTs found in study 2's loss condition compared to the gain condition.

Parameter recovery exercises for the Hybrid aDDM are presented in Fig. B.10 and suggest that it is capable of successfully recovering (d, σ, θ), while maintaining a positive correlation between estimated and true η.

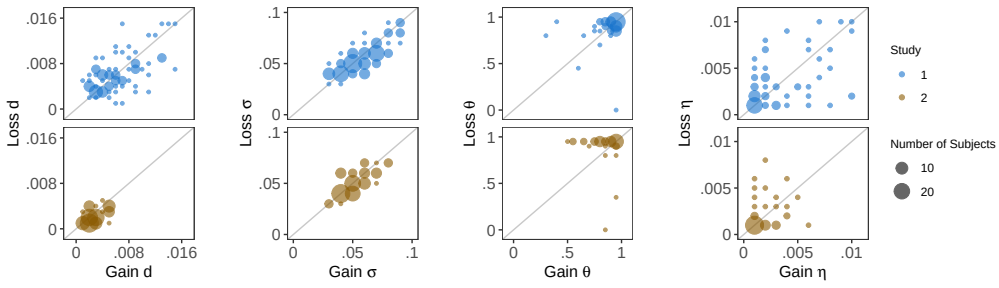

Figure 2.7: Participant-Level Hybrid aDDM Estimates in the Joint Dataset. Estimates of the drift rate (d), noise (σ), multiplicative attentional bias (θ), and additive attentional bias (η) in the loss condition as a function of their counterpart in the gain condition. Grey lines mark equality across the two conditions.

Out-of-Sample Simulations

To verify that the Hybrid aDDM is fitting our participants properly, we simulate behavior out-of-sample using the best-fitting parameters for each participant, separately for each condition. In Fig. 2.8, we compare simulated with observed out-of-sample behavior with a series of four plots: (a) the psychometric choice curve, as in Fig. 2.3 row 1; (b) the RT curve, as in Fig. 2.3 row 2; (c) the net fixation bias curve, as in Fig. 2.5 row 1; and (d) the last fixation bias curve, as in Fig. 2.5 row 2. Our simulations show that the Hybrid aDDM is able to qualitatively match the average quality of choices and the degree of net and last fixation bias exhibited in our out-of-sample data. It predicts RTs well in study 1, however it struggles to predict

RTs when value differences are small in study 2. This is typically remedied in the DDM by imposing a collapsing decision boundary (Hawkins et al., 2015; Glickman and Usher, 2019; Malhotra et al., 2017; Tajima, Drugowitsch, and Pouget, 2016; P. L. Smith and Ratcliff, 2022). Unfortunately, the current version of the ADDM Toolbox in Julia does not accurately recover the rate of collapsing bounds, so we have opted to exclude this feature from our model and leave the RT predictions as is.

Figs. B.11 and B.12 show that the RaDDM and AddDDM are also capable of qualitatively matching out-of-sample behavior.

2.4 Discussion

To investigate the role of attention in the aversive choice process, we ran two binary risky choice experiments with choices in gains and in losses. We found that the average quality of choices did not differ across conditions in either study, and that fixations were slower in losses than in gains in study 2, but not in study 1. This result is consistent with the "attention-aversion gap", where people invest more attentional resources in evaluating choices between losses than between gains, regardless of their preferences (Lejarraga et al., 2019). Attentional choice biases (net fixation bias and last fixation bias) were identical in the two conditions. This was surprising given that the standard aDDM predicts attentional choice biases would undergo a reversal when switching from the gain condition to the loss condition. The identical biases across conditions suggest that regardless of whether the choice is appetitive or aversive, attention is pulling choices in the direction of the attended item. In order to explain these findings, we propose the Hybrid aDDM with the Minimum-Outcome reference-point rule. This model transforms the outcomes of the lotteries using a reference-point set to the minimum possible value in a condition, then uses these reference-dependent outcomes to compute expected values for the value comparison process. It also splits the role of attention in the choice process into two parts: (a) a value-dependent role via an attentional discounting parameter, and (b) a value-independent role via an additive attentional parameter. Behavioral signatures (i.e. the relationship between RT and overall value) and model comparison indicate that the Hybrid aDDM is best suited to explain our data. This suggests that sequential integration is occurring over reference-dependent (e.g. range-normalized) values, and that attention may play multiple roles in the choice process.

These results have important implications for any decision involving aversive out-

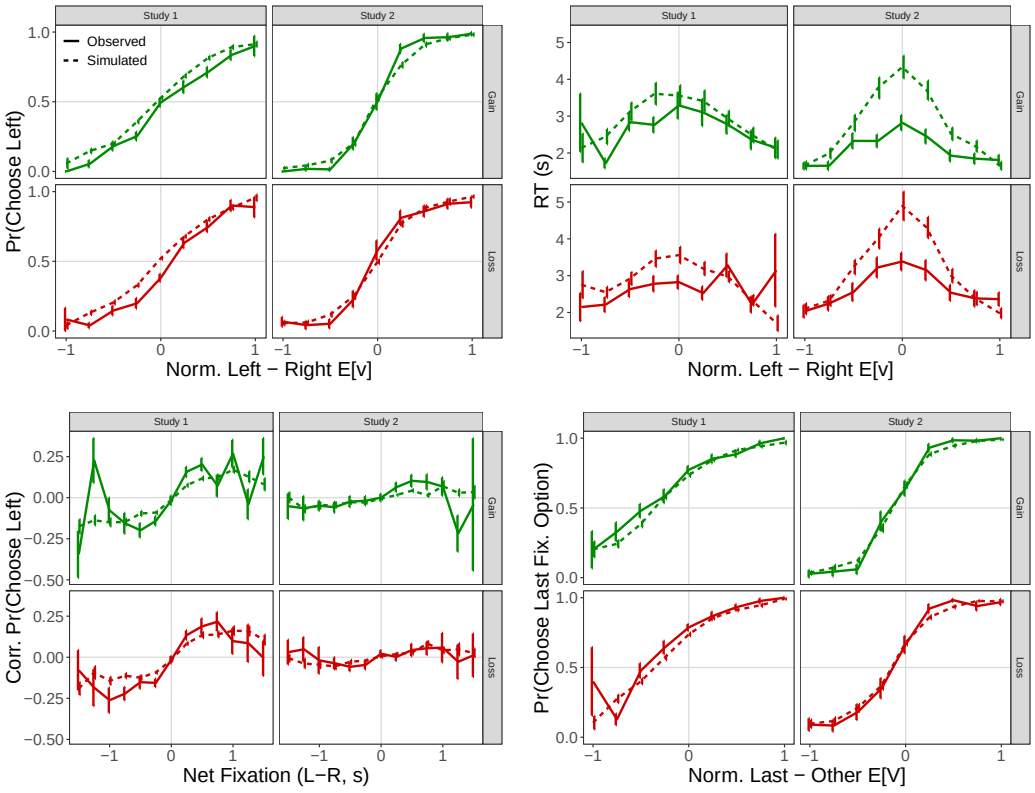

Figure 2.8: Hybrid aDDM Out-of-Sample Predictions. (Top Left) The probability of choosing left as a function of the normalized relative expected value of the left option, as in Row 1 of Fig. 2.3. (Top Right) Response time as a function of the normalized best minus worst expected value, as in Row 2 of Fig. 2.3. (Bottom Left) Corrected probability of choosing the left option as a function of the net fixation time to the left option, as in Row 1 of Fig. 2.5. The corrected probability is computed by subtracting from each choice observation (coded as 1 if left chosen, and 0 otherwise) the proportion with which left is chosen at each relative expected value. (Bottom Right) The probability of choosing the last fixated option as a function of the normalized net expected value of the left option, as in Row 2 of Fig. 2.5. Rows separate the data by condition, and columns separate by study. Out-of-sample data consists of all trials divisible by 10 from all participants in joint dataset. 10 simulated datasets per participant. Error bars represent the standard error of the mean across all simulations for all participants.

comes. Lejarraga et al. (2019) posit an evolutionary explanation for why individuals might pay more attention to losses. In an uncertain environment, the next loss could be deadly, whereas the next gain does not have an equivalent critical threshold. In other words, we pay more attention to losses because historically, choices between losses have had more severe consequences. Understanding how attention plays a role in these important, but aversive, decisions might suggest interventions that

choice architects can employ to help individuals make better decisions (Johnson et al., 2012).

Previous work has looked at choices between aversive food options and found that attention biased choices away from the fixated option (Armel, Beaumel, and Rangel, 2008). Related work has also found that attention to negative attributes biased choices away from the fixated option (Krajbich, Lu, et al., 2012; Fisher, 2021). Our results contradict these findings, though we are not the first to find evidence of last fixation bias in aversive risky choices. Brandstätter and Körner (2014) found that the last fixated option and choice coincided in nearly 78% of trials in their loss condition, but that this had nothing to do with the last fixated attribute. Taken together, these results suggest that the role of attention in aversive choices might be more nuanced than in gains, and that the role may change depending on the choice domain (e.g. food vs. product vs. intertemporal vs. risky) or the decision-maker's goals. Our proposed Hybrid aDDM is well suited to explore these nuances, but that is outside the scope of this paper.

Interpreting our modeling results, we believe that attention affects choices through both a value-dependent and a value-independent channel, and that value signals are transformed into reference-dependent value signals prior to comparison. This is consistent with previous computational and neural studies. By including the minimum possible value in a context as a potential reference-point, the Hybrid aDDM nests the hypothesis that attention aids in the accumulation of *goal-relevant* evidence that is transformed according to the goal at-hand (Sepulveda et al., 2020). There is also neural evidence that suggests avoiding aversive outcomes are encoded as rewards in the orbitofrontal cortex, a region of the human brain that has been previously implicated in encoding stimulus reward values (Kim, Shimojo, and O'Doherty, 2006). In other words, avoiding a bad outcome is encoded as a positive value in the brain. In behavioral economics, these results fall into the vast literature on reference-dependence (O'Donoghue and Sprenger, 2018; Tversky and Kahneman, 1992), though they deviate from popular models of expectations-based reference-points (Kőszegi and Rabin, 2009; Kőszegi and Rabin, 2006; Kőszegi and Rabin, 2007). Instead, they suggest a reference-point rule that takes into consideration the range of values in a context and selects the reference-point to be the minimum. Related reference-point rules have been proposed in the past (Soltani, De Martino, and Camerer, 2012). Operationally, this behaves a lot like the MaxMin reference-point rule, which Baillon, Bleichrodt, and Spinu (2020) found to be the

most common reference-point used in risky choices. In the DDM, the Minimum-Outcome reference-point rule is mathematically equivalent to range normalization of the outcomes of the lotteries, since the range in the denominators are factored into the drift rate parameter. For a review of range normalization, its biological foundations, and neural evidence, see Rangel and Clithero (2012).

There are other variations of the Hybrid aDDM that we did not consider in this paper. One of the most apparent modifications to consider would be to introduce collapsing decision boundaries in order to improve the predictions of the RT distribution in study 2. Our current model predictions indicate that the Hybrid aDDM is struggling to make difficult decisions, but does not account for any "impatience" or "urgency" that humans may exhibit when faced with the same choices (Churchland, Kiani, and Shadlen, 2008; Hawkins et al., 2015; Malhotra et al., 2017; Glickman and Usher, 2019; P. L. Smith and Ratcliff, 2022). Unfortunately, collapsing decision boundaries are not accurately recoverable in the current version of the ADDM Toolbox in Julia; therefore we opted not to include it in our model. Second, we chose not to model non-decision time (NDT) as a free parameter. This is because we manipulated the location of the fixation cross in study 2 to overlap with one of the lottery locations in two-thirds of the trials. We could allow NDT to vary trial-by-trial according to the location of the fixation cross, however it is not clear how this relationship should be modeled. Instead, we opted to fix NDT and allow for latency to first fixation to account for differences in NDT across trials. Note that in study 2, the latency to first fixation is 0 when the fixation cross overlaps with one of the lottery locations. Lastly, we opted to exclude the bias parameter from the model since our basic psychometric results did not suggest any leftward or rightward bias in participants' choices, and each additional parameter exponentially increases the size of the parameter grid and computation time.

There are other limitations to the generalizability of our findings. First, while our study is limited to binary choice, many decisions require one or more selections from 3 or more options. Second, the role of attention in the choice process may vary depending on the complexity of the environment in which a decision is made (Oprea, 2023; Gonçalves, 2024). For instance, attentional discounting of the non-fixated option might be much more severe if lottery options contained multiple outcomes instead of just two options. Third, we chose to focus on risky choices, but aversive choice extends to any form of stimulus (e.g. food, social decisions, experiences, voting). Further work is needed to ensure that these results extend to choices with

more options, more complex options, and with different stimuli.

In summary, our results describe the role of attention in the aversive choice process. Contrary to our initial hypothesis, attention to an aversive option nudges choices towards that option, despite its negative value. This is consistent with the hypothesis that individuals accumulate evidence over range-normalized (or reference-dependent) value signals during our choice process. Our computational results also suggest that attention is contributing both a value-dependent and a value-independent effect in the choice process.

References

Alsharawy, Abdelaziz et al. (2021). *Incentives Affect the Process of Risky Choice*. Rochester, NY. DOI: 10.2139/ssrn.3943681.

Armel, Carrie, Aurelie Beaumel, and Antonio Rangel (2008). "Biasing simple choices by manipulating relative visual attention". In: *Judgment and Decision Making* 3.5.

Baillon, Aurélien, Han Bleichrodt, and Vitalie Spinu (2020). "Searching for the Reference Point". In: *Management Science* 66.1, pp. 93–112. DOI: 10.1287/mnsc.2018.3224.

Basu, Shankha and Krishna Savani (2017). "Choosing one at a time? Presenting options simultaneously helps people make more optimal decisions than presenting options sequentially". In: *Organizational Behavior and Human Decision Processes* 139, pp. 76–91. DOI: 10.1016/j.obhdp.2017.01.004.

– (2019). "Choosing Among Options Presented Sequentially or Simultaneously". In: *Current Directions in Psychological Science* 28.1, pp. 97–101. DOI: 10.1177/0963721418806646.

Bates, Douglas et al. (2015). "Fitting Linear Mixed-Effects Models Using lme4". In: *Journal of Statistical Software* 67.1. DOI: 10.18637/jss.v067.i01.

Brandstätter, Eduard and Christof Körner (2014). "Attention in risky choice". In: *Acta Psychologica* 152, pp. 166–176. DOI: 10.1016/j.actpsy.2014.08.008.

Cavanagh, James F. et al. (2014). "Eye Tracking and Pupillometry are Indicators of Dissociable Latent Decision Processes". In: *Journal of experimental psychology. General* 143.4, pp. 1476–1488. DOI: 10.1037/a0035813.

Churchland, Anne K., Roozbeh Kiani, and Michael N. Shadlen (2008). "Decision-making with multiple alternatives". In: *Nature Neuroscience* 11.6, pp. 693–702. DOI: 10.1038/nn.2123.

Enkavi, Zeynep et al. (2024). *aDDM-Toolbox: Estimating time-varying drift rates with Julia and GPUs*.

Eum, Brenden, Stephanie Dolbier, and Antonio Rangel (2023). "Peripheral Visual Information Halves Attentional Choice Biases". In: *Psychological Science* 34.9, pp. 984–998. DOI: 10.1177/09567976231184878.

Fiedler, Susann and Andreas Glöckner (2012). "The Dynamics of Decision Making in Risky Choice: An Eye-Tracking Analysis". In: *Frontiers in Psychology* 3. DOI: 10.3389/fpsyg.2012.00335.

Fisher, Geoffrey (2021). "Intertemporal Choices Are Causally Influenced by Fluctuations in Visual Attention". In: *Management Science* 67.8, pp. 4961–4981. DOI: 10.1287/mnsc.2020.3732.

Frydman, Cary and Lawrence J Jin (2021). "Efficient Coding and Risky Choice". In: *The Quarterly Journal of Economics* 137.1, pp. 161–213. DOI: 10.1093/qje/qjab031.

Glickman, Moshe and Marius Usher (2019). "Integration to boundary in decisions between numerical sequences". In: *Cognition* 193, p. 104022. DOI: 10.1016/j.cognition.2019.104022.

Gluth, Sebastian et al. (2020). "Value-based attention but not divisive normalization influences decisions with multiple alternatives". In: *Nature Human Behaviour* 4.6, pp. 634–645. DOI: 10.1038/s41562-020-0822-0.

Gold, Joshua I. and Michael N. Shadlen (2007). "The Neural Basis of Decision Making". In: *Annual Review of Neuroscience* 30.1, pp. 535–574. DOI: 10.1146/annurev.neuro.29.051605.113038.

Gonçalves, Duarte (2024). *Speed, Accuracy, and Complexity*. DOI: 10.48550/arXiv.2403.11240.

Hawkins, Guy E. et al. (2015). "Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making". In: *The Journal of Neuroscience* 35.6, pp. 2476–2484. DOI: 10.1523/JNEUROSCI.2410-14.2015.

Johnson, Eric J. et al. (2012). "Beyond nudges: Tools of a choice architecture". In: *Marketing Letters* 23.2, pp. 487–504. DOI: 10.1007/s11002-012-9186-1.

Kim, Hackjin, Shinsuke Shimojo, and John P. O'Doherty (2006). "Is Avoiding an Aversive Outcome Rewarding? Neural Substrates of Avoidance Learning in the Human Brain". In: *PLOS Biology* 4.8, e233. DOI: 10.1371/journal.pbio.0040233.

Kleiner, M et al. (2007). "What's new in psychtoolbox-3". In: *Perception* 36.14, pp. 1–16.

Kőszegi, Botond and Matthew Rabin (2006). "A Model of Reference-Dependent Preferences". In: *The Quarterly Journal of Economics* 121.4, pp. 1133–1165.

– (2007). "Reference-Dependent Risk Attitudes". In: *American Economic Review* 97.4, pp. 1047–1073. DOI: 10.1257/aer.97.4.1047.

Kőszegi, Botond and Matthew Rabin (2009). "Reference-Dependent Consumption Plans". In: *American Economic Review* 99.3, pp. 909–936. DOI: 10.1257/aer.99.3.909.

Krajbich, Ian, Carrie Armel, and Antonio Rangel (2010). "Visual fixations and the computation and comparison of value in simple choice". In: *Nature Neuroscience* 13.10, pp. 1292–1298. DOI: 10.1038/nn.2635.

Krajbich, Ian, Dingchao Lu, et al. (2012). "The Attentional Drift-Diffusion Model Extends to Simple Purchasing Decisions". In: *Frontiers in Psychology* 3. DOI: 10.3389/fpsyg.2012.00193.

Krajbich, Ian and Antonio Rangel (2011). "Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions". In: *Proceedings of the National Academy of Sciences* 108.33, pp. 13852–13857. DOI: 10.1073/pnas.1101328108.

Lejarraga, Tomás et al. (2019). "The attention–aversion gap: how allocation of attention relates to loss aversion". In: *Evolution and Human Behavior* 40.5, pp. 457–469. DOI: 10.1016/j.evolhumbehav.2019.05.008.

Lombardi, Gaia and Todd Hare (2021). *Piecewise constant averaging methods allow for fast and accurate hierarchical Bayesian estimation of drift diffusion models with time-varying evidence accumulation rates*. preprint. DOI: 10.31234/osf.io/5azyx.

Malhotra, Gaurav et al. (2017). "Overcoming indecision by changing the decision boundary." In: *Journal of Experimental Psychology: General* 146.6, pp. 776–805. DOI: 10.1037/xge0000286.

O'Donoghue, Ted and Charles Sprenger (2018). "Reference-Dependent Preferences". In: *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 1. Elsevier, pp. 1–77. DOI: 10.1016/bs.hesbe.2018.07.003.

Oprea, Ryan (2023). *Simplicity Equivalents*.

Pachur, Thorsten et al. (2018). "Prospect theory reflects selective allocation of attention." In: *Journal of Experimental Psychology: General* 147.2, pp. 147–169. DOI: 10.1037/xge0000406.

Peirce, Jonathan et al. (2019). "PsychoPy2: Experiments in behavior made easy". In: *Behavior Research Methods* 51.1, pp. 195–203. DOI: 10.3758/s13428-018-01193-y.

R Core Team (2023). *R: A Language and Environment for Statistical Computing*. Version 4.3.1.

Rangel, Antonio and John A Clithero (2012). "Value normalization in decision making: theory and evidence". In: *Current Opinion in Neurobiology*. Decision making 22.6, pp. 970–981. DOI: 10.1016/j.conb.2012.07.011.

Ratcliff, Roger and Gail McKoon (2008). "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks". In: *Neural Computation* 20.4, pp. 873–922. DOI: 10.1162/neco.2008.12-06-420.

Sepulveda, Pradyumna et al. (2020). "Visual attention modulates the integration of goal-relevant evidence and not value". In: *eLife* 9, e60705. DOI: 10.7554/eLife.60705.

Shevlin, Blair R. K. et al. (2022). "High-value decisions are fast and accurate, inconsistent with diminishing value sensitivity". In: *Proceedings of the National Academy of Sciences* 119.6, e2101508119. DOI: 10.1073/pnas.2101508119.

Smith, Philip L. and Roger Ratcliff (2022). "Modeling Evidence Accumulation Decision Processes Using Integral Equations: Urgency-gating and Collapsing Boundaries". In: *Psychological review* 129.2, pp. 235–267. DOI: 10.1037/rev0000301.

Smith, Stephanie M. and Ian Krajbich (2018). "Attention and choice across domains." In: *Journal of Experimental Psychology: General* 147.12, pp. 1810–1826. DOI: 10.1037/xge0000482.

– (2019). "Gaze Amplifies Value in Decision Making". In: *Psychological Science* 30.1, pp. 116–128. DOI: 10.1177/0956797618810521.

Smith, Stephanie M., Ian Krajbich, and Ryan Webb (2019). "Estimating the dynamic role of attention via random utility". In: *Journal of the Economic Science Association* 5.1, pp. 97–111. DOI: 10.1007/s40881-019-00062-4.

Soltani, Alireza, Benedetto De Martino, and Colin Camerer (2012). "A Range-Normalization Model of Context-Dependent Choice: A New Model and Evidence". In: *PLoS Computational Biology* 8.7. Ed. by Laurence T. Maloney, e1002607. DOI: 10.1371/journal.pcbi.1002607.

Stewart, Neil, Frouke Hermens, and William J. Matthews (2016). "Eye Movements in Risky Choice". In: *Journal of Behavioral Decision Making* 29.2, pp. 116–136. DOI: 10.1002/bdm.1854.

Tajima, Satohiro, Jan Drugowitsch, and Alexandre Pouget (2016). "Optimal policy for value-based decision-making". In: *Nature Communications* 7.1, p. 12400. DOI: 10.1038/ncomms12400.

Tavares, Gabriela, Pietro Perona, and Antonio Rangel (2017). "The Attentional Drift Diffusion Model of Simple Perceptual Decision-Making". In: *Frontiers in Neuroscience* 11, p. 468. DOI: 10.3389/fnins.2017.00468.

Thomas, Armin W. et al. (2019). "Gaze bias differences capture individual choice behaviour". In: *Nature Human Behaviour* 3.6, pp. 625–635. DOI: 10.1038/s41562-019-0584-8.

Towal, R. Blythe, Milica Mormann, and Christof Koch (2013). "Simultaneous modeling of visual saliency and value computation improves predictions of economic choice". In: *Proceedings of the National Academy of Sciences* 110.40. DOI: 10.1073/pnas.1304429110.

Tversky, Amos and Daniel Kahneman (1992). “Advances in Prospect Theory: Cumulative Representation of Uncertainty”. In: *Journal of Risk and Uncertainty* 5.4, pp. 297–323.

Chapter 3

SEQUENTIAL INTEGRATION IN RISKY CHOICES FROM EXPERIENCE RESULTS IN MEAN-VARIANCE PREFERENCES

3.1 Introduction

There is evidence that sequential integration models, like the Drift-Diffusion-Model ("DDM"), provide a reasonable explanation of the psychometrics of choices and response times ("RTs") in a wide variety of economic decision-making tasks (Ratcliff, P. L. Smith, et al., 2016; Forstmann, Ratcliff, and Wagenmakers, 2016), and can even account for the role of fixations (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Krajbich, Lu, et al., 2012). There has even been some preliminary evidence linking these algorithms to neural activity in dmPFC and parietal cortex (Hare et al., 2011; Basten et al., 2010).

However, a limitation of these studies is that while the computational models assume that choices are made by integrating noisy samples, the experimenters do not have a measure of the instantaneous samples processed by the participants, which limits the ability to pin down the underlying algorithms at work.

A key goal of this study is to carry out some deeper tests of the assumptions underlying sequential-sampling algorithms, like the DDM, in canonical simple choice tasks.

Overall, our experimental design is inspired by and builds on previous experiments designed to test how individuals make risky choices from experience. In a typical experiment, participants have to choose between one of two slot-machines and get a payoff equal to a random draw from the selected machine. Before making the decision, participants are shown two parallel streams of samples taken from both machines to inform the decision, with the number of samples either fixed by the experimenter, or like here, controlled by the participant (Glickman and Usher, 2019; Tsetsos, Moran, et al., 2016; Tsetsos, Chater, and Usher, 2012).

This type of experimental paradigm has also been used extensively in behavioral economics to study the differences between risky decisions from description and risky decisions from experienced samples (Hertwig, Barron, et al., 2004; Hertwig and Erev, 2009; Wulff, Mergenthaler-Canseco, and Hertwig, 2018; Hau, Pleskac,

and Hertwig, 2010). Although the experimental paradigms have much in common, there are two important distinctions. First, in this other literature, the focus has been on lotteries with a small number of outcomes, often taken from the canonical work of Kahneman and Tversky (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). Second, the goal in this literature has been to characterize the systematic differences in choices between decisions from experience and description, and there has been less emphasis in characterizing the underlying algorithms at work at the time of decisions.

This type of task is also related to experiments in perceptual decision-making and numerical comparisons, in which participants are shown samples of numbers or dot fields and have to decide which one is larger or whether they are above or below a threshold (Irwin, W. A. S. Smith, and Mayfield, 1956; Irwin and W. A. S. Smith, 1957; Katzin, Rosenbaum, and Usher, 2021; Prat-Carrabin and Woodford, 2022a; Clarmann von Clarenau et al., 2024).

A number of previous studies have used sequential integration frameworks, including the DDM, to investigate the extent to which they can account for psychometric data (Glickman and Usher, 2019; Zeigenfuse, Pleskac, and Liu, 2014; Turkakin et al., 2023; Busemeyer, 1985).

A key aspect of our experimental design is the ability to control the value of each sample received by the participant. This allows for deeper tests of the algorithm used to make decisions, like the presence of leakage in the sequential integration process, or the use of collapsing barriers (Glickman and Usher, 2019; Malhotra et al., 2017; Trueblood et al., 2021; P. L. Smith and Ratcliff, 2022).

It also allows us to look for evidence of non-linear sample weighting. Several studies have investigated the extent to which evidence is sampled linearly, or is over- or under-weighted depending on its distance to its most likely outcome. Spitzer, Waschke, and Summerfield (2017), Clarmann von Clarenau et al. (2024), and Tickle et al. (2023) all find evidence that extreme samples are differentially weighted. Tsetsos, Moran, et al. (2016) provide some conditions in which this type of non-linear weighting of samples can improve the quality of decisions. For related results in perceptual decision-making, see Cheadle et al. (2014). For a review, see Usher, Tsetsos, et al. (2019).

To preview the results, we found that sequential sampling models can provide reasonable explanation of the psychometrics of choices and number of samples

seen in experience-based risky choices. We find evidence of either leakage or late collapsing boundaries, as well as evidence for a delay in sample integration roughly between 300 and 900 ms. We also find evidence of value compression (or inflation) that changes over the course of a trial. Finally, we establish a link between the DDM and a Modified Probit model that yields estimates for fixed decision boundaries in the space of cumulative samples.

3.2 Methods

Task

To test the assumptions of how sequential information is integrated to make risky choice, we ran an online study involving a series of slot machines (see Fig. 3.1). On each trial, one out of five slot machines was selected, and participants had to decide whether to play or skip the machine based on samples drawn from its outcome distribution. These samples did not directly affect earnings. Instead, participants had to integrate this information over time until they felt they had enough information to make a choice.

Each trial began with a 500 ms fixation cross, followed by a series of samples, each displayed one-at-a-time as a colored dot in the center of the screen for 300 ms each. Participants were free to choose play or skip using the arrow keys at any time within 12 s. If they failed to decide within this time limit, a screen would appear displaying "Faster!" for 1.5 s and no response would be recorded. If participants chose to Skip, then the reward for the trial was recorded as 0. If they chose to Play, then the reward for the trial was recorded as a random draw from the machine's outcome distribution.

Each of the five slot machines had a censored Gaussian outcome distribution with a standard deviation of 2. Their means were {-2, -1, 0, 1, 2}. Censoring occurred at ±2.875. Once a sample was drawn from the outcome distribution, two things occurred in the background before it was shown on the screen. First, its sign was used to determine the color of the dot (positive → green, negative → red). Second, its absolute value was used to determine the size of the dot. We binned absolute values to ensure dot sizes were discriminable. During piloting, we converged on 12 bins which corresponded to step sizes of 0.25 between $[0.125, 2.875]$.

Participants completed 300 trials, split into 6 blocks of 50 trials each. Each slot machine was shown exactly 60 times during the experiment, in random order. It was possible for the same slot machine to be shown in two subsequent trials. Participants

were guaranteed a show-up fee of $10. Bonus earnings were determined based on the outcomes of two randomly selected trials from each block (totalling 12 trials) and could not drop below $0. For a participant that chooses optimally, expected bonus earnings is $6.

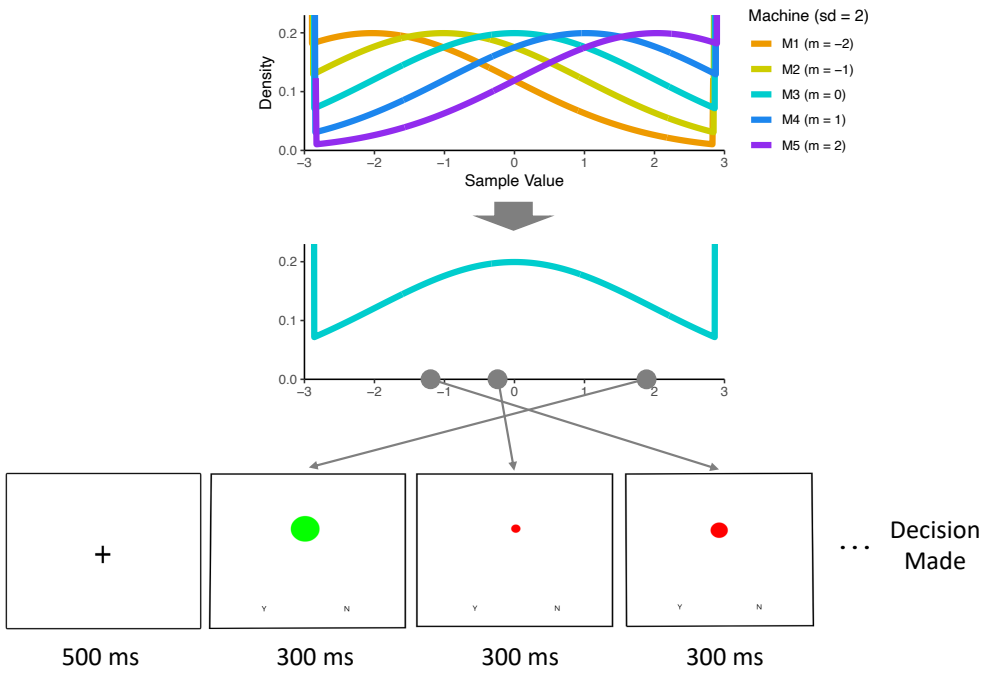

Figure 3.1: Task. Participants choose to either play or skip a slot machine based on a series of draws, representing outcomes from that machine. These draws do not affect their payoff and merely serve as examples. At the start of each trial, the computer randomly selects a machine and a fixation cross is shown for 500 ms. The participant does not know which machine was selected. Then, each draw is shown for 300 ms. The participant has free response time up to 12 seconds, after which a "Faster" screen will be shown and no response will be recorded for that trial. There are 5 machines with normally distributed values, each with a standard deviation of 2 and means in {-2, -1, . . . , 2}. Sample value distributions are censored at ±2.875. Samples are represented as green or red dots of various sizes. Color represents positive (green) or negative (red) value, and size represents magnitude (binned into 12 discrete magnitudes).

Participants

25 participants from Prolific participated in this study. We required participants to (a) be between ages 18 and 65, (b) live outside the EU or China, (c) be proficient in English, (d) have full color vision, (e) have normal or corrected-to-normal vision, and (f) have no history of epilepsy or seizures.

As a pre-commitment to high quality data, we ran participants through a quality filter immediately after the experiment finished. If a participant's data did not pass the filter, it was excluded from the final dataset. The components of our filter were: (1) participants must have responded on at least 95% of trials, (2) participants must have performed significantly better than chance, and (3) a participant's average response time must be longer than the average non-decision time under time pressure, 350 ms (estimated from Milosavljevic et al. (2010)).

In order to minimize error in stimulus display time, we required participants to complete the experiment on a laptop (i.e. no mobile devices or tablets). The typical laptop has a minimum screen refresh rate of 60 hz, which translates to 16.67 ms and represents the maximum possible error in stimulus display time. Expected error is 8.34 ms.

We selected the number of participants and trials based on previous studies that have shown that this sample size is sufficient for reliable estimation of parameters and effects of interest. Participants provided digital consent through Prolific and `Gorilla.sc`, and experiments were approved by the Caltech IRB.

Model

We reference the Drift-Diffusion-Model ("DDM") with moment-to-moment evidence, which is a sequential sampling model that has previously been successful in characterizing the relationship between choices and response times (Ratcliff and McKoon, 2008; Gold and Shadlen, 2007; Ratcliff, P. L. Smith, et al., 2016). Samples are integrated into an accumulator, RDV_t, that evolves over time until a decision boundary is hit, initiating a choice. Boundary heights are symmetric about the origin and are determined by parameter a. Hitting the upper boundary indicates a decision to Play, and hitting the lower boundary indicates a decision to Skip. The accumulator initiates at $RDV_0 = b$, incorporating a bias towards one of the choices if $b \neq 0$. The evolution of the accumulator follows the diffusion process:

$$RDV(t) = RDV(t-1) + \mu(t) + \epsilon(t) \tag{3.1}$$

where $\epsilon(t)$ is i.i.d. Gaussian white noise with variance fixed to 1. Evidence in the process is equal to the sample value at time t, $V(t)$, scaled by the drift rate parameter which controls the speed of integration:

$$\mu(t) = dV(t) \tag{3.2}$$

where d is the drift rate parameter. In the standard DDM, μ is not a function of time since evidence is typically time-homogenous. In our experiment, evidence is rapidly changing moment-to-moment ("time-heterogenous"), therefore μ must change with the evidence. Other studies investigating the choice process with time-heterogenous evidence have also made similar assumptions (Glickman and Usher, 2019; Tsetsos, Chater, and Usher, 2012; Tsetsos, Moran, et al., 2016; Tsetsos, Usher, and McClelland, 2011; Holmes, Trueblood, and Heathcote, 2016; P. L. Smith and Ratcliff, 2022).

Hierarchical Regressions

Regressions were implemented as hierarchical regressions in R version 4.3.1 (R Core Team, 2023) using the rstan package version 2.32.5 (Stan Development Team, 2024). Posterior distributions were estimated using 4 chains with 4000 warm-up iterations and 4000 sampling iterations, for a total of 16,000 samples from each of the posteriors. Priors for coefficients were normally distributed, centered at 0, with a standard deviation scaled depending on the units of the data. Priors for errors in linear regressions were Cauchy-distributed, with location = 0 and scale = 1, and strictly bounded below by 0.

Code and data for this paper can be found at `rnl.caltech.edu`.

3.3 Results

Psychometrics

Our goal in this section is to establish whether sequential sampling models are applicable to risky choices from evidence.

The top two rows of Fig. 3.2 plot the psychometric curve as a function of the machine mean and the average sample seen, respectively. Table C.1 contains associated regressions. Choices were sensitive to both the machine mean and the average sample seen. In fact, when average sample seen was not equal to 0, participants made the risk-neutral correct choice roughly 89% of the time, suggesting that participants' performance was remarkably good given the information available to them. In row 1, the psychometric curve falls below 0.5 when the machine mean is 0, which initially led us to believe that participants exhibited some form of risk aversion. However, upon closer look at row 2, the psychometric curve is not significantly away from 0.5

when average sample seen equals 0, suggesting that participants' risk attitudes were actually risk neutral.

The bottom two rows of Fig. 3.2 plot the number of samples seen in a trial as a function of the absolute machine mean and the absolute average sample seen, respectively. In row 3, the number of samples seen is decreasing in the absolute machine mean. This is mirrored in row 4 as the number of samples seen decreases in the absolute average of the sample seen. Response times can be approximated by multiplying the number of samples by 300 ms. These results are consistent with previous findings that response times (and number of fixations) increase with the difficulty of a choice (Krajbich, Armel, and Rangel, 2010; Eum, Dolbier, and Rangel, 2023).

Overall, our participants' behavior exhibits patterns that are consistent with sequential sampling models, suggesting that they are applicable to risky choices from experience.

Evidence Weighting: Temporal Dynamics

Now that we know sequential sampling models are applicable to experience-based risky choices, our next goal is to test the temporal dynamics of evidence weighting. In a DDM with fixed decision boundaries, evidence at every time point is weighted equally by the drift rate parameter. With either collapsing decision boundaries or leakage, we should see a recency effect, where later evidence is weighted more than early evidence. Opposite to this, the Bayesian observer model would predict that early samples shift posterior beliefs by much more than later samples. This would induce a primacy bias, where early evidence is weighted more than later evidence. Note that this is true despite the exchangeability property in Bayesian updating.

To test the temporal dynamics of evidence weighting, we first run the following logistic regression:

$$\text{choice} \sim \alpha_0 + \sum_{i=0}^{16} \beta_i V_{S-i} \tag{3.3}$$

where choice is encoded as (1=Play, 0=Skip), and S is the total samples seen in a trial, making V_S the last sample in a trial. If the trial did not have at least 17 samples, then V_{S-i} was set to 0 for all samples that did not exist.

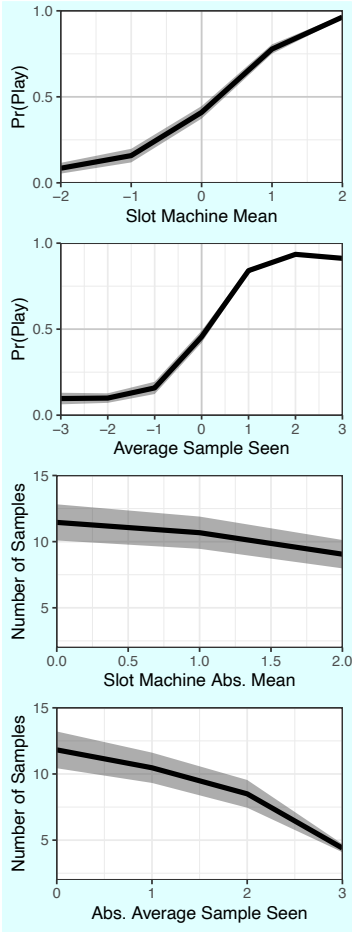

Figure 3.2: Psychometrics. Row 1 depicts the probability of choosing play as a function of the machine mean. Row 2 repeats this, except with the average sample instead of machine mean. Row 3 depicts the average number of samples seen in a trial as a function of the absolute machine mean. Row 4 repeats this, except with the absolute average sample instead of absolute machine mean. Error ribbons show the standard error of the mean across participants.

In Panel A of Fig. 3.3, we plot the posterior distributions for α_0 and all β_i from Eq. 3.3. There are a few things to take away from this plot. First, the evidence weight for the Last Sample is 0. This suggests that either (a) evidence takes time to process before being integrated into the accumulator, (b) that actions take time to be carried out even after the accumulation process has crossed a decision boundary, or (c) both. These peripheral processes are typically modeled with a non-decision time parameter in the DDM. Since our samples are only displayed for 300 ms at a time and the evidence weight for the last sample is 0, it seems that non-decision time must be bounded below by at least 300 ms.

Second, late samples (Last Sample - 3 through Last Sample - 1) have significantly higher evidence weights than samples that came before them. This violates the assumptions of static evidence weights over time in the standard DDM, but is achievable with either of two mechanisms: (1) collapsing decision boundaries or (2) leakage. With decision bounds that collapse over time, each additional sample has an increasing impact on the state of the accumulator, since non-zero evidence will shift the accumulator closer to the bounds the later it is integrated. With leakage, evidence is hyperbolic discounted according to the amount of time that has elapsed since it was integrated. This means that new samples have an exponentially larger impact on the state of the accumulator compared to older samples.

Third, the evidence weight for Last Sample - 1 is smaller than for Last Sample - 2. While at first this might seem inconsistent with collapsing bounds or leakage, this result is possible under the assumption of a delay in sample integration (as mentioned above). In Panel B of Fig. 3.3, we plot a noise-less example of the accumulation process with various amounts of delay in sample integration. Black, "Sample On" lines indicate when the last samples are displayed on the screen. Red, "Small Delay" lines indicate when each sample would enter the integration process with a small delay. Blue, "Large Delay" lines indicate the same thing, except with a larger delay. Notice that if a decision is reached during the last sample (marked by the dark grey dashed line), it is possible that this was the result of integrating Last Sample - 1 or Last Sample - 2, depending on the size of the delay. A previous study estimated this delay to be roughly 500 ms (Holmes, Trueblood, and Heathcote, 2016), which would look approximately like the large delay in Panel B. This means that some proportion of the choices are the result of integrating Last Sample - 2 (exemplified by the dark grey dashed line), and the remaining choices are the result of integrating Last Sample - 1 (exemplified by the light grey dashed line). The relative evidence weight of Last Sample - 2 compared to Last Sample - 1 will therefore depend on the proportion of choices made as a result of integrating Last Sample - 2 compared to Last Sample - 1. The longer the delay in sample integration, the larger the evidence weight for Last Sample - 2 compared to Last Sample - 1, up to a theoretical limit of 900 ms. This places limits on the duration of the delay in sample integration between 300 and 900 ms.

In Panel C of Fig. 3.3, we plot evidence weights for the sum of the first 3 samples, the sum of the last 3 samples, and the sum of remaining middle samples. In order to calculate these weights, we first subsetted our dataset to trials containing at least

7 samples. Then, we calculated these sums and weighted them by their proportion of the total number of samples seen in their respective trial. Lastly, we regressed choice on these variables, along with a constant.

Once again, we find that the last 3 samples before the last have the largest evidence weight compared to all other samples. However, we also find that the first 3 samples have a larger evidence weight compared to middle samples. This suggests a primacy effect also exists in the sampling process, although it is weaker than the recency effect. Fig. C.1 repeats the same regression as Eq. 3.3, except using the *first* 16 samples in a trial instead of the last, and including Last Sample - 3 through Last Sample - 1. It replicates that early samples have a slightly larger evidence weight compared to middle samples, but that late samples have the largest evidence weights.

Overall, these results suggest that both primacy and recency effects occur during evidence-based risky choice, but the strength of recency is larger than of primacy. They also suggest that there is a delay in sample integration between 300 and 900 ms.

Evidence Weighting: Transformations

In this section, our goal is to test if sample values are undergoing transformations during integration (Clarmann von Clarenau et al., 2024). In order to determine this, we run a series of three logistic regressions to estimate the weight of each level of sample value.

For the first regression, we round sample values to {-2.5, -1.5, -0.5, 0.5, 1.5, 2.5}, corresponding to (negative, positive) $\times$ (small, medium, large) values. Next, we generate indicator variables for each rounded value ("value") and interact these with indicator variables for early, middle, and late samples ("time"):

$$\text{choice} \sim \sum_{\text{value}} \sum_{\text{time}} \beta_{\text{value, time}} \mathbb{1}_{\{\text{sample value = value}\}} \mathbb{1}_{\{\text{sample time = time}\}} \tag{3.4}$$

Note that we exclude the constant to avoid the "dummy variable trap". The coefficient for each term estimates the sample weight during integration when it is integrated at that time. For instance, $\beta_{\text{-2.5, early}}$ is the effect that a sample of -2.5 had on the final choice when it occurred within the first three samples in a trial.

In row 1 of Fig. 3.4, we plot every $\beta_{\text{value, time}}$ from Eq. 3.4. The slope for late samples is steeper than the slope for early or middle samples, suggesting either that

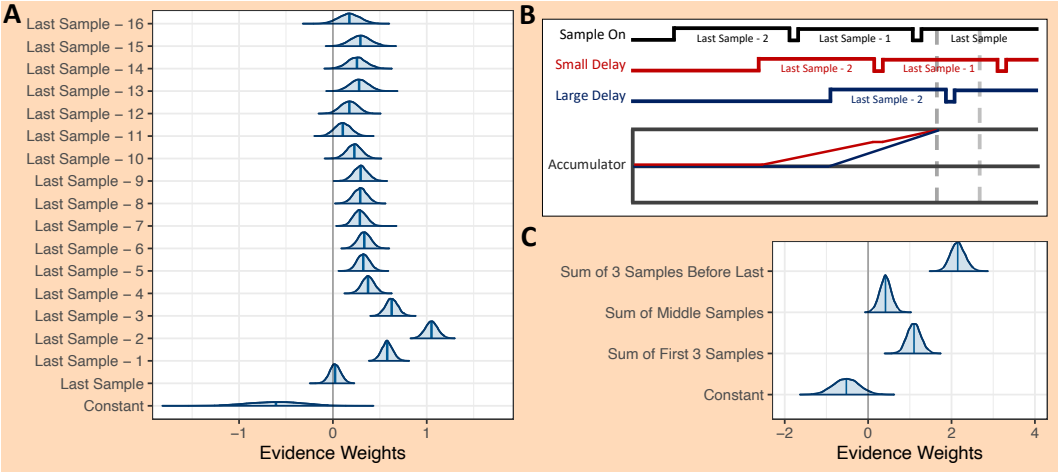

Figure 3.3: Evidence Weighting: Temporal Dynamics. (Panel A) Evidence weights for the last 17 samples. (Panel B) Visualization of how delay in sample processing can affect evidence weights. Assume all samples have the same value. "Sample On" plots when the sample is displayed on the screen. "Small Delay" plots when each sample is integrated into the accumulator, with a small delay < 300 ms. "Large Delay" plots when each sample is integrated with a large delay between 300 to 600 ms. Red accumulator line shows how noise-less accumulator would evolve according to small delay. Blue accumulator line shows how noise-less accumulator would evolve according to large delay. Dark grey dashed line indicates the time of the choice in the accumulator example. Light grey dashed line indicates the time of an example choice made as the result of integrating Last Sample - 1 under Large Delay. (Panel C) Evidence weights for the sum of the first 3 samples, the sum of the middle samples, and the sum of the last 3 samples. Sums are scaled by their proportion of the total number samples in a trial.

(a) early and middle samples undergo value compression, (b) late samples undergo value inflation, or (c) both.

For the second regression, we calculate the distance from the machine mean to the sample and round these distances to {-2.5, -1.5, -0.5, 0.5, 1.5, 2.5}, corresponding to (negative, positive) × (small, medium, large) distances. Similar to the first regression in this section, we generate indicator variables for each distance ("distance") and interact these with indicator variables for Machines 1, 2, 4, and 5 ("machine"):

$$\text{correct} \sim \sum_{\text{distance}} \sum_{\text{machine}} \beta_{\text{distance, machine}} \mathbb{1}_{\{\text{distance}\}} \mathbb{1}_{\{\text{machine}\}} \tag{3.5}$$

We excluded data from trials with Machine 3 (mean = 0), since ex-ante, there was no objectively correct decision in these trials. The coefficient for each term measures

how the rarity of a sample impacts correct choices.

In row 2 of Fig. 3.4, we plot each coefficient from Eq. 3.5 as a function of the sample distance from the machine mean. There is a cluster of sample weights that do not significantly differ from 1 towards the center of the plot. This cluster suggests that as long as the sample maintained the same sign as the machine mean, it had the same positive impact on the probability of making the correct choice, regardless of its value or the machine. Sample weights below 0 towards the sides of the plot indicate the opposite. Once a sample took on the opposite sign as the machine mean, it had a negative impact on the probability of making the correct choice. Despite mean estimates of this impact scaling in the magnitude of the distance, the scaling is not significant.

Lastly, for the third regression, we repeat most of the steps from the second regression. We subset the data to just trials with Machine 3 (mean = 0) and change the dependent variable to choice (1=Play, 0=Skip). Coefficients from this regression measure the impact of each sample distance from the machine mean (= 0) on the decision to Play.

In row 3 of Fig. 3.4, we plot the coefficients from this regression as a function of the sample distance from machine mean. The slope remains relatively constant across negative and positive distances, suggesting that any transformation to sample value is symmetric across the origin.

Overall, these results suggest that some form of value compression or inflation is occurring, separately at different times during a trial, and is symmetric about the origin. Regardless, the impact of these transformed values on making the correct decision is the same as long as the sample shares the same sign as the machine mean.

Termination Rules and Policies

The goal in this section is to determine the shape of the decision boundaries. As a bonus, we also establish a connection between the Probit model and fixed boundaries in the DDM.

To start, for every sample in a participant's trial, we rounded cumulative sample to the nearest fifth, generating roughly 25 bins. Then, for each participant, we looked at whether a choice was made once cumulative sample reached one of these bins on a given sample number. In other words, (cumulative sample, sample number) defines a state, and we ask: "Of the times that the participant has reached this state, what is the proportion with which the participant made a choice to Play or Skip in

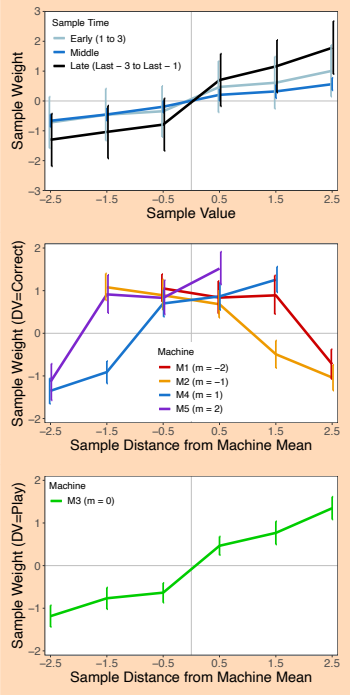

Figure 3.4: Evidence Weighting: Transformations. (Row 1) Sample weight as a function of sample value, separately for early, middle, and late samples. (Row 2) Sample weight as a function of the distance between the sample and the machine mean, for Machines 1, 2, 4, and 5. (Row 3) Sample weight as a function of the distance between the sample and the machine mean, for Machine 3 only. Error bars denote the standard error of the mean across participant means.

this state?" In Panel A of Fig. 3.5, we plot these proportions over states for four randomly sampled participants, excluding states with less than five trials. In Panel B, we repeat this exercise, except we ask: "Of the times that the participant has made a choice in this state, what is the proportion with which the participant made a choice to Play?" This time, we only exclude states with 0 trials. Panels C and D repeat these exercises from Panels A and B, respectively, except they use the average sample seen instead of cumulative sample.

In Panel A, we see that participants needed less cumulative evidence to reach a decision as sample number grew larger. This hints that participants may have grown impatient or felt some urgency during their choice process. This can be modeled in the DDM with a collapsing decision boundary.

In Panel C, we find evidence that more samples were required to make a choice when average sample was closer to 0, consistent with the results in row 4 of Fig. 3.2.

In Panels B and D, we see that participants were very accurate in their decisions,

regardless of the cumulative sample, average sample, or number of samples seen.

See Figs. C.2, C.3, C.4, and C.5 for these plots across all participants.

In each plot of Fig. 3.5, we also include estimates of the decision boundaries for the DDM, calculated using a Modified Probit regression. We're able to do this because the *RDV* from the DDM operates in the same (cumulative sample, sample number) space shown in the plots. In fact, if the DDM was noiseless, the diffusion clouds in Panel A would mimic the accumulators of the DDM across trials. Because they operate in the same space, it is possible to define an approximated fixed decision boundary using probabilistic choice models (Luce, 1959; McFadden, 1974). Below, we show a direct link between a Modified Probit model and fixed decision bounds in the DDM.

Theorem 1. *Fixed decision boundaries, a, in a DDM with i.i.d. white Gaussian noise, $\epsilon_t \sim N(0, \sigma^2)$, can be estimated via the following Modified Probit model with a dependent variable Y_t that is encoded as 1 if the participant chooses to Play or Skip at sample t and 0 if the participant chooses to Continue Sampling at sample t:*

$$\Pr(Y_t = 1|x_1, x_2, \ldots, x_t, t) = 1 - \Phi\left(\frac{a}{\sigma\sqrt{t}} - \frac{\bar{x}_t\sqrt{t}}{\sigma}\right) + \Phi\left(-\frac{a}{\sigma\sqrt{t}} - \frac{\bar{x}_t\sqrt{t}}{\sigma}\right)$$

Proof. Suppose a choice is made using the DDM on cumulative signals, $\sum_t x_t$, with a fixed decision boundary, a, and i.i.d. white Gaussian noise, $\epsilon_t \sim N(0, \sigma^2)$. Then, the decision-maker ("DM") will choose to sample another piece of evidence if:

$$-a < \sum_t x_t + \sum_t \epsilon_t < a$$

In other words, the probability that the decision maker will choose to Play or Skip at time t is given by:

$$\begin{aligned}
\Pr(Y_t = 1|x_1, \ldots, x_t, t) &= \Pr\left(\left|\sum_t x_t + \sum_t \epsilon_t\right| \geq a\right) \\
&= \Pr\left(\sum_t x_t + \sum_t \epsilon_t \geq a\right) + \Pr\left(\sum_t x_t + \sum_t \epsilon_t \leq -a\right) \\
&= \Pr\left(\sum_t \epsilon_t \geq a - \sum_t x_t\right) + \Pr\left(\sum_t \epsilon_t \leq -a - \sum_t x_t\right)
\end{aligned}$$

We assumed that $\epsilon_t \sim N(0, \sigma^2)$. Then, $\sum_t \epsilon_t \sim N(0, t\sigma^2) \Rightarrow \frac{\sum_t \epsilon_t}{\sigma\sqrt{t}} \sim N(0, 1)$. Applying this transformation to both sides of the inequalities yields

$$\begin{aligned}\Pr(Y_t = 1|x_1, \ldots, x_t, t) &= \Pr\left(\frac{\sum_t \epsilon_t}{\sigma\sqrt{t}} \geq \frac{a}{\sigma\sqrt{t}} - \frac{\bar{x}_t\sqrt{t}}{\sigma}\right) + \Pr\left(\frac{\sum_t \epsilon_t}{\sigma\sqrt{t}} \leq -\frac{a}{\sigma\sqrt{t}} - \frac{\bar{x}_t\sqrt{t}}{\sigma}\right) \\ &= 1 - \Phi\left(\frac{a}{\sigma\sqrt{t}} - \frac{\bar{x}_t\sqrt{t}}{\sigma}\right) + \Phi\left(-\frac{a}{\sigma\sqrt{t}} - \frac{\bar{x}_t\sqrt{t}}{\sigma}\right)\end{aligned}$$

□

Running the regression in Theorem 1 produces two estimated coefficients: $\hat{\gamma}_1 = a/\sigma$ and $\hat{\gamma}_2 = 1/\sigma$. In order to calculate the estimate for the fixed decision bounds, we need to multiply σ out: $\hat{a} = \hat{\gamma}_1\hat{\gamma}_2^{-1} = a$.

We generate estimates of fixed decision boundaries for each participant, $\hat{a}_j \in (0, 60)$. In Panels A and B of Fig. 3.5, we plot $\hat{a}_j$ as horizontal red lines. The estimated boundaries for each participant approximately contain the set of all cumulative samples at which a decision was made for each participant.

Note that when we switch to (average sample, sample number) space in Panels C and D, the fixed boundary parameter, a, must be divided by sample number. Therefore, we observe exponential collapsing bounds when states are in terms of average sample.

3.4 Discussion

To establish if sequential sampling models apply to risky decisions from experience and to test some of the underlying assumptions of these models, we ran an online study in which participants chose to Play or Skip a slot machine, based on a stream of samples drawn from its outcome distribution.

First, we found that choices were very sensitive to the average of the samples seen and that the number of samples increased the closer machine mean was to 0. Our psychometric results are consistent with sequential sampling models, which suggest that they are applicable to risky choices from experience.

Second, we ran a series of logistic regressions to generate evidence weights for each sample in a trial. We found that the first three samples were weighted more than middle samples, but that the last three samples before the last sample were weighted the most. This is consistent with evidence for both primacy and recency effects in

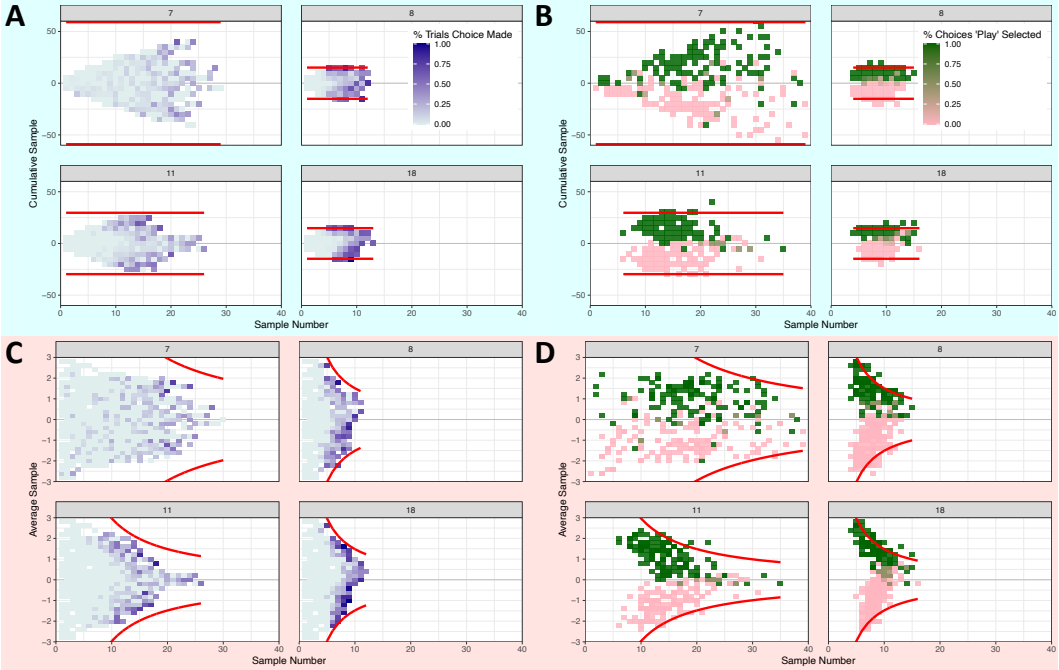

Figure 3.5: Termination Rules and Policies for Participants 7, 8 ,11, and 18. (Panel A) Define a state by the combination of cumulative sample and sample number. For every state, we plot the proportion of trials in which a choice was made in this state, conditional on having reached this state. If less than 5 trials reached a state, that state was excluded. (Panel B) Define a state by the combination of cumulative sample and sample number. For every state, we plot the proportion of trials in which Play was selected, conditional on a choice being made in this state. (Panel C) Define a state by the combination of average sample and sample number. Repeat the same calculation as in Panel A. (Panel D) Define a state by the combination of average sample and sample number. Repeat the same calculation as in Panel C. Red lines plot the estimated decision boundaries from Theorem 1.

sequential sampling (see Ratcliff, P. L. Smith, et al. (2016) for a discussion). Last samples did not exhibit any effect on the choice. Taken altogether, these results provide evidence for leaky integration, collapsing boundaries, and a delay in sample integration.

Third, we ran another series of logistic regressions to generate sample weights for the size of each sample in a trial. We found evidence of value (anti-)compression that occurs at different times during a trial. Regardless, the impact of these transformations on making the correct decision is the same as long as the sign of the sample and the machine mean match.

Lastly, we establish a link between the DDM with fixed decision boundaries and a

Modified Probit model which allows for estimation of the decision boundary in the space of cumulative samples without any computational modeling.

These results are important for many reasons. They establish that risky choices from experience fall under a class of decisions that can be described as a trade-off between accuracy and speed (Forstmann, Ratcliff, and Wagenmakers, 2016). They also identify non-linear sample weighting in the integration process. Although these are not modeled by the standard DDM, various sequential sampling models have incorporated parts. For example: the leaky competing accumulator incorporates leakage (Usher and McClelland, 2001); variations of the DDM with collapsing bounds exist (Hawkins et al., 2015; Glickman and Usher, 2019); and the piecewise linear ballistic accumulator incorporates a delay in sample integration (Holmes, Trueblood, and Heathcote, 2016).

Our evidence for collapsing boundaries is consistent with Bayesian models of optimal value-based decision-making (Tajima, Drugowitsch, and Pouget, 2016), though our experiment design was not set up to induce a sense of urgency in making the choice (Glickman and Usher, 2019). We also found that late samples were weighted more than early and middle samples, which follows the pattern of early under-reaction and late over-reaction documented by Prat-Carrabin and Woodford (2022b) in their probabilistic inference task.

Webb (2018) establishes a link between the choice probabilities implied by the DDM and the logit model. Rather than model *which* choice was made, we look at *whether* a choice was made at time t. This allows us to prove a link between the DDM and a Modified Probit, and use this to estimate fixed decision bounds in cumulative sample space.

References

Basten, Ulrike et al. (2010). "How the brain integrates costs and benefits during decision making". In: *Proceedings of the National Academy of Sciences* 107.50, pp. 21767–21772. DOI: 10.1073/pnas.0908104107.

Busemeyer, Jerome R. (1985). "Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models." In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11.3, pp. 538–564. DOI: 10.1037/0278-7393.11.3.538.

Cheadle, Samuel et al. (2014). "Adaptive Gain Control during Human Perceptual Choice". In: *Neuron* 81.6, pp. 1429–1441. DOI: 10.1016/j.neuron.2014.01.020.

Clarmann von Clarenau, Verena et al. (2024). "Over- and underweighting of extreme values in decisions from sequential samples". In: *Journal of Experimental Psychology: General* 153.3, pp. 814–826. DOI: 10.1037/xge0001530.

Eum, Brenden, Stephanie Dolbier, and Antonio Rangel (2023). "Peripheral Visual Information Halves Attentional Choice Biases". In: *Psychological Science* 34.9, pp. 984–998. DOI: 10.1177/09567976231184878.

Forstmann, Birte U., Roger Ratcliff, and Eric-Jan Wagenmakers (2016). "Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions". In: *Annual Review of Psychology* 67.1, pp. 641–666. DOI: 10.1146/annurev-psych-122414-033645.

Glickman, Moshe and Marius Usher (2019). "Integration to boundary in decisions between numerical sequences". In: *Cognition* 193, p. 104022. DOI: 10.1016/j.cognition.2019.104022.

Gold, Joshua I. and Michael N. Shadlen (2007). "The Neural Basis of Decision Making". In: *Annual Review of Neuroscience* 30.1, pp. 535–574. DOI: 10.1146/annurev.neuro.29.051605.113038.

Hare, Todd A. et al. (2011). "Transformation of stimulus value signals into motor commands during simple choice". In: *Proceedings of the National Academy of Sciences* 108.44, pp. 18120–18125. DOI: 10.1073/pnas.1109322108.

Hau, Robin, Timothy J. Pleskac, and Ralph Hertwig (2010). "Decisions from experience and *statistical probabilities* : Why they trigger different choices than a priori probabilities". In: *Journal of Behavioral Decision Making* 23.1, pp. 48–68. DOI: 10.1002/bdm.665.

Hawkins, Guy E. et al. (2015). "Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making". In: *The Journal of Neuroscience* 35.6, pp. 2476–2484. DOI: 10.1523/JNEUROSCI.2410-14.2015.

Hertwig, Ralph, Greg Barron, et al. (2004). "Decisions from Experience and the Effect of Rare Events in Risky Choice". In: *Psychological Science* 15.8, pp. 534–539. DOI: 10.1111/j.0956-7976.2004.00715.x.

Hertwig, Ralph and Ido Erev (2009). "The description–experience gap in risky choice". In: *Trends in Cognitive Sciences* 13.12, pp. 517–523. DOI: 10.1016/j.tics.2009.09.004.

Holmes, William R., Jennifer S. Trueblood, and Andrew Heathcote (2016). "A new framework for modeling decisions about changing information: The Piecewise Linear Ballistic Accumulator model". In: *Cognitive Psychology* 85, pp. 1–29. DOI: 10.1016/j.cogpsych.2015.11.002.

Irwin, Francis W. and W. A. S. Smith (1957). "Value, cost, and information as determiners of decision." In: *Journal of Experimental Psychology* 54.3, pp. 229–232. DOI: 10.1037/h0049137.

Irwin, Francis W., W. A. S. Smith, and Jane F. Mayfield (1956). “Tests of two theories of decision in an "expanded judgment" situation.” In: *Journal of Experimental Psychology* 51.4, pp. 261–268. DOI: 10.1037/h0041911.

Kahneman, Daniel and Amos Tversky (1979). “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47.2, p. 263. DOI: 10.2307/1914185.

Katzin, Naama, David Rosenbaum, and Marius Usher (2021). “The averaging of numerosities: A psychometric investigation of the mental line”. In: *Attention, Perception, & Psychophysics* 83.3, pp. 1152–1168. DOI: 10.3758/s13414-020-02140-w.

Krajbich, Ian, Carrie Armel, and Antonio Rangel (2010). “Visual fixations and the computation and comparison of value in simple choice”. In: *Nature Neuroscience* 13.10, pp. 1292–1298. DOI: 10.1038/nn.2635.

Krajbich, Ian, Dingchao Lu, et al. (2012). “The Attentional Drift-Diffusion Model Extends to Simple Purchasing Decisions”. In: *Frontiers in Psychology* 3. DOI: 10.3389/fpsyg.2012.00193.

Krajbich, Ian and Antonio Rangel (2011). “Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions”. In: *Proceedings of the National Academy of Sciences* 108.33, pp. 13852–13857. DOI: 10.1073/pnas.1101328108.

Luce, R.D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley.

Malhotra, Gaurav et al. (2017). “Overcoming indecision by changing the decision boundary.” In: *Journal of Experimental Psychology: General* 146.6, pp. 776–805. DOI: 10.1037/xge0000286.

McFadden, Daniel (1974). “Conditional Logit Analysis of Qualitative Choice Behavior”. In: *Frontiers in Econometrics*. Academic Press New York, pp. 105–142.

Milosavljevic, Milica et al. (2010). “The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure”. In: *Judgment and Decision Making* 5.6.

Prat-Carrabin, Arthur and Michael Woodford (2022a). “Efficient coding of numbers explains decision bias and noise”. In: *Nature Human Behaviour* 6.8, pp. 1142–1152. DOI: 10.1038/s41562-022-01352-4.

– (2022b). *Imprecise Probabilistic Inference from Sequential Data*. preprint. PsyArXiv. DOI: 10.31234/osf.io/xn5mk.

R Core Team (2023). *R: A Language and Environment for Statistical Computing*. Version 4.3.1.

Ratcliff, Roger and Gail McKoon (2008). “The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks”. In: *Neural Computation* 20.4, pp. 873–922. DOI: 10.1162/neco.2008.12-06-420.

Ratcliff, Roger, Philip L. Smith, et al. (2016). "Diffusion Decision Model: Current Issues and History". In: *Trends in Cognitive Sciences* 20.4, pp. 260–281. DOI: 10.1016/j.tics.2016.01.007.

Smith, Philip L. and Roger Ratcliff (2022). "Modeling Evidence Accumulation Decision Processes Using Integral Equations: Urgency-gating and Collapsing Boundaries". In: *Psychological review* 129.2, pp. 235–267. DOI: 10.1037/rev0000301.

Spitzer, Bernhard, Leonhard Waschke, and Christopher Summerfield (2017). "Selective overweighting of larger magnitudes during noisy numerical comparison". In: *Nature Human Behaviour* 1.8, pp. 1–8. DOI: 10.1038/s41562-017-0145.

Stan Development Team (2024). *RStan: the R interface to Stan*. Version 2.32.5.

Tajima, Satohiro, Jan Drugowitsch, and Alexandre Pouget (2016). "Optimal policy for value-based decision-making". In: *Nature Communications* 7.1, p. 12400. DOI: 10.1038/ncomms12400.

Tickle, Hannah et al. (2023). "Human optional stopping in a heteroscedastic world". In: *Psychological Review* 130.1, pp. 1–22. DOI: 10.1037/rev0000315.

Trueblood, Jennifer S. et al. (2021). "Urgency, leakage, and the relative nature of information processing in decision-making." In: *Psychological Review* 128.1, pp. 160–186. DOI: 10.1037/rev0000255.

Tsetsos, Konstantinos, Nick Chater, and Marius Usher (2012). "Salience driven value integration explains decision biases and preference reversal". In: *Proceedings of the National Academy of Sciences* 109.24, pp. 9659–9664. DOI: 10.1073/pnas.1119569109.

Tsetsos, Konstantinos, Rani Moran, et al. (2016). "Economic irrationality is optimal during noisy decision making". In: *Proceedings of the National Academy of Sciences* 113.11, pp. 3102–3107. DOI: 10.1073/pnas.1519157113.

Tsetsos, Konstantinos, Marius Usher, and James L. McClelland (2011). "Testing Multi-Alternative Decision Models with Non-Stationary Evidence". In: *Frontiers in Neuroscience* 5. DOI: 10.3389/fnins.2011.00063.

Turkakin, Esin et al. (2023). *Should I Sample or Should I Go? An approximately optimal model for deciding when to stop sampling information*. DOI: 10.31234/osf.io/tfe94.

Tversky, Amos and Daniel Kahneman (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty". In: *Journal of Risk and Uncertainty* 5.4, pp. 297–323.

Usher, Marius and James L. McClelland (2001). "The time course of perceptual choice: The leaky, competing accumulator model." In: *Psychological Review* 108.3, pp. 550–592. DOI: 10.1037/0033-295X.108.3.550.

Usher, Marius, Konstantinos Tsetsos, et al. (2019). "Selective Integration: An Attentional Theory of Choice Biases and Adaptive Choice". In: *Current Directions in Psychological Science* 28.6, pp. 552–559. DOI: 10.1177/0963721419862277.

Webb, Ryan (2018). "The (Neural) Dynamics of Stochastic Choice". In: *Management Science* 65.1, pp. 230–255. DOI: 10.1287/mnsc.2017.2931.

Wulff, Dirk U., Max Mergenthaler-Canseco, and Ralph Hertwig (2018). "A meta-analytic review of two modes of learning and the description-experience gap." In: *Psychological Bulletin* 144.2, pp. 140–176. DOI: 10.1037/bul0000115.

Zeigenfuse, Matthew D., Timothy J. Pleskac, and Taosheng Liu (2014). "Rapid decisions from experience". In: *Cognition* 131.2, pp. 181–194. DOI: 10.1016/j.cognition.2013.12.012.

Appendix A

SUPPLEMENTARY MATERIALS FOR CHAPTER I

A.1 Figures

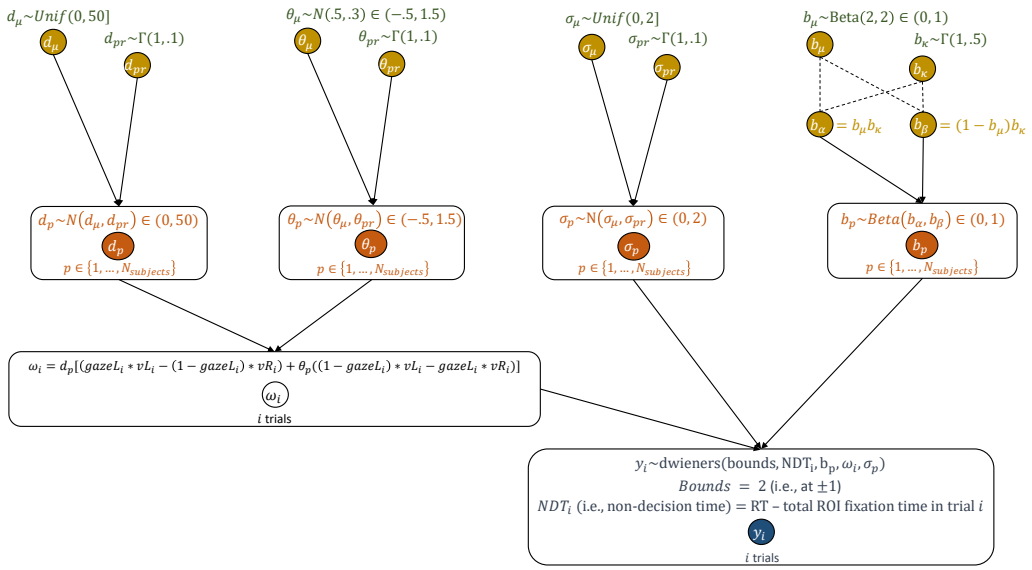

Figure A.1: Directed Acyclic Graph of Hierarchical aDDM with Uncorrelated Priors. The hierarchical model estimates group and individual parameters for the aDDM with uncorrelated priors. The 10 group parameters are depicted in the top row of yellow circles. The four individual parameters estimated for each subject are depicted in orange in the middle row. The distribution of individual parameters as a function of the group parameters is specified using a transformation of some of the parameters, denoted by the dashed lines. The choice and RT outcome y_i of trial i for a subject p is modeled as a Drift-Diffusion-Model with bounds at ± 1, non-decision time NDT_i, bias b_p, trial specific slope ω_i, and noise σ_p. The trial specific slope ω_i depends on the subject's drift rate parameter d_p, attentional bias parameter θ_p, gaze data for the trial $gazeL_i$, and item liking ratings for the foods used in the trial (vL_i,vR_i). $gazeL_i$ denotes the proportion of time spent fixating on the left item during the trial. The hyperpriors for the group parameters are described at the top of the graph. "$\in (X,Y)$" indicates truncation to bounds. $N_{subjects} = 25$ in the exploratory and confirmatory datasets, and 50 in the joint dataset.

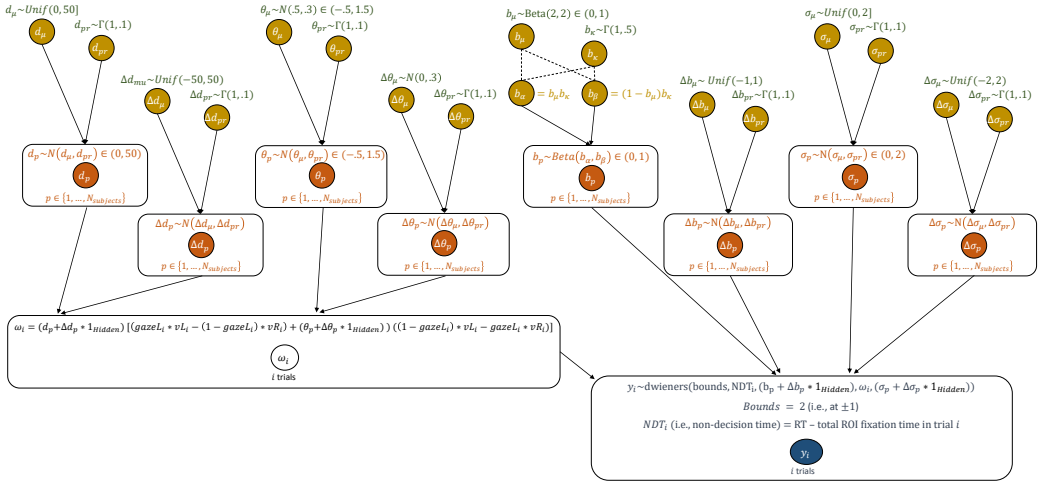

Figure A.2: Directed Acyclic Graph of Hierarchical aDDM with Correlated Priors. The hierarchical model estimates group and individual parameters for the aDDM with correlated priors. The 18 group parameters are depicted in the top row of yellow circles. The 8 individual parameters estimated for each subject are depicted in orange in the middle row. The distribution of individual parameters as a function of the group parameters is specified using a transformation of some of the parameters, denoted by the dashed lines. The choice and RT outcome y_i of trial i for a subject p is modeled as a Drift-Diffusion-Model with bounds at ±1, non-decision time NDT_i, bias b_p, conditional difference in bias Δb_p, trial specific slope ω_i, noise σ_p, and conditional difference in noise $\Delta\sigma_p$. The trial specific slope ω_i depends on the subject's drift rate parameter d_p, conditional difference in drift rate parameter Δd_p, attentional bias parameter θ_p, conditional difference in attentional bias parameter $\Delta\theta_p$, gaze data for the trial $gazeL_i$, and item liking ratings for the foods used in the trial (vL_i,vR_i). $gazeL_i$ denotes the proportion of time spent fixating on the left item during the trial. The hyperpriors for the group parameters are described at the top of the graph. "$\in (X, Y)$" indicates truncation to bounds. $N_{subjects} = 25$ in the exploratory and confirmatory datasets, and 50 in the joint dataset.

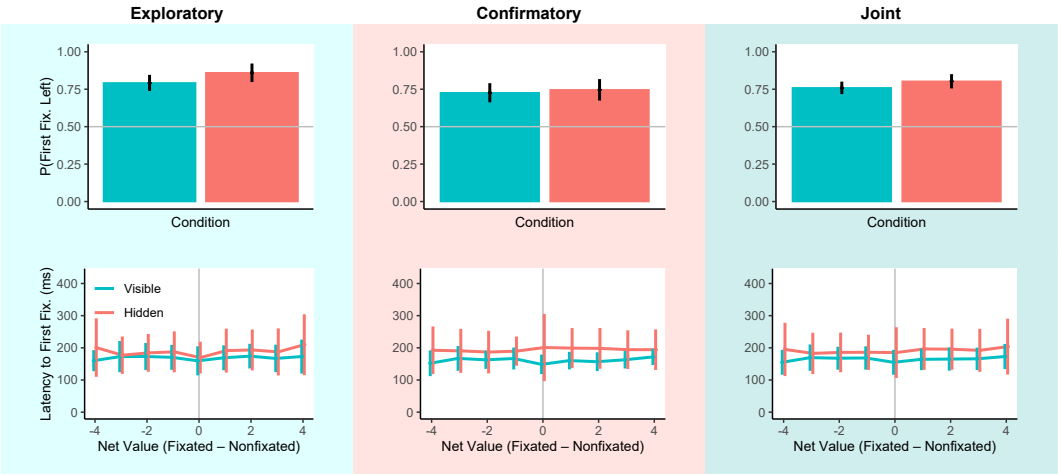

Figure A.3: Additional Fixation Properties. (Top) Probability of first fixation to the left item in the two conditions. (Bottom) Latency to first fixation as a function of the relative rating of the fixated item. Columns indicate which dataset generated the figures. Black error bars show standard errors of the mean across participants. Color error bars show standard deviation across participants.

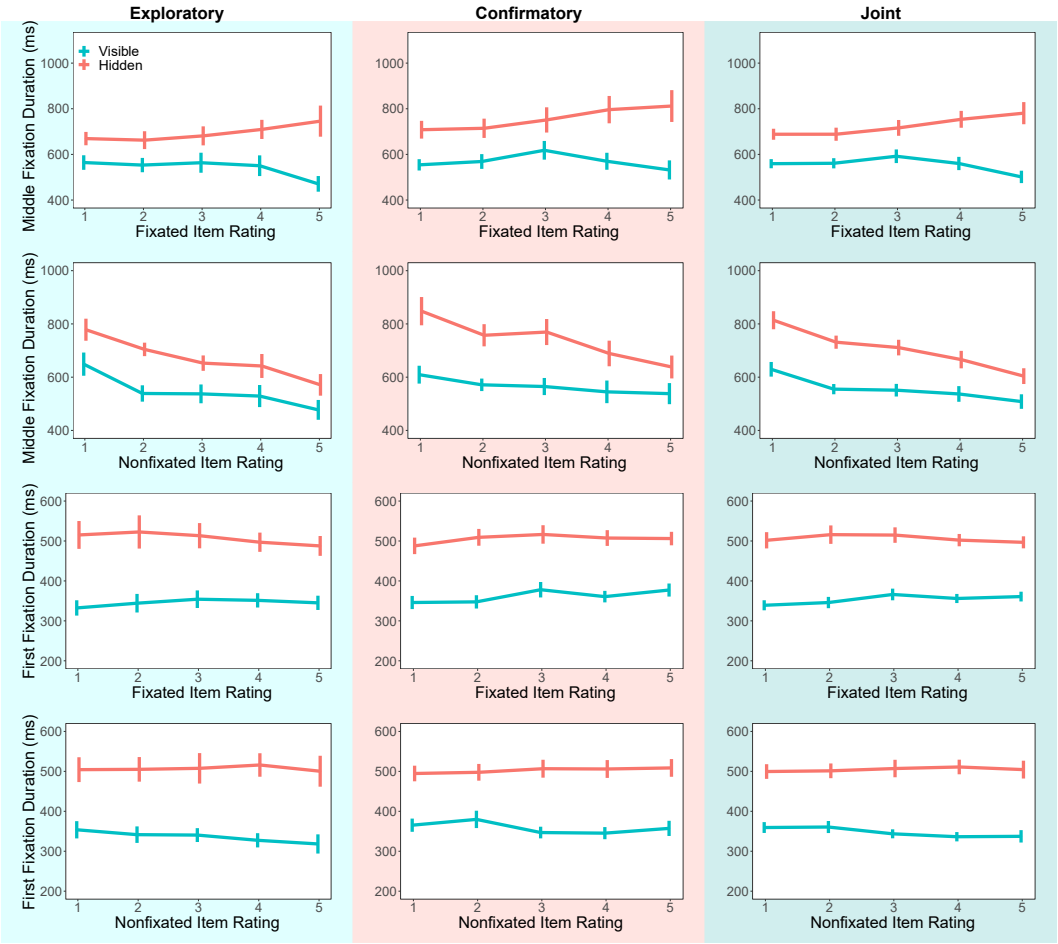

Figure A.4: Fixation Durations. (Row 1) Middle fixation duration as a function of the fixated item rating. (Row 2) Middle fixation duration as a function of the nonfixated item rating. (Row 3) First fixation duration as a function of the fixated item rating. (Row 4) First fixation duration as a function the nonfixated item rating. Columns indicate which dataset generated the figures. Error bars show standard errors of the mean across participants.

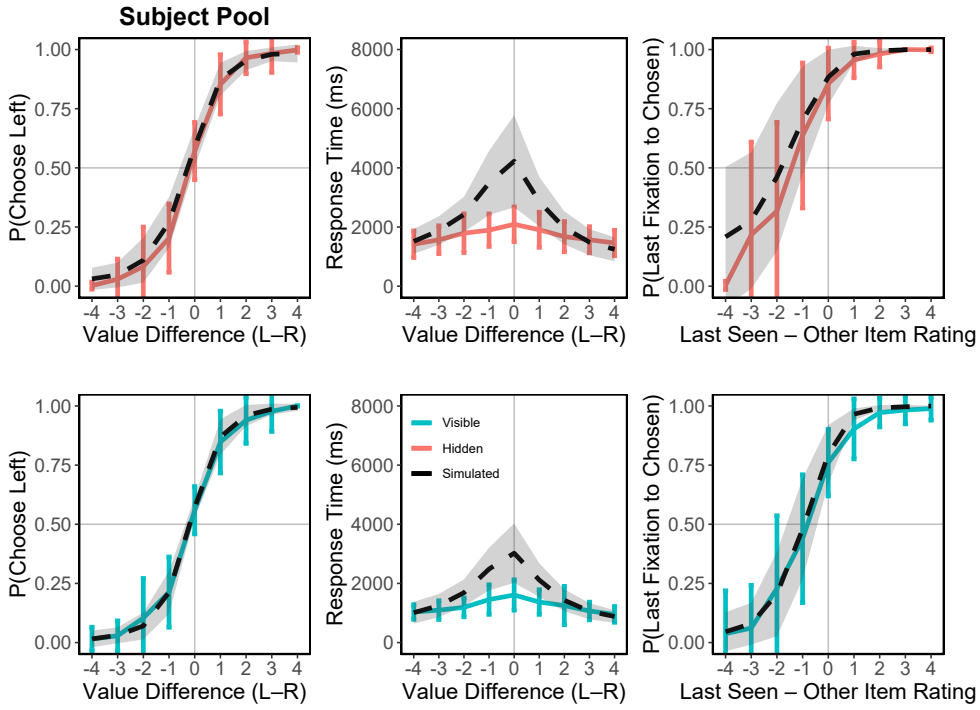

Figure A.5: Group-level Predictions in the Joint Dataset. We use the estimates of the hierarchical aDDM in the odd trials to make predictions out-of-sample, in the even trials, separately for each subject and condition. For each subject, we simulate 10 observations per trial, and compare the simulated and observed data. *Blue lines:* Behavior in the visible condition. *Red lines:* Behavior in the hidden condition. *Black dashed lines and grey areas:* Simulated behavior for the respective condition (dash = mean, grey = SD). Error bars show standard deviations across subjects.

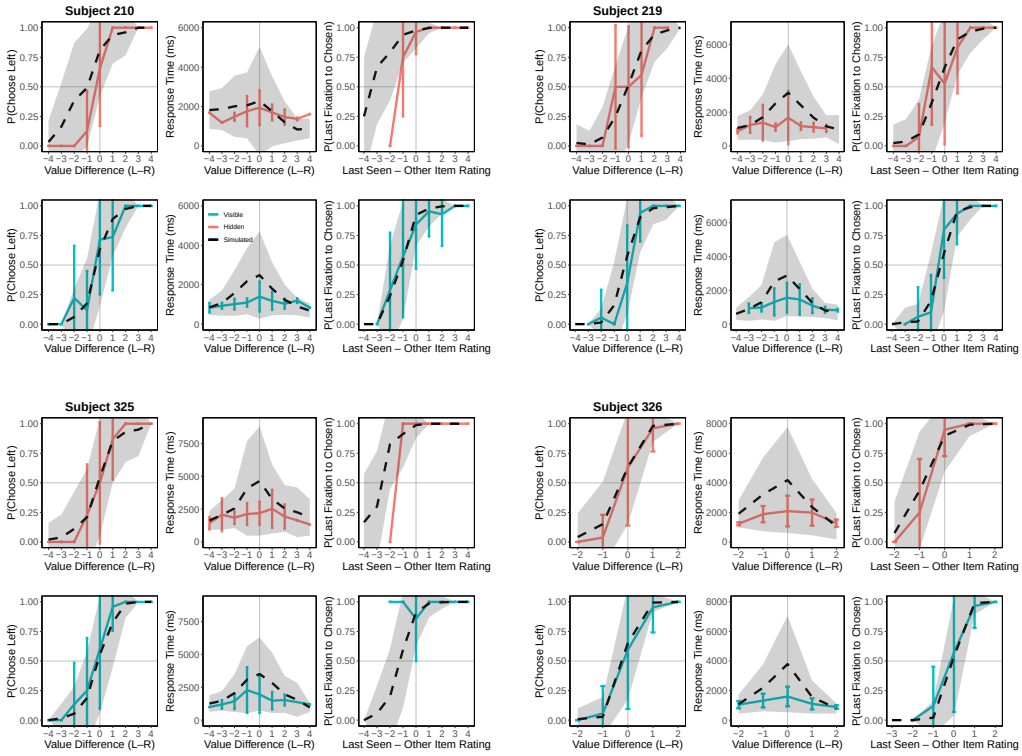

Figure A.6: Subject-level Predictions. Out-of-sample predictions versus data for four randomly selected subjects. See Figure A.5 and Supplementary Methods for details. *Blue lines:* Behavior in the visible condition. *Red lines:* Behavior in the hidden condition. *Black dashed lines and grey areas:* Simulated behavior for the respective condition (dash = mean, grey = SD). Error bars show standard deviations across subjects.

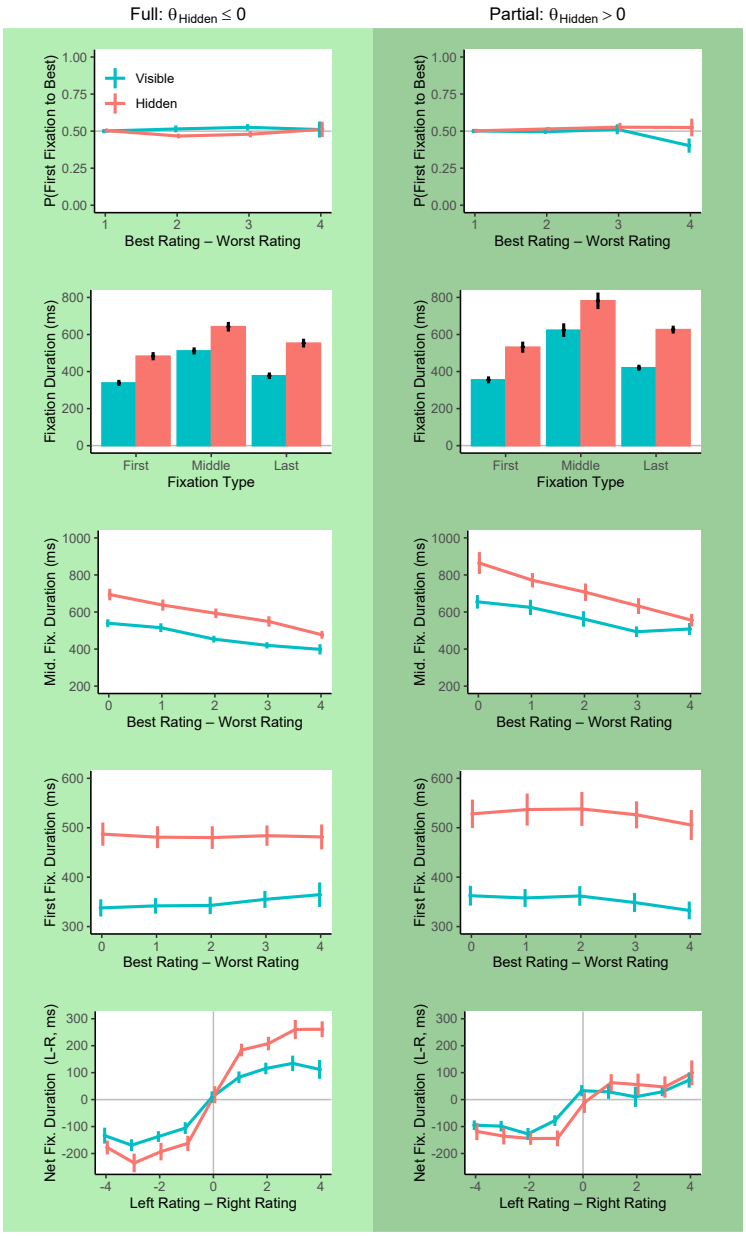

Figure A.7: Fixation Process by Attentional Discounting Group. (Row 1) The probability that the first fixation is to the best item as a function of choice difficulty. (Row 2) Fixation durations by fixation type. (Row 3) Middle fixation duration as a function of choice difficulty. (Row 4) First fixation duration as a function of choice difficulty. (Row 5) Net fixation duration to the left item as a function of its relative value. Columns indicate which dataset generated the figures: "$\theta_{Hidden} \leq 0$" indicates subjects with full attentional discounting ($N = 7$), "$\theta_{Hidden} > 0$" indicates subjects with partial attentional discounting ($N = 43$). Data is from the joint dataset. Error bars show standard errors of the mean across participants.

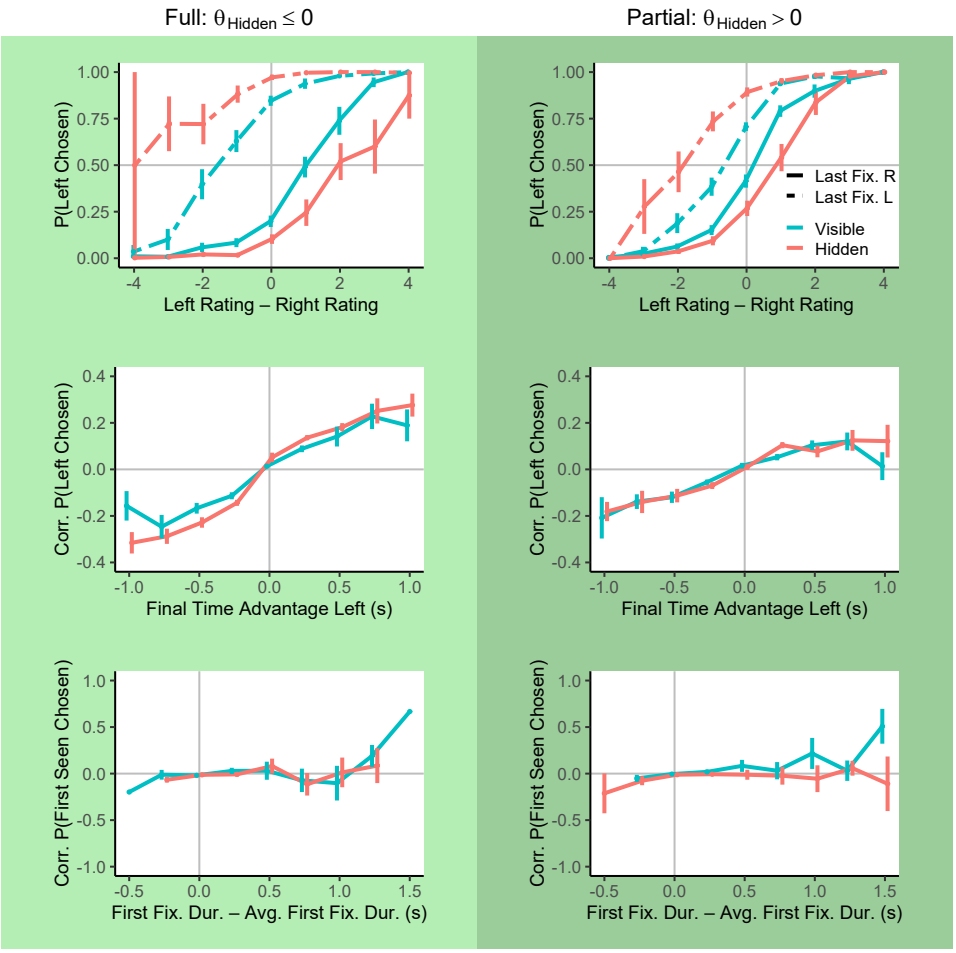

Figure A.8: Choice Biases by Attentional Discounting Group. (Top) The probability of choosing the left item as a function of its relative value, conditional on last fixation location. (Middle) The corrected probability of choosing the left item as a function of the net fixation time to the left item. (Bottom) Corrected probability that the first seen item is chosen as a function of the excess first fixation duration, defined as first fixation duration minus mean first fixation duration (computed for each subject). Columns indicate which dataset generated the figures: "$\theta_{Hidden} \leq 0$" indicates subjects with full attentional discounting ($N = 7$), "$\theta_{Hidden} > 0$" indicates subjects with partial attentional discounting ($N = 43$). Data is from the joint dataset. Error bars show standard errors of the mean across participants.

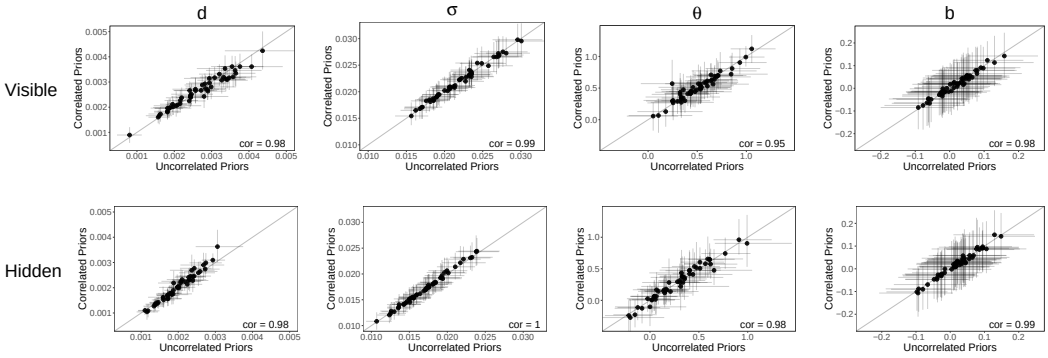

Figure A.9: Comparison of Subject-level aDDM MAP Parameter Estimates Between Models with Uncorrelated and Correlated Priors. Parameter estimate from the model with correlated priors as a function of its respective estimate from the model with uncorrelated priors. Rows separate by condition, columns separate by parameter. Data from joint data set. Black lines denote 95% HDIs. "cor" denotes Pearson correlation coefficient.

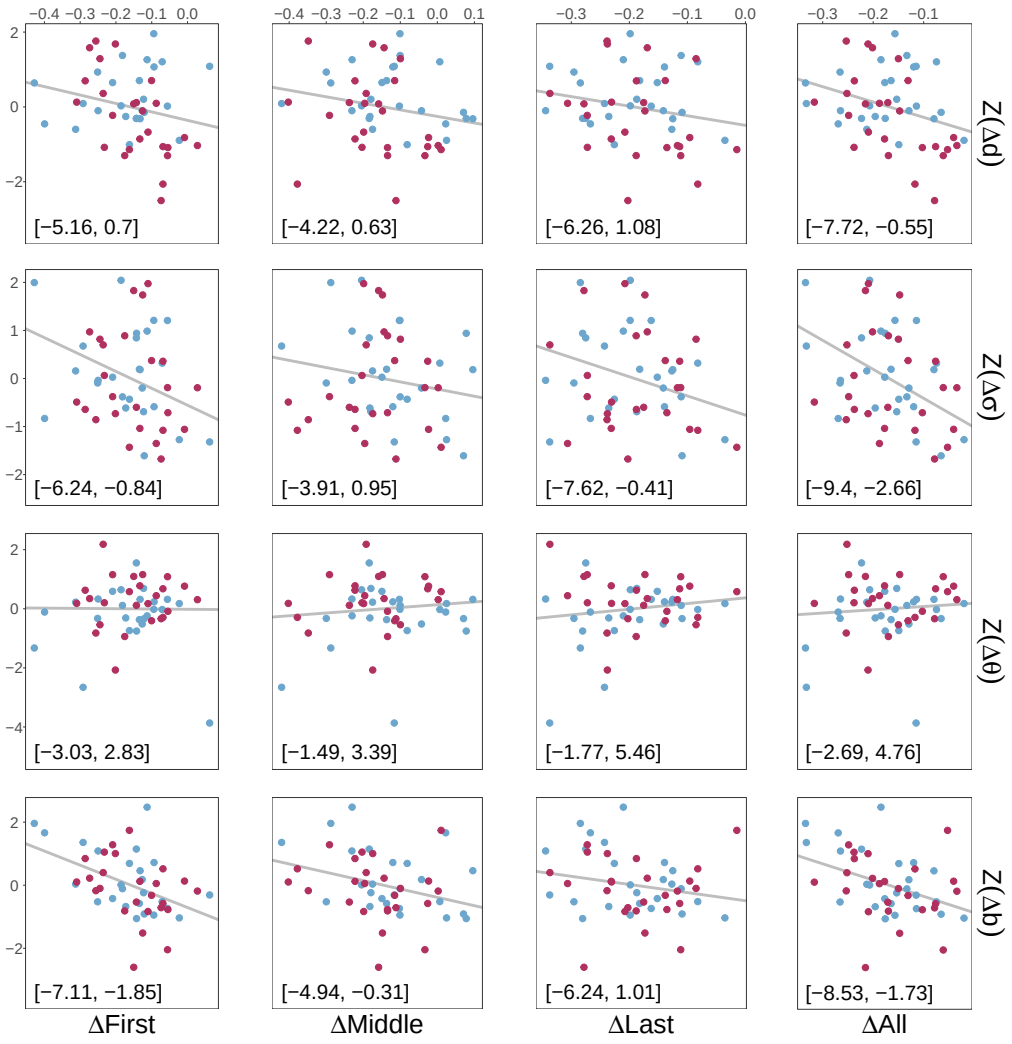

Figure A.10: Linking Fixation Patterns with aDDM Estimates. Each point represents a subject (blue = exploratory, red = confirmatory). ΔFirst is the mean first fixation duration in the visible condition minus the mean in the hidden condition. ΔMiddle, ΔLast, and ΔAll are the same, except for middle, last, and all fixations, respectively. Δd is the MAP estimate of drift rate in the visible condition minus the MAP estimate in the hidden condition. $\Delta\sigma$, $\Delta\theta$, and Δb are the same, except for the noise, attentional discounting, and bias parameters, respectively. Δparameters have been Z-scored. Data is from the joint dataset. Grey lines are univariate linear regression predictions. Black text presents the 95% HDI for the slope of the regression line.

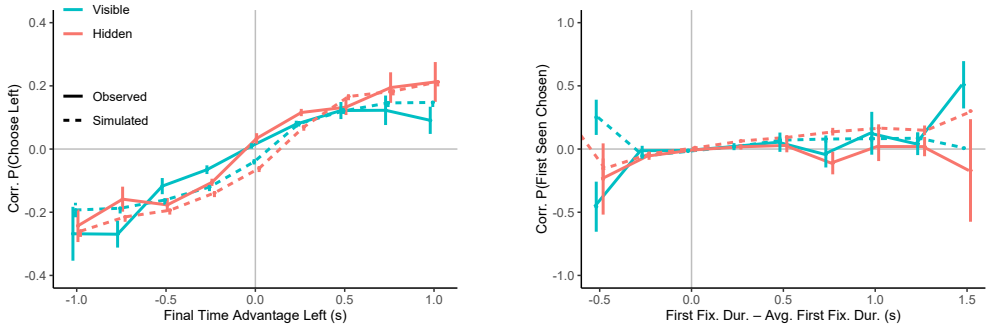

Figure A.11: Other Choice Bias Simulations. (Left) Corrected probability of choosing the left item as a function of net fixation time to the left item, in observed (solid) and simulated (dashed) out-of-sample data. The corrected probability is computed by subtracting from each choice observation (coded as 1 if left chosen, and 0 otherwise) the proportion with which left is chosen at each relative value. (Right) Corrected probability that the first seen item is chosen as a function of the excess first fixation duration, defined as first fixation duration minus mean first fixation duration (computed for each subject), in observed and simulated out-of-sample data. Blue lines indicate data in the visible condition. Red lines indicate data in the hidden condition. Data is from the joint dataset. Error bars show standard errors of the mean across participants.

A.2 Tables

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Left chosen	Intercept	0.11	0.07		0.03	0.04		0.07	0.04	
(Logistic)	L - R rating	1.58	0.13	*	1.69	0.15	*	1.64	0.10	*
(Top)	Hidden	-0.03	0.09		0.09	0.06		0.03	0.06	
	Interaction	0.19	0.10		0.08	0.08		0.12	0.06	*
RT	Intercept	1491.27	99.61	*	1759.54	121.52	*	1635.85	80.61	*
(Linear)	B - W rat.	-172.28	25.95	*	-213.22	26.02	*	-194.33	18.05	*
(Middle)	Hidden	542.61	76.61	*	497.94	96.90	*	520.08	61.05	*
	Interaction	-13.62	18.93		-33.38	30.68		-22.47	17.17	
# of fix.	Intercept	2.83	0.09	*	3.19	0.10	*	3.01	0.07	*
(Linear)	B - W rat.	-0.20	0.02	*	-0.27	0.02	*	-0.23	0.02	*
(Bottom)	Hidden	-0.02	0.07		-0.21	0.10	*	-0.11	0.06	
	Interaction	0.05	0.02	*	0.09	0.03	*	0.07	0.02	*

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
"B - W rat." Best - Worst Rating.

Table A.1: Regressions associated with the basic psychometric results in Fig. 1.3.

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
1st fix. best	Intercept	0.14	0.06	*	-0.10	0.05		0.01	0.04	
(Logistic)	Best - worst rat.	-0.04	0.03		0.04	0.03		0.00	0.02	
(Row 1)	Hidden	-0.16	0.09		0.13	0.08		0.00	0.06	
	Interaction	0.05	0.05		-0.06	0.04		-0.02	0.03	
Mid. fix. dur.	Intercept	618.35	36.49	*	595.78	36.27	*	612.92	23.59	*
(Linear)	Best - worst rat.	-55.23	8.09	*	-47.48	8.37	*	-51.20	5.84	*
(Row 3)	Hidden	103.82	27.53	*	129.57	33.42	*	141.50	22.48	*
	Interaction	-7.04	11.68		-7.87	15.14		-13.19	9.65	
1st fix. dur.	Intercept	341.64	21.31	*	353.30	16.29	*	348.37	13.23	*
(Linear)	Best - worst rat.	-0.83	2.43		1.50	2.38		0.32	1.66	
(Row 4)	Hidden	159.51	25.57	*	152.78	20.04	*	158.73	15.57	*
	Interaction	0.50	3.84		-4.78	3.39		-2.18	2.44	
Net fix. dur.	Intercept	14.45	20.80		-12.18	30.78		4.44	17.81	
(Linear)	Net. Val. > 0 (A)	-56.13	59.87		-38.46	69.75		-54.50	44.02	
(Row 5)	Net. Val. < 0 (B)	50.64	59.66		14.11	67.26		24.83	42.89	
	A : Net. Val. (C)	9.78	20.15		16.86	21.88		14.55	14.77	
	B : Net. Val. (D)	63.20	20.40	*	47.00	21.64	*	54.48	14.66	*
	Hidden (E)	-22.05	39.01		-21.88	42.16		-23.83	27.89	
	A:E	68.83	41.34		81.30	43.28		77.78	29.09	*
	B:E	-108.67	39.16	*	-40.40	43.83		-70.26	28.91	*
	C:E	3.08	12.73		8.90	14.65		5.49	9.84	
	D:E	-35.43	13.05	*	-6.01	13.85		-20.20	9.35	*

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.

Table A.2: Regressions associated with the fixation results in Fig. 1.4.

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Lat. to 1st fix.	Intercept	167.16	8.89	*	157.22	6.89	*	162.48	162.48	*
(Linear)	Fix. - nonfix. rat.	-0.97	3.26		1.45	4.38		0.30	0.30	
	Hidden	3.74	4.42		5.06	4.75		8.20	8.20	
	Interaction	1.12	3.97		0.81	4.59		2.17	2.17	

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
"Lat." Latency. "rat." rating.

Table A.3: Regressions associated with additional fixation results in Fig. A.3.

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Mid. fix. dur.	Intercept	587.88	36.06	*	567.63	27.92	*	576.42	21.60	*
(Linear)	Fix. rating	-11.83	7.66		1.43	8.44		-3.96	5.66	
(Row 1)	Hidden	61.91	28.76	*	85.31	29.15	*	78.08	20.22	*
	Interaction	24.21	10.49	*	29.37	11.11	*	26.19	7.51	*
Mid. fix. dur.	Intercept	672.53	38.31	*	654.00	28.25	*	661.13	23.87	*
(Linear)	Nonfix. rating	-42.07	6.73	*	-29.98	8.84	*	-34.24	5.74	*
(Row 2)	Hidden	143.89	30.69	*	218.65	38.68	*	193.83	25.95	*
	Interaction	-8.57	9.43		-21.26	9.77	*	-16.84	6.62	*
1st fix. dur.	Intercept	329.36	21.44	*	340.71	20.31	*	334.77	14.75	*
(Linear)	Fix. rating	4.17	2.23		5.71	3.47		5.19	2.00	*
(Row 3)	Hidden	189.71	30.63	*	156.43	23.03	*	176.84	19.61	*
	Interaction	-10.27	4.56	*	-3.45	3.56		-7.31	3.07	*
1st fix. dur.	Intercept	357.71	21.32	*	374.03	17.20	*	367.47	13.59	*
(Linear)	Nonfix. rating	-6.73	2.15	*	-6.63	2.01	*	-6.52	1.46	*
(Row 4)	Hidden	129.55	24.69	*	122.36	19.34	*	130.51	15.53	*
	Interaction	10.20	3.18	*	7.81	2.94	*	8.80	2.12	*

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.

Table A.4: Regressions associated with the fixation duration results in Fig. A.4.

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Left chosen	Intercept	-0.68	0.16	*	-1.59	0.26	*	-1.11	0.16	*
(Logistic)	Left - right rating (A)	1.54	0.14	*	1.75	0.18	*	1.64	0.11	*
(Top)	Last fix. loc. (B)	1.83	0.31	*	3.20	0.47	*	2.46	0.30	*
	Hidden (C)	-1.20	0.30	*	-1.46	0.36	*	-1.32	0.23	*
	A:B	0.15	0.12		-0.14	0.13		-0.01	0.08	
	A:C	0.28	0.15	*	-0.18	0.18		0.08	0.09	
	B:C	3.74	0.63	*	3.82	0.63	*	3.69	0.43	*
	A:B:C	-0.59	0.23	*	0.14	0.25		-0.26	0.15	
Corr. left chosen	Intercept	0.00	0.01		0.00	0.00		0.00	0.00	
(Linear)	Net fixation left (s)	0.25	0.04	*	0.24	0.04	*	0.24	0.03	*
(Middle)	Hidden	0.00	0.01		0.01	0.01		0.01	0.01	
	Interaction	0.04	0.04		0.04	0.03		0.04	0.02	
Corr. 1st seen	Intercept	-0.01	0.01		-0.02	0.01	*	-0.01	0.00	*
chosen	Excess first fix. dur. (s)	0.21	0.07	*	0.21	0.08	*	0.22	0.05	*
(Linear)	Hidden	-0.01	0.01		0.00	0.01		0.00	0.01	
(Bottom)	Interaction	-0.21	0.07	*	-0.13	0.06	*	-0.18	0.05	*

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
Last fixation location coded as (1=L, 0=R).

Table A.5: Regressions associated with choice bias results in Fig. 1.5.

	Exploratory		**Confirmatory**		**Joint**	
	H	V	H	V	H	V
d	0.002	0.003	0.002	0.003	0.002	0.003
	[0.002, 0.002]	[0.002, 0.003]	[0.002, 0.002]	[0.002, 0.003]	[0.002, 0.002]	[0.002, 0.003]
σ	0.018	0.023	0.016	0.021	0.017	0.022
	[0.015, 0.020]	[0.021, 0.025]	[0.014, 0.019]	[0.019, 0.023]	[0.016, 0.019]	[0.021, 0.023]
θ	0.38	0.54	0.20	0.49	0.28	0.52
	[0.22, 0.54]	[0.42, 0.66]	[0.06, 0.35]	[0.37, 0.62]	[0.18, 0.39]	[0.44, 0.61]
b	0.02	0.03	0.01	0.00	0.02	0.02
	[-0.10, 0.14]	[-0.06, 0.11]	[-0.10, 0.14]	[-0.09, 0.09]	[-0.04, 0.08]	[-0.03, 0.06]

MAP estimate and 95% HDI of group-level mean.

Table A.6: Group-level MAP Parameter Estimates for Model with Correlated Priors across Datasets and Conditions.

Dept. Var.	Indept. Var.	Est.	SE	
Corr. left chosen	Intercept	-0.02	0.00	*
(Linear)	Net fixation left (s)	0.13	0.02	*
(Left)	Hidden	0.00	0.00	
	Interaction	0.08	0.02	*
Corr. 1st seen chosen	Intercept	-0.01	0.00	*
(Linear)	Excess first fix. dur. (s)	0.06	0.02	*
(Right)	Hidden	0.03	0.00	*
	Interaction	0.06	0.02	*

* indicates significance in all data sets at the 95% confidence level.

Table A.7: Regressions associated with choice bias simulations in Fig. A.11.

Appendix B

SUPPLEMENTARY MATERIALS FOR CHAPTER II

B.1 Figures

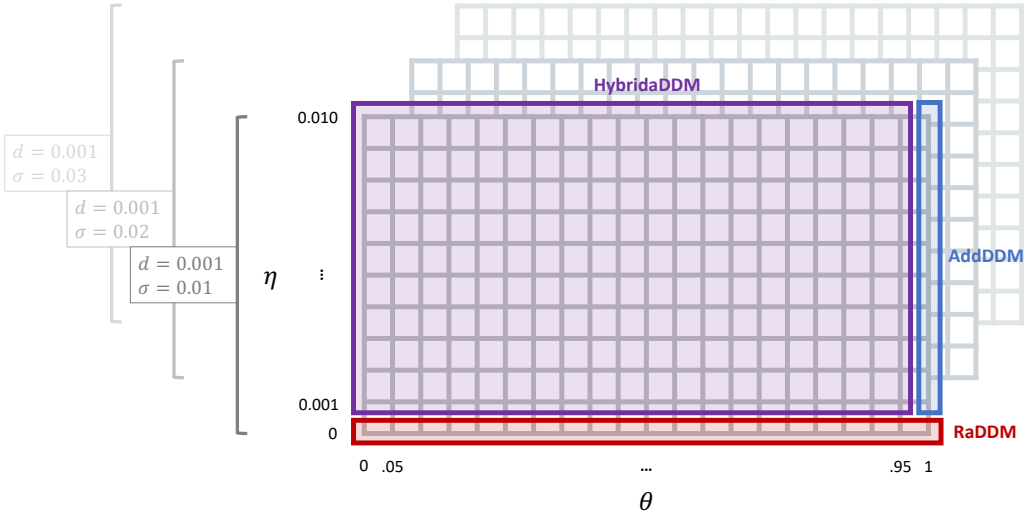

Figure B.1: Parameter Grid for Model Fitting. Grey grids represent every possible combination of (η, θ), given a combination of (d, σ). The Hybrid aDDM is defined on every possible combination of $\eta \in \{0.001, 0.002, \ldots, 0.010\}$ and $\theta \in \{0.00, 0.05, \ldots, 0.95\}$, implying some level of both additive and multiplicative attentional bias. The AddDDM fixes $\theta = 1$ and lets η vary along $\{0.001, 0.002, \ldots, 0.010\}$. The RaDDM fixes $\eta = 0$ and lets θ vary along $\{0.00, 0.05, \ldots, 1.00\}$. Note that the RaDDM nests the standard DDM when $\theta = 1$. $d \in \{0.001, 0.002, \ldots, 0.015\}$ and $\sigma \in \{0.01, 0.02, \ldots, 0.09\}$ for all models.

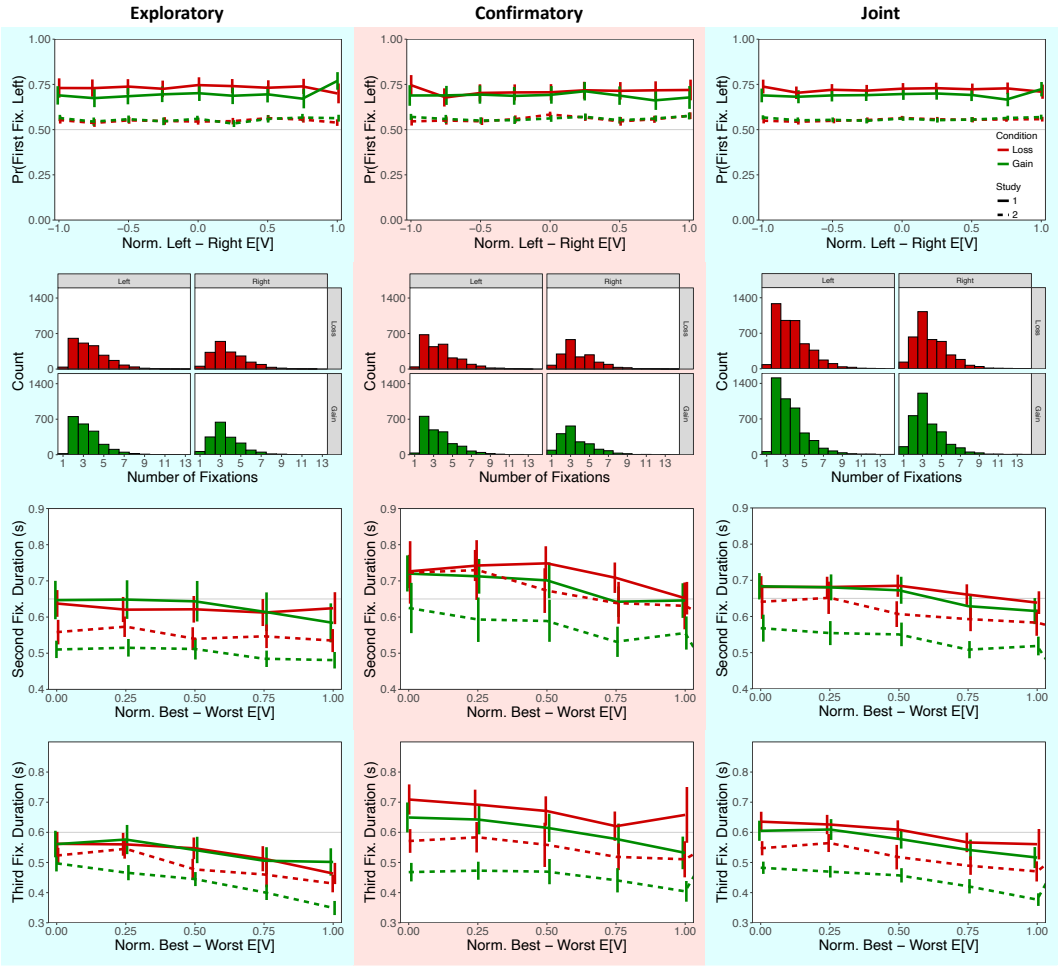

Figure B.2: Additional Fixation Properties. The first row depicts the probability of first fixating left as a function of choice difficulty. The second row pools over all subjects in Study 2 and depicts histograms of the number of fixations in a trial, separately by condition and the location of first fixation. The third row depicts the second fixation duration as a function of the normalized relative value. The fourth row depicts the third fixation duration as a function of the normalized relative value. Error bars denote standard error of the mean across participants.

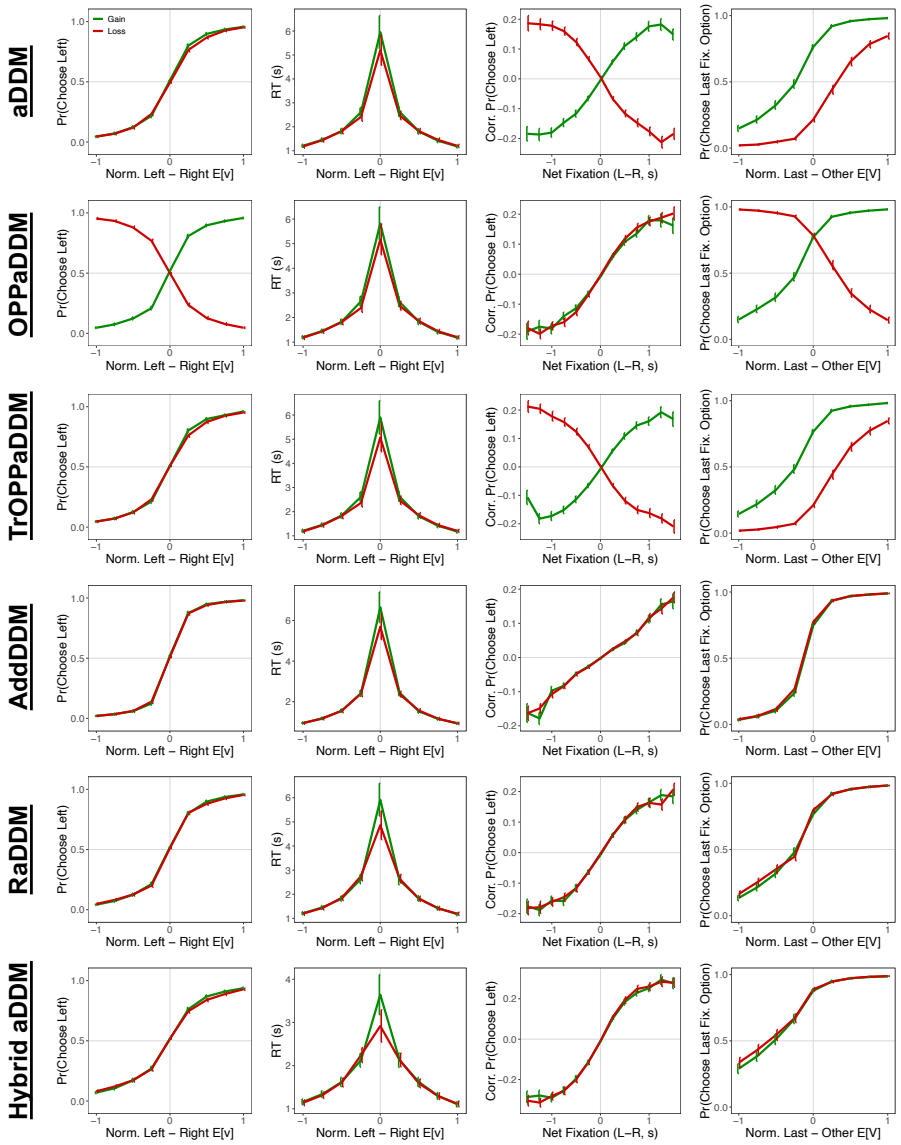

Figure B.3: Simulations of Attentional Choice Biases. For each model and condition, we generated a simulated dataset of 200 trials for each of 25 simulated subjects, using the best-fitting estimates from the exploratory dataset. (Column 1) the psychometric choice curve, (Column 2) the RT curve, (Column 3) the net fixation bias curve, and (Column 4) the last fixation bias curve. For details, see Figs. 2.3 and 2.5. (Row 1) Data simulated using the aDDM with $\theta \in (0, 1)$. (Row 2) Data simulated using aDDM with final choice = opposite of barrier hit. (Row 3) Data simulated using aDDM with absolute value signals and final choice = opposite of barrier hit. (Row 4) Data simulated using AddDDM with $\eta \in [0.001, 0.010]$. (Row 5) Data simulated using RaDDM with $\theta \in (0, 1)$. (Row 6) Data simulated using Hybrid aDDM with $\theta \in (0, 1)$ and $\eta \in [0.001, 0.010]$. Error bars denote the standard error of the mean across simulated subjects.

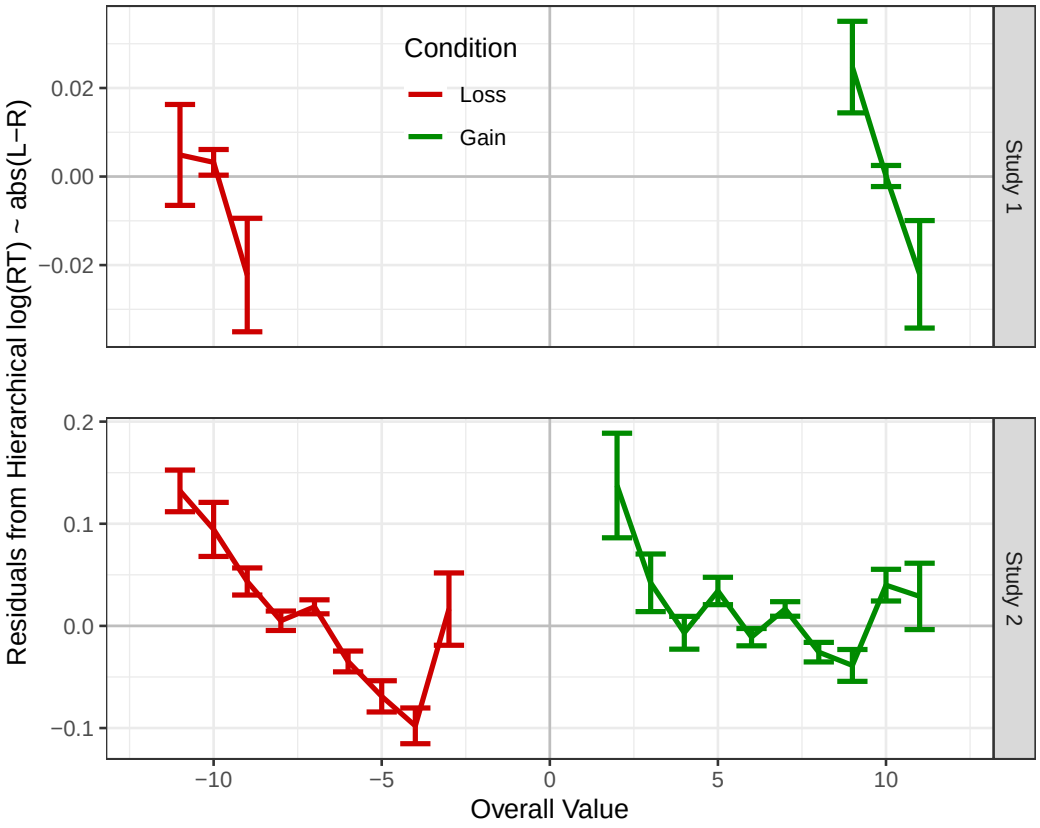

Figure B.4: Relationship between Response Time and Overall Value. We run a mixed effects regression of log(RT) on the absolute difference between the left and right expected values. Then, we plot the residuals as a function of overall value of the choice, separately for study 1 and 2. Error bars denote the standard error of the mean across participant means. Data from the joint dataset.

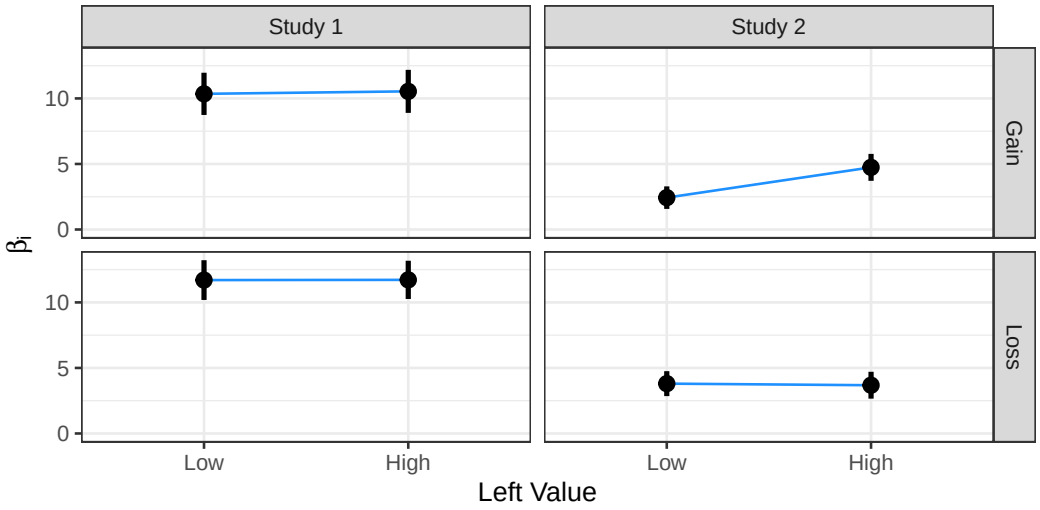

Figure B.5: The Effects of the Interaction of Between Attention and Value on Choice. To measure the effect of the value-attention interaction on choice, Smith and Krajbich (2019) regress choice $\sim \alpha_0 + \alpha_1(V_L - V_R) + \alpha_2(V_L + V_R) + \beta_{Low}(\text{Left Dwell Proportion}|V_L < \text{median}(V_L)) + \beta_{High}(\text{Left Dwell Proportion}|V_L \geq \text{median}(V_L)$. Here, we repeat this analysis, separately by study and condition, and plot β_{Low} and β_{High} for comparison. Bars denote 95% confidence intervals.

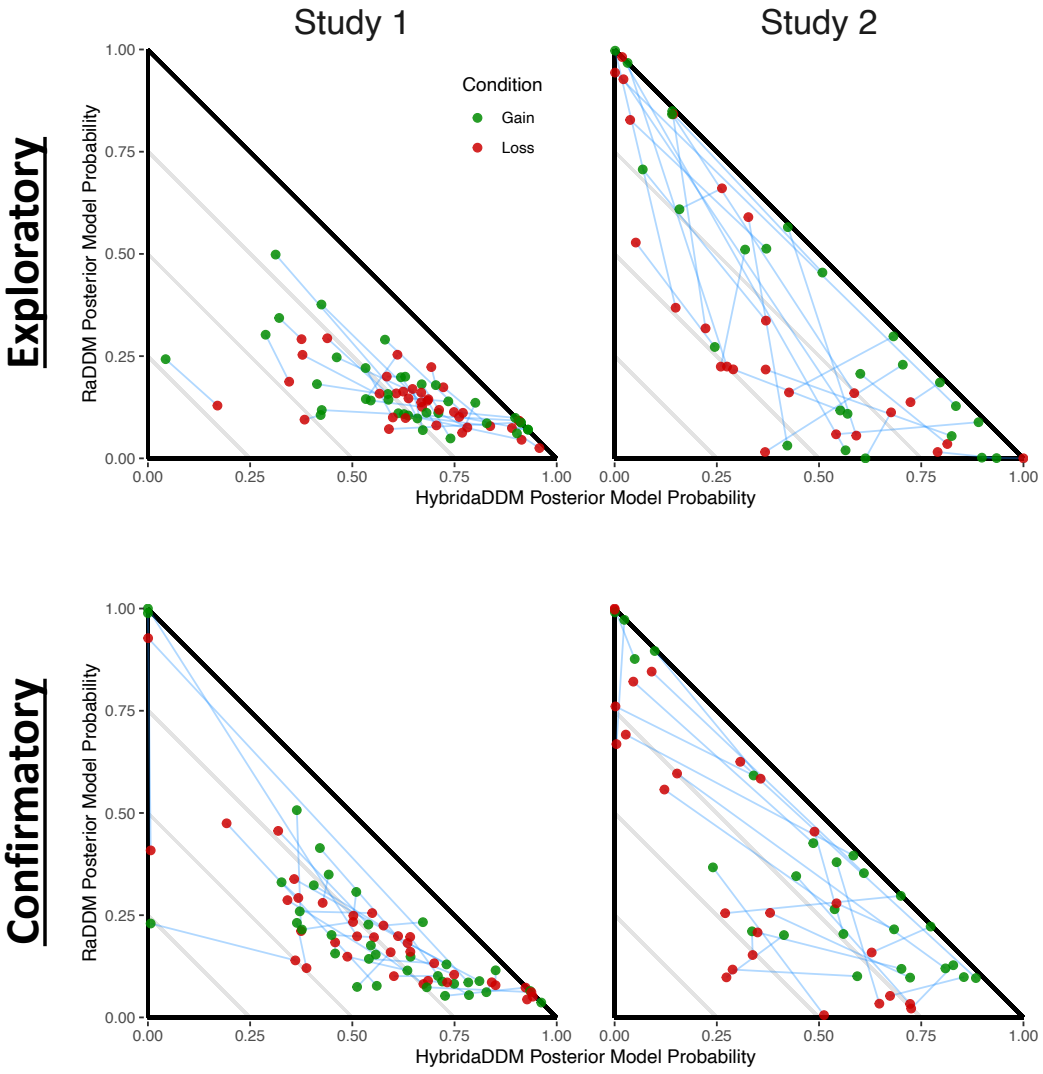

Figure B.6: Model Comparison Between RaDDM, HybridaDDM, and AddDDM, Plotted as a Probability Simplex, Separately by Dataset. For each participant, the posterior model probability of the RaDDM is plotted as a function of the posterior model probability of the HybridaDDM. The posterior model probability of the AddDDM for each individual is therefore 1 minus the sum of their posterior model probabilities for the RaDDM and HybridaDDM. Points closer to the origin indicate participants that were best-fit by the AddDDM, points closer to the top indicate participants that were best-fit by the RaDDM, and points closer to the right indicate participants that were best-fit by the HybridaDDM. Blue lines connect participants across conditions.

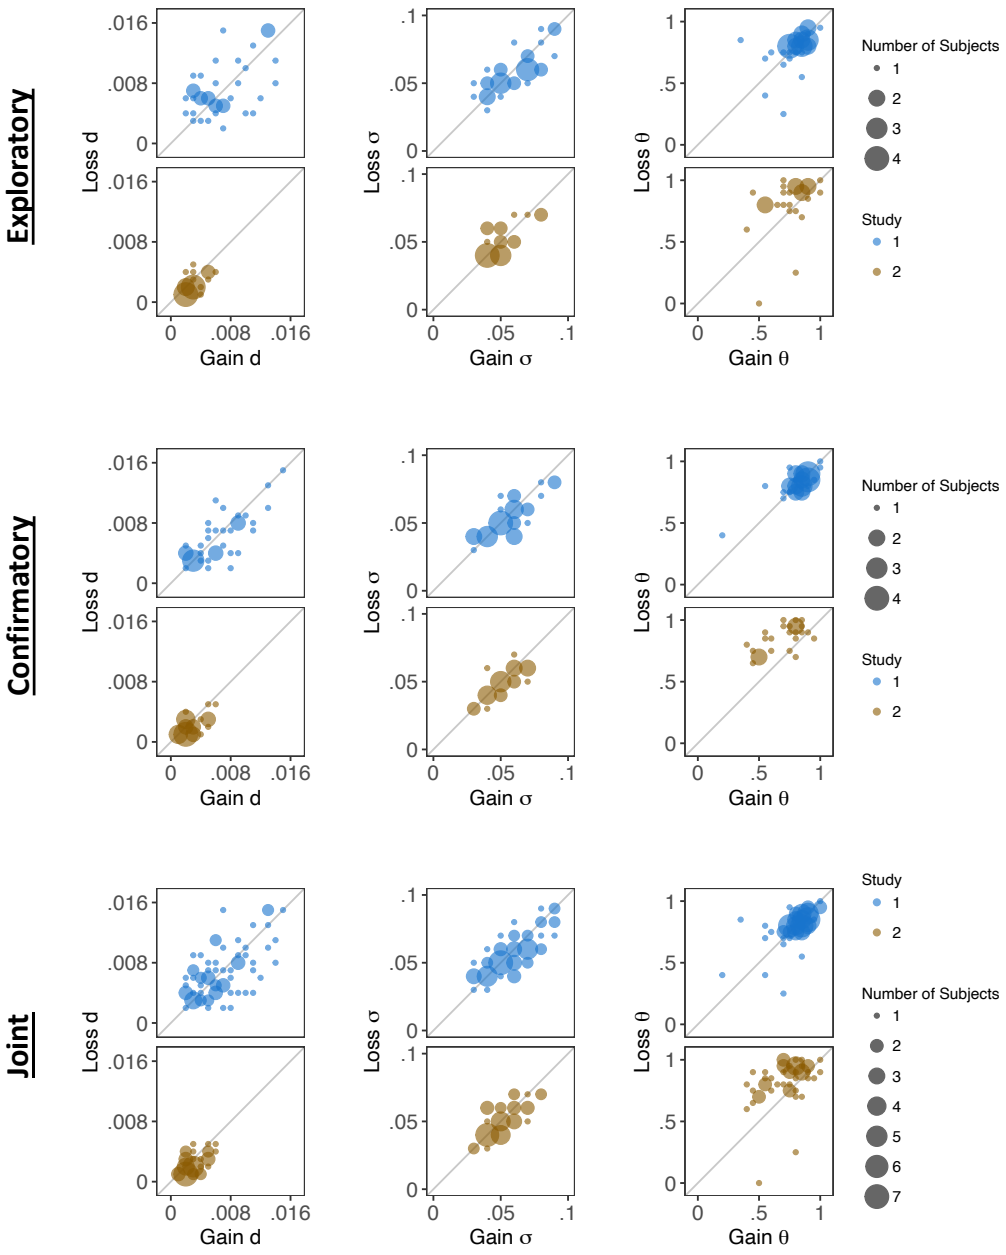

Figure B.7: Participant-Level RaDDM Estimates, Split by Dataset. Estimates of the drift rate (d), noise (σ), and attentional bias (θ) in the loss condition as a function of their counterpart in the gain condition. Grey lines mark equality across the two conditions.

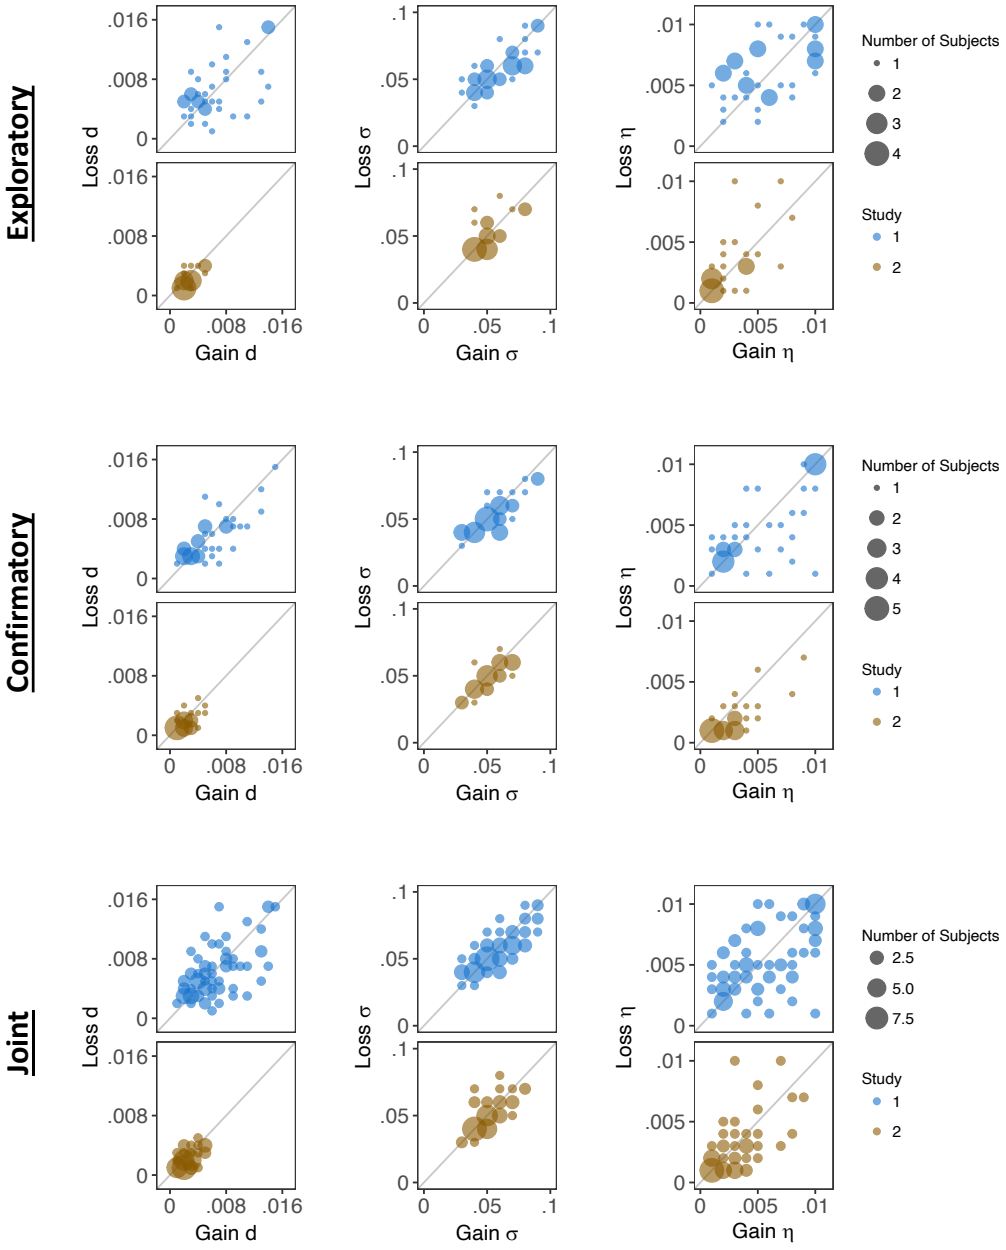

Figure B.8: Participant-Level AddDDM Estimates, Split by Dataset. Estimates of the drift rate (d), noise (σ), and additive attentional bias (η) in the loss condition as a function of their counterpart in the gain condition. Grey lines mark equality across the two conditions.

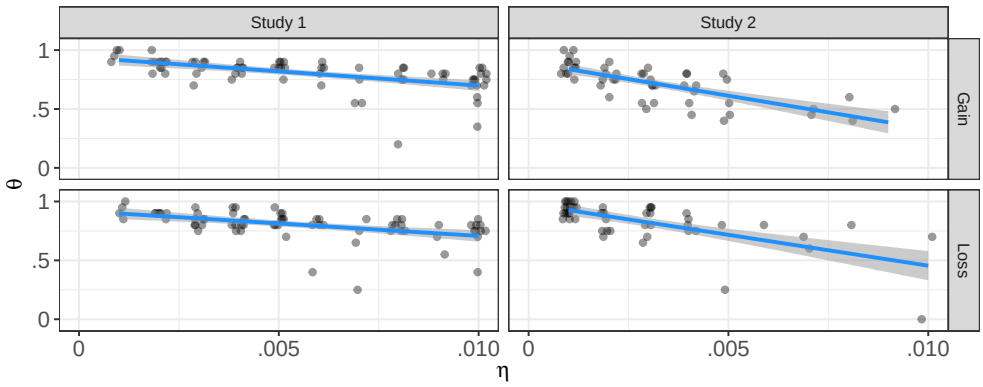

Figure B.9: Comparing Attentional Biases Across Models. Plots depict fitted θ from the RaDDM as a function of the fitted η from the AddDDM, separately by study (columns) and condition (rows). Note that smaller θ's and larger η's correspond to more attentional bias. Blue lines indicate OLS regression line, with grey 95% confidence intervals. η's are slightly jittered for visual clarity.

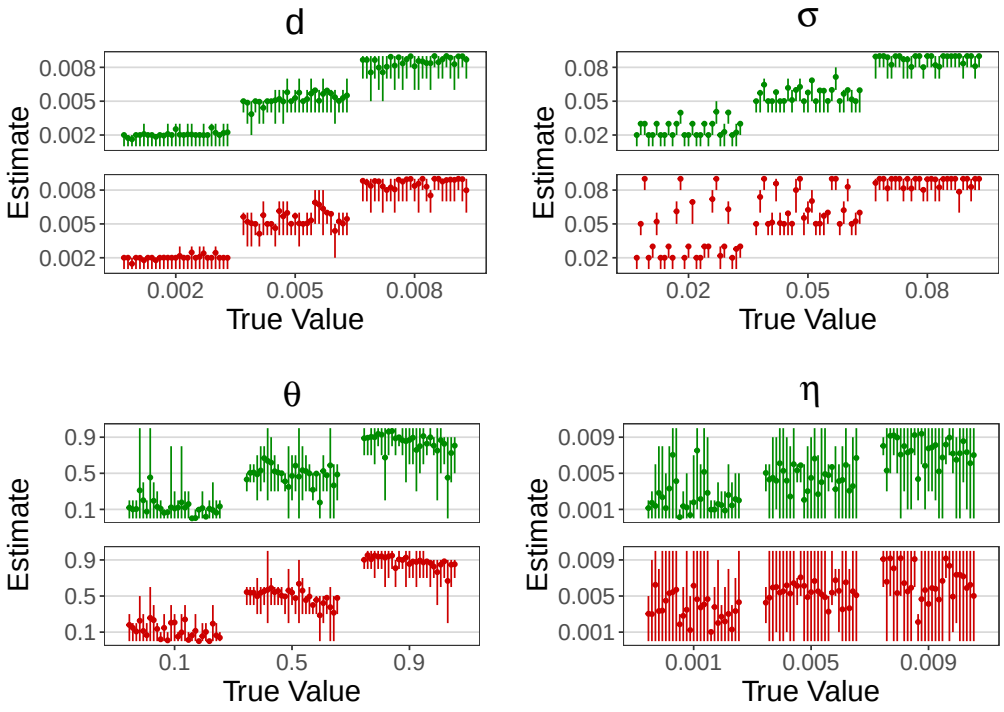

Figure B.10: Hybrid aDDM Parameter Recovery. Estimated parameter estimates as a function of the true data-generating parameters in a parameter recovery exercise for the Hybrid aDDM. Data was simulated using the Hybrid aDDM with parameters from every combination of: $d \in \{0.002, 0.005, 0.008\}$, $\sigma \in \{0.02, 0.05, 0.08\}$, $\theta \in \{0.1, 0.5, 0.9\}$, $\eta \in \{0.001, 0.005, 0.009\}$. Dots represent the mean of the posterior. Lines represent the ≥95% HDI: Posteriors are discrete distributions; therefore, we include the discrete values just outside of the mass containing 95% of the density in order to achieve a minimum of 95% confidence.

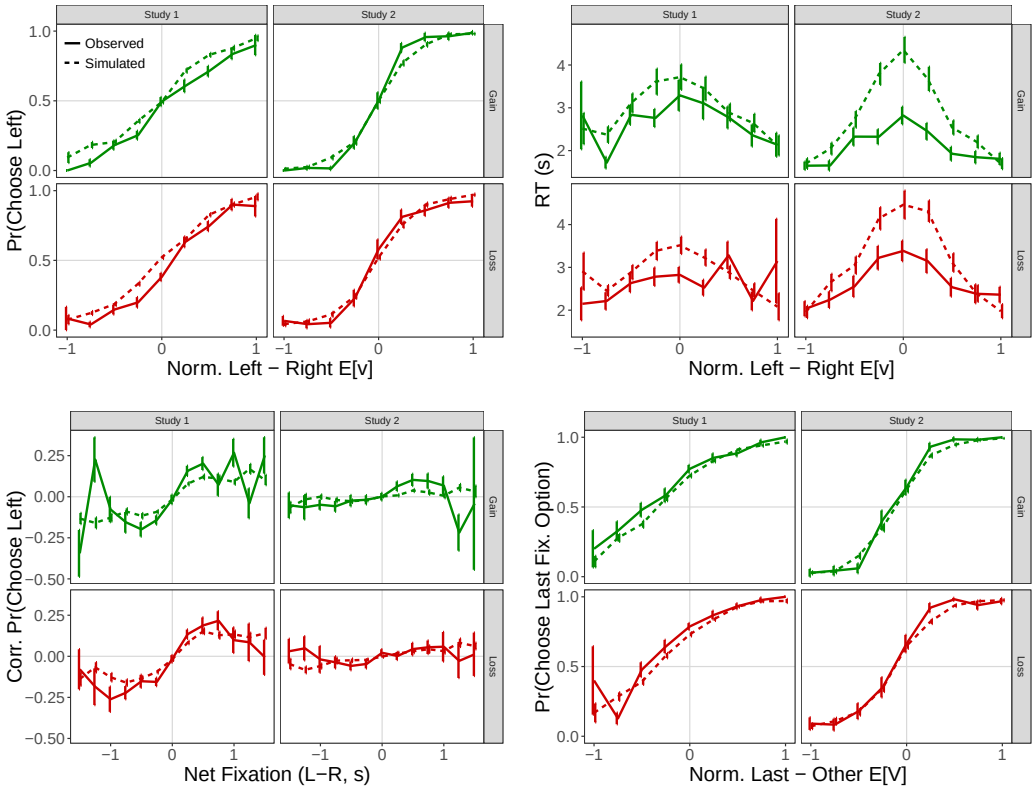

Figure B.11: RaDDM Out-of-Sample Predictions. (Top Left) The probability of choosing left as a function of the normalized relative expected value of the left option, as in Row 1 of Fig. 2.3. (Top Right) Response time as a function of the normalized best minus worst expected value, as in Row 2 of Fig. 2.3. (Bottom Left) Corrected probability of choosing the left option as a function of the net fixation time to the left option, as in Row 1 of Fig. 2.5. The corrected probability is computed by subtracting from each choice observation (coded as 1 if left chosen, and 0 otherwise) the proportion with which left is chosen at each relative expected value. (Bottom Right) The probability of choosing the last fixated option as a function of the normalized net expected value of the left option, as in Row 2 of Fig. 2.5. Rows separate the data by condition, and columns separate by study. Out-of-sample data consists of all trials divisible by 10 from all participants in joint dataset. 10 simulated datasets per participant. Error bars represent the standard error of the mean across all simulations for all participants.

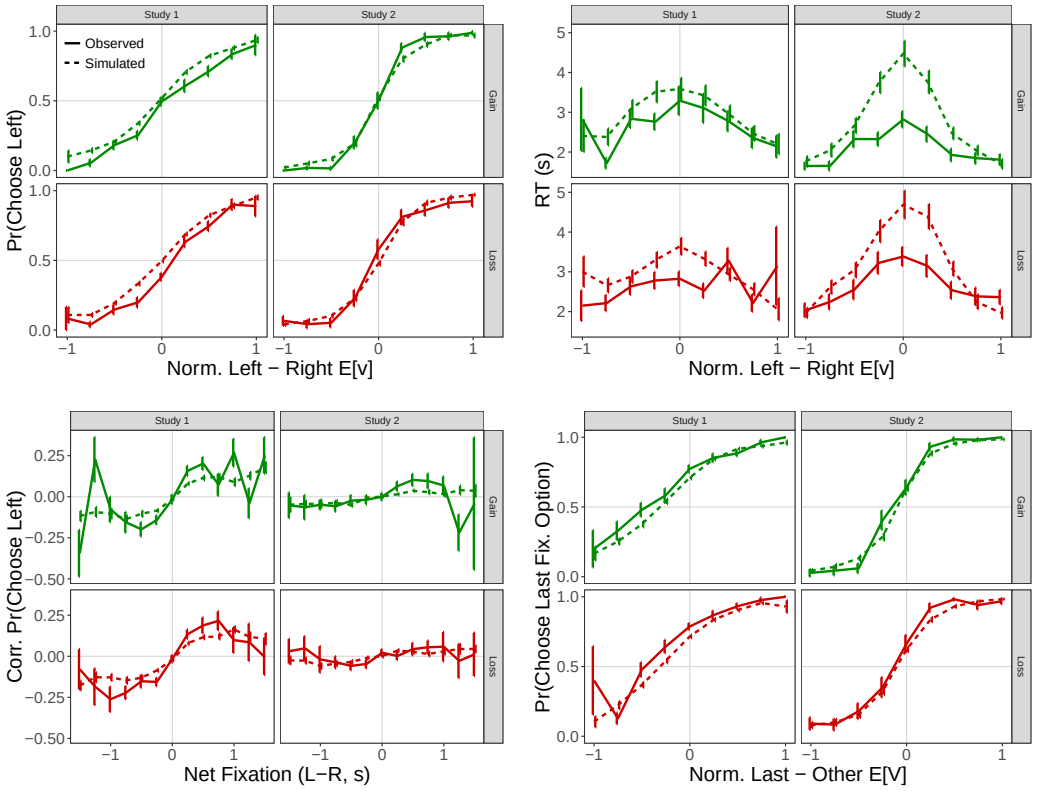

Figure B.12: AddDDM Out-of-Sample Predictions. (Top Left) The probability of choosing left as a function of the normalized relative expected value of the left option, as in Row 1 of Fig. 2.3. (Top Right) Response time as a function of the normalized best minus worst expected value, as in Row 2 of Fig. 2.3. (Bottom Left) Corrected probability of choosing the left option as a function of the net fixation time to the left option, as in Row 1 of Fig. 2.5. The corrected probability is computed by subtracting from each choice observation (coded as 1 if left chosen, and 0 otherwise) the proportion with which left is chosen at each relative expected value. (Bottom Right) The probability of choosing the last fixated option as a function of the normalized net expected value of the left option, as in Row 2 of Fig. 2.5. Rows separate the data by condition, and columns separate by study. Out-of-sample data consists of all trials divisible by 10 from all participants in joint dataset. 10 simulated datasets per participant. Error bars represent the standard error of the mean across all simulations for all participants.

B.2 Tables

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
				Study 1						
Left Chosen	Intercept	-0.17	0.11		-0.40	0.08	*	-0.28	0.07	*
(Logistic)	N. L - R E[V]	2.64	0.13	*	3.05	0.18	*	2.85	0.11	*
(Top)	Loss Condition	-0.16	0.11		0.07	0.12		-0.05	0.08	
	Interaction	0.34	0.17	*	0.14	0.13		0.24	0.10	*
RT (s)	Intercept	2.85	0.31	*	3.56	0.49	*	3.20	0.29	*
(Linear)	N. B - W E[V]	-0.62	0.16	*	-1.26	0.36	*	-0.94	0.20	*
(Middle)	Loss Condition	-0.23	0.24		-0.14	0.37		-0.18	0.22	
	Interaction	0.14	0.15		0.42	0.28		0.28	0.16	
# of Fix.	Intercept	4.16	0.29	*	4.43	0.32	*	4.30	0.21	*
(Linear)	N. B - W E[V]	-0.65	0.14	*	-0.97	0.19	*	-0.81	0.12	*
(Bottom)	Loss Condition	-0.16	0.17		-0.07	0.17		-0.12	0.12	
	Interaction	0.16	0.14		0.12	0.14		0.14	0.09	
				Study 2						
Left Chosen	Intercept	0.12	0.06		0.14	0.06	*	0.11	0.04	*
(Logistic)	N. L - R E[V]	7.00	0.50	*	6.26	0.48	*	6.62	0.35	*
(Top)	Loss Condition	0.10	0.09		-0.07	0.09		0.03	0.07	
	Interaction	-0.49	0.79		0.19	0.66		-0.10	0.52	
RT (s)	Intercept	2.56	0.16	*	2.90	0.29	*	2.73	0.16	*
(Linear)	N. B - W E[V]	-1.02	0.12	*	-1.03	0.15	*	-1.03	0.10	*
(Middle)	Loss Condition	0.70	0.16	*	0.64	0.18	*	0.67	0.12	*
	Interaction	-0.39	0.11	*	-0.25	0.11	*	-0.32	0.08	*
# of Fix.	Intercept	4.21	0.19	*	4.37	0.26	*	4.29	0.16	*
(Linear)	N. B - W E[V]	-1.39	0.16	*	-1.42	0.16	*	-1.41	0.11	*
(Bottom)	Loss Condition	0.57	0.16	*	0.32	0.21		0.44	0.13	*
	Interaction	-0.39	0.12	*	-0.13	0.14		-0.26	0.09	*

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
"N.": Normalized. "B - W rat.": Best - Worst Rating.

Table B.1: Regressions associated with the basic psychometric results in Fig. 2.3.

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Study 1										
1st Fix. Best	Intercept	-0.10	0.06		0.05	0.06		-0.02	0.05	
(Logistic)	N. B - W E[V]	0.22	0.12		-0.07	0.12		0.08	0.09	
(Row 1)	Loss Condition	0.11	0.10		0.02	0.09		0.07	0.07	
	Interaction	-0.32	0.18		-0.06	0.17		-0.19	0.12	
Mid. Fix. Dur.	Intercept	0.71	0.04	*	0.81	0.05	*	0.76	0.03	*
(Linear)	N. B - W E[V]	-0.04	0.02	*	-0.10	0.03	*	-0.07	0.02	*
(Row 3)	Loss Condition	-0.01	0.04		0.00	0.04		0.00	0.03	
	Interaction	0.00	0.02		0.05	0.04		0.02	0.03	
1st Fix. Dur.	Intercept	0.55	0.06	*	0.64	0.06	*	0.59	0.04	*
(Linear)	N. First E[V]	0.02	0.03		-0.01	0.04		0.01	0.02	
(Row 4)	Loss Condition	0.01	0.05		0.01	0.06		0.02	0.04	
	Interaction	-0.01	0.04		-0.01	0.06		-0.02	0.03	
Net Fix. Dur.	Intercept	0.04	0.03		0.03	0.04		0.03	0.02	
(Linear)	N. L - R E[V]	0.18	0.02	*	0.14	0.03	*	0.16	0.02	*
(Row 5)	Loss Condition	-0.06	0.03	*	0.05	0.04		-0.01	0.03	
	Interaction	0.02	0.03		0.15	0.07	*	0.08	0.04	*
Study 2										
1st Fix. Best	Intercept	-0.02	0.08		0.03	0.08		0.00	0.06	
(Logistic)	N. B - W E[V]	0.04	0.12		-0.03	0.12		0.01	0.08	
(Row 1)	Loss Condition	0.05	0.12		-0.05	0.11		0.00	0.08	
	Interaction	-0.07	0.17		0.08	0.17		0.00	0.12	
Mid. Fix. Dur.	Intercept	0.52	0.02	*	0.54	0.03	*	0.53	0.02	*
(Linear)	N. B - W E[V]	-0.07	0.02	*	0.00	0.02		-0.03	0.01	*
(Row 3)	Loss Condition	0.04	0.02		0.08	0.02	*	0.06	0.01	*
	Interaction	0.03	0.02		-0.03	0.02		0.00	0.02	
1st Fix. Dur.	Intercept	0.43	0.02	*	0.47	0.03	*	0.45	0.02	*
(Linear)	N. First E[V]	0.02	0.01		0.03	0.02		0.02	0.01	
(Row 4)	Loss Condition	0.07	0.02	*	0.12	0.03	*	0.09	0.02	*
	Interaction	0.00	0.02		0.00	0.03		0.00	0.01	
Net Fix. Dur.	Intercept	-0.16	0.04	*	-0.20	0.04	*	-0.18	0.02	*
(Linear)	N. L - R E[V]	0.14	0.02	*	0.14	0.03	*	0.14	0.02	*
(Row 5)	Loss Condition	-0.05	0.03		-0.09	0.04	*	-0.07	0.03	*
	Interaction	0.02	0.03		0.01	0.03		0.01	0.02	

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
"N.": Normalized. "B - W rat.": Best - Worst Rating.

Table B.2: Regressions associated with the fixation process results in Fig. 2.4.

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Study 1										
1st Fix. L	Intercept	-1.23	0.28	*	1.29	0.35	*	1.26	0.22	*
(Logistic)	N. L - R E[V]	-0.07	0.10		0.05	0.08		0.07	0.06	
(Top)	Loss Condition	-0.32	0.14	*	0.01	0.16		0.17	0.11	
	Interaction	0.16	0.12		0.05	0.11		-0.05	0.08	
2nd Fix. Dur.	Intercept	0.66	0.05	*	0.73	0.05	*	0.69	0.04	*
(Linear)	N. B - W E[V]	-0.05	0.01	*	-0.08	0.02	*	-0.07	0.01	*
(Middle)	Loss Condition	-0.02	0.04		0.01	0.03		-0.01	0.02	
	Interaction	0.02	0.02		0.06	0.03		0.04	0.02	*
3rd Fix. Dur.	Intercept	0.58	0.05	*	0.66	0.05	*	0.62	0.03	*
(Linear)	N. B - W E[V]	-0.08	0.02	*	-0.10	0.03	*	-0.09	0.02	*
(Bottom)	Loss Condition	-0.01	0.04		0.06	0.04		0.03	0.03	
	Interaction	0.01	0.03		-0.01	0.04		0.00	0.03	
Study 2										
1st Fix. L	Intercept	0.22	0.07	*	0.24	0.06	*	0.24	0.04	*
(Logistic)	N. L - R E[V]	0.01	0.05		0.01	0.05		0.01	0.04	
(Top)	Loss Condition	-0.03	0.05		0.00	0.05		-0.01	0.04	
	Interaction	-0.01	0.07		0.01	0.07		0.00	0.05	
2nd Fix. Dur.	Intercept	0.52	0.03	*	0.62	0.07	*	0.57	0.04	*
(Linear)	N. B - W E[V]	-0.04	0.02	*	-0.08	0.04	*	-0.06	0.02	*
(Middle)	Loss Condition	0.05	0.02	*	0.12	0.04	*	0.09	0.02	*
	Interaction	0.01	0.03		-0.04	0.03		-0.02	0.02	
3rd Fix. Dur.	Intercept	0.51	0.02	*	0.49	0.03	*	0.50	0.02	*
(Linear)	N. B - W E[V]	-0.15	0.03	*	-0.06	0.03	*	-0.11	0.02	*
(Bottom)	Loss Condition	0.04	0.02		0.10	0.04	*	0.07	0.02	*
	Interaction	0.03	0.03		-0.01	0.04		0.01	0.02	

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
"N." : Normalized. "B - W rat.": Best - Worst Rating.

Table B.3: Regressions associated with the additional fixation properties results in Fig. B.2.

		Exploratory			Confirmatory			Joint		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
				Study 1						
Corr. L Chosen	Intercept	0.01	0.01		0.00	0.01		0.00	0.00	
(Linear)	Net Fix. Left	0.37	0.05	*	0.25	0.04	*	0.31	0.04	*
(Top)	Loss Condition	0.00	0.01		0.00	0.01		0.00	0.00	
	Interaction	0.02	0.04		-0.02	0.03		0.00	0.03	
Last Fix. Chosen	Intercept	1.27	0.17	*	1.54	0.17	*	1.40	0.12	*
(Logistic)	N. Last - Other E[V]	2.55	0.14	*	2.83	0.17	*	2.69	0.11	*
(Middle)	Loss Condition	0.05	0.08		-0.03	0.08		0.01	0.06	
	Interaction	0.25	0.19		0.21	0.16		0.23	0.13	
1st Fix. Chosen	Intercept	0.18	0.10		-0.09	0.08		0.04	0.06	
(Logistic)	N. 1st - Other E[V]	2.61	0.12	*	2.94	0.17	*	2.77	0.11	*
(Bottom)	Loss Condition	-0.39	0.11	*	-0.14	0.08		-0.26	0.07	*
	Interaction	0.36	0.10	*	0.14	0.13		0.26	0.10	*
				Study 2						
Corr. L Chosen	Intercept	0.01	0.00	*	0.01	0.00		0.01	0.00	*
(Linear)	Net Fix. Left	0.09	0.02	*	0.06	0.01	*	0.07	0.01	*
(Top)	Loss Condition	0.00	0.01		0.00	0.01		0.00	0.00	
	Interaction	-0.02	0.02		-0.02	0.01		-0.02	0.01	
Last Fix. Chosen	Intercept	0.59	0.13	*	0.74	0.10	*	0.67	0.08	*
(Logistic)	N. Last - Other E[V]	7.04	0.52	*	6.23	0.50	*	6.62	0.36	*
(Middle)	Loss Condition	0.10	0.13		-0.03	0.10		0.03	0.08	
	Interaction	-0.73	0.75		0.20	0.65		-0.23	0.50	
1st Fix. Chosen	Intercept	-0.03	0.06		-0.02	0.05		-0.03	0.04	
(Logistic)	N. 1st - Other E[V]	6.95	0.48	*	6.24	0.48	*	6.60	0.35	*
(Bottom)	Loss Condition	0.08	0.09		-0.12	0.08		-0.03	0.06	
	Interaction	-0.52	0.79		0.08	0.62		-0.20	0.51	

* indicates significance in all data sets at the 95% confidence level.
* indicates a significant effect that was not present in all three data sets.
"N.": Normalized. "B - W rat.": Best - Worst Rating.

Table B.4: Regressions associated with the attentional choice biases results in Fig. 2.5.

Loss Condition

		RaDDM			AddDDM			HybridaDDM		
Dept. Var.	Indept. Var.	Est.	SE		Est.	SE		Est.	SE	
Study 1	β_0 Intercept	0.56	0.11	*	0.73	0.10	*	0.63	0.10	*
$log(RT)$	β_1 N. Abs. Net Value	-0.79	0.06	*	-0.81	0.07	*	-0.74	0.07	*
(Linear)	β_2 N. Overall Value	-0.20	0.08	*	-0.02	0.08		-0.12	0.07	
Study 2	β_0 Intercept	1.10	0.10	*	1.14	0.10	*	1.10	0.09	*
$log(RT)$	β_1 N. Abs. Net Value	-0.76	0.07	*	-0.76	0.07	*	-0.76	0.07	*
(Linear)	β_2 N. Overall Value	-0.13	0.02	*	-0.02	0.02		-0.10	0.03	*

Gain Condition

Study 1	β_0 Intercept	1.05	0.10	*	0.86	0.11	*	0.89	0.12	*
$log(RT)$	β_1 N. Abs. Net Value	-0.70	0.07	*	-0.70	0.07	*	-0.68	0.07	*
(Linear)	β_2 N. Overall Value	-0.35	0.08	*	-0.13	0.07		-0.18	0.08	*
Study 2	β_0 Intercept	1.21	0.09	*	1.15	0.09	*	1.18	0.08	*
$log(RT)$	β_1 N. Abs. Net Value	-0.95	0.06	*	-0.97	0.06	*	-0.95	0.05	*
(Linear)	β_2 N. Overall Value	-0.08	0.03	*	0.02	0.03		-0.04	0.03	

* indicates significance in all data sets at the 95% confidence level.
"N.": Normalized. "Abs.": Absolute.

Table B.5: Model Comparison Regression Test in Simulated Data Using RaDDM, AddDDM, and Hybrid aDDM

B.3 Additional Text

Attentional Manipulation in Study 2

Using the forced fixation cross to manipulate the location of the first fixation in study 2, we nudged choices away from the target option by a small but significant amount in the loss condition, but not in the gain condition (Fig. B.13 top row). RTs did not significantly change in response to attentional manipulations (Fig. B.13 bottom row).

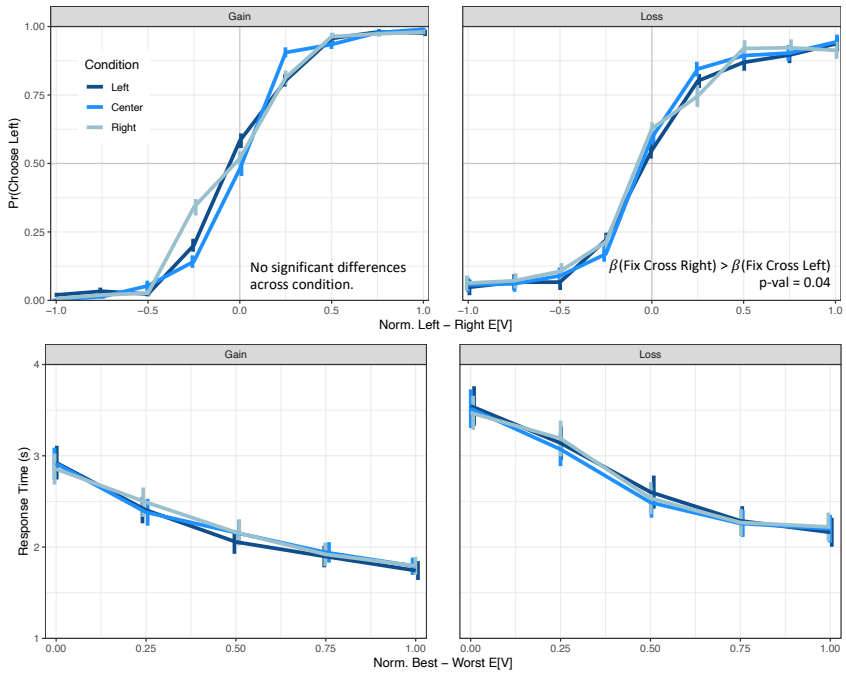

Figure B.13: Basic Psychometrics with Attentional Manipulations. The top row depicts the probability of choosing the left lottery as a function of the normalized expected value difference between the left and right lottery. Value differences are normalized by the maximum possible value difference. The bottom row depicts response times as a function of the normalized choice difficulty. Choice difficulty is measured by the difference between the best and worst expected value and is normalized by the maximum possible choice difficulty. Columns denote the location of the fixation cross. Error bars denote standard error of the mean across participants.

In study 2, we used counterbalancing to ensure that altering the location of the forced fixation cross did not affect the probability of first fixating on the best option (Fig. B.14 top row). This attentional manipulation increased first fixation durations by roughly 100 ms on average without altering middle fixation durations (Fig. SB.14 middle and bottom row). This biased net fixation durations in the direction of the manipulation (Fig. B.14 fourth row), partially explaining the right-side attentional

favoritism seen in the bottom row of Fig. 2.4.

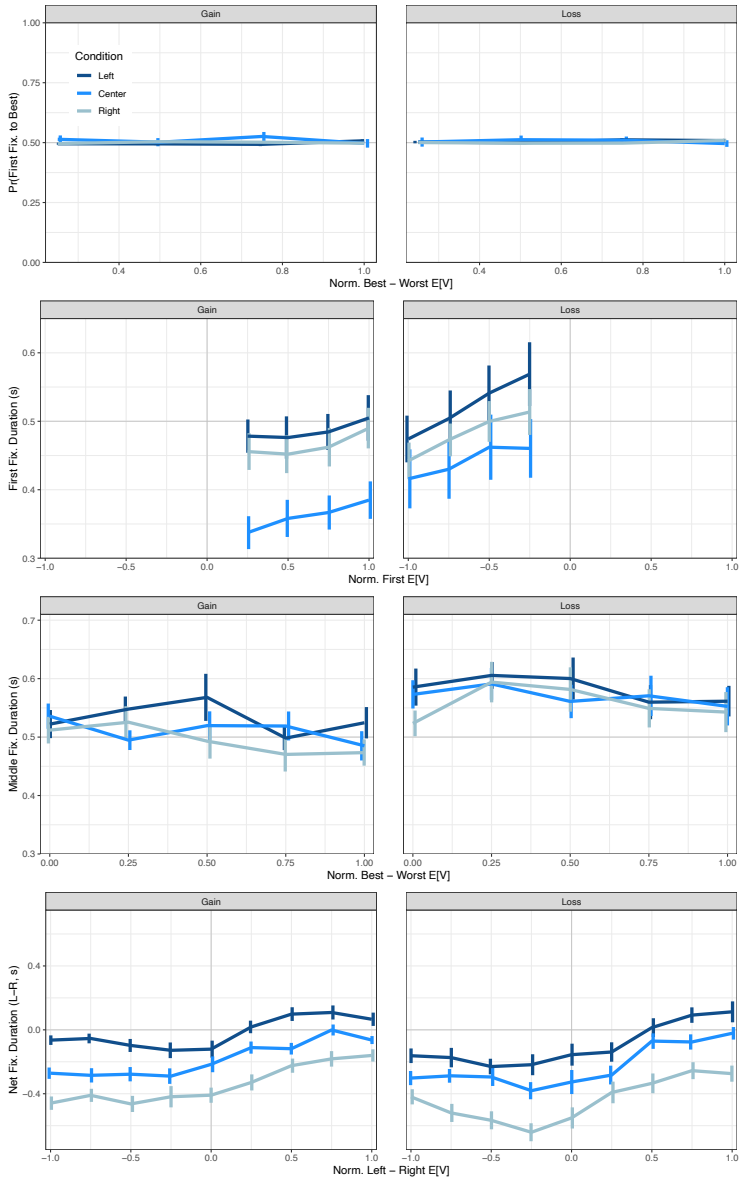

Figure B.14: Fixation Process with Attentional Manipulations. The first row depicts the probability of first fixating on the best lottery as a function of normalized choice difficulty. The second row depicts the first fixation durations as a function of normalized expected value of the first seen lottery. The third row depicts average middle fixation durations as a function of normalized choice difficulty. The fourth row depicts the net fixation duration to the left lottery as a function of the normalized expected value difference. Columns denote the location of the fixation cross which manipulates the location of first fixation. Error bars denote standard error of the mean across participants. Data from the joint dataset.

Manipulating the location of the first fixation using a forced fixation cross at the

start of each trial did not affect the magnitude of net fixation bias (Fig. B.15 first row). Non-centered fixation crosses reduced the accuracy with which the first and last fixated item was chosen when relative values were close to 0 (Fig. B.15 second row). In aversive choices, when the relative value of the first fixated item was equal to 0, the attentional manipulation biased induced a first fixation bias in the direction of the forced fixation cross (Fig. B.15 third row). In choices between gains, this is true for left fixation crosses, but not for right.

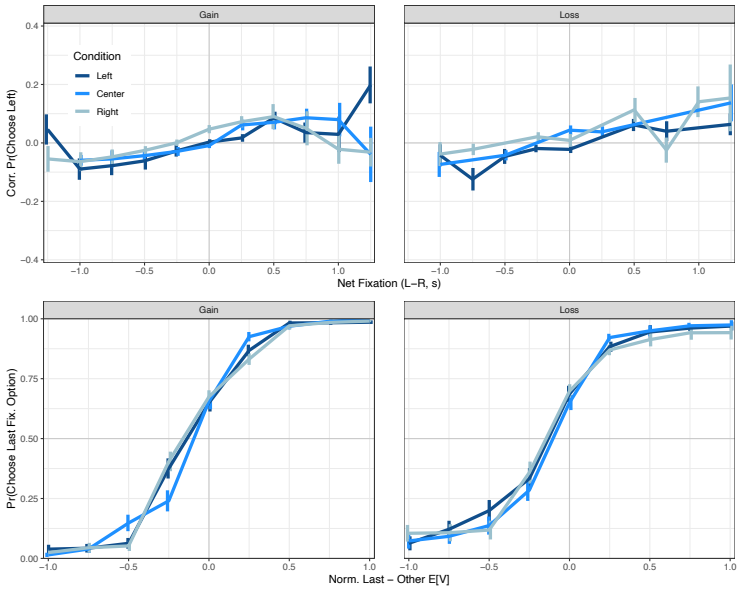

Figure B.15: Attentional Choice Biases with Attentional Manipulations. The top row depicts the corrected probability of choosing the left lottery as a function of the net fixation time to the left lottery. Corrected probabilities are calculated by taking the choice observation (1=left, 0=right) and subtracting the proportion of times left was chosen at each value difference. The bottom row depicts the probability of choosing the last fixated option as a function of the normalized relative expected value of the last fixated lottery. Columns denote the location of the fixation cross which manipulates the location of first fixation. Error bars denote standard error of the mean across participants. Data from joint dataset.

Appendix C

SUPPLEMENTARY MATERIALS FOR CHAPTER III

C.1 Figures

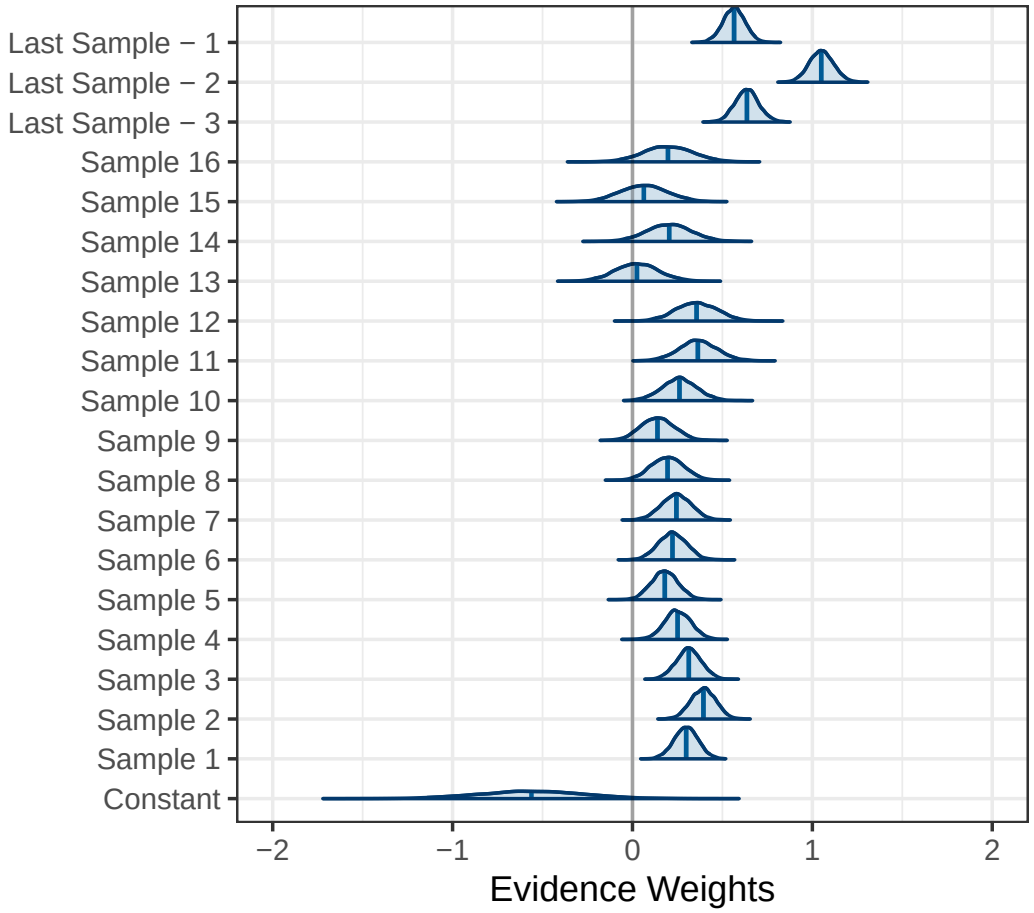

Figure C.1: Evidence Weights for Early vs. Late. Evidence weights for the first 16 samples in a trial and Last Sample - 3 through Last Sample - 1. Evidence weights were calculated in the same manner as shown in Eq. 3.3 and plotted in Fig. 3.3.

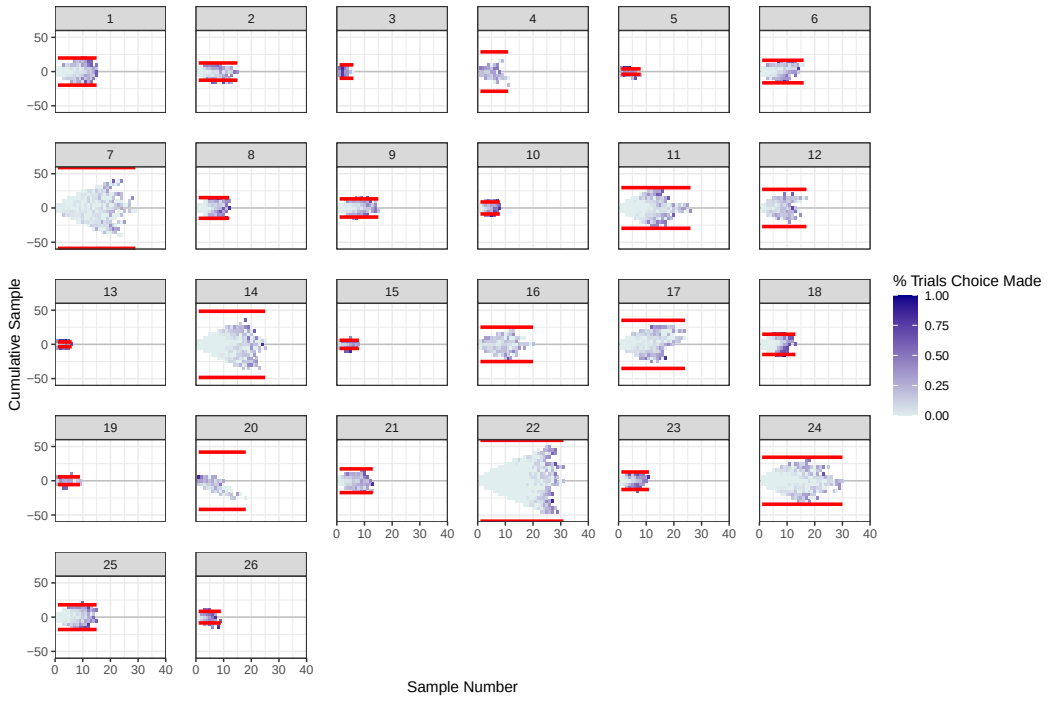

Figure C.2: Termination Rule over Cumulative Sample for Every Participant. We efine a state by the combination of cumulative sample and sample number. For every state, we plot the proportion of trials in which a choice was made in this state, conditional on having reached this state. If less than 5 trials reached a state, that state was excluded. Red lines plot the estimated decision boundaries from Theorem 1.

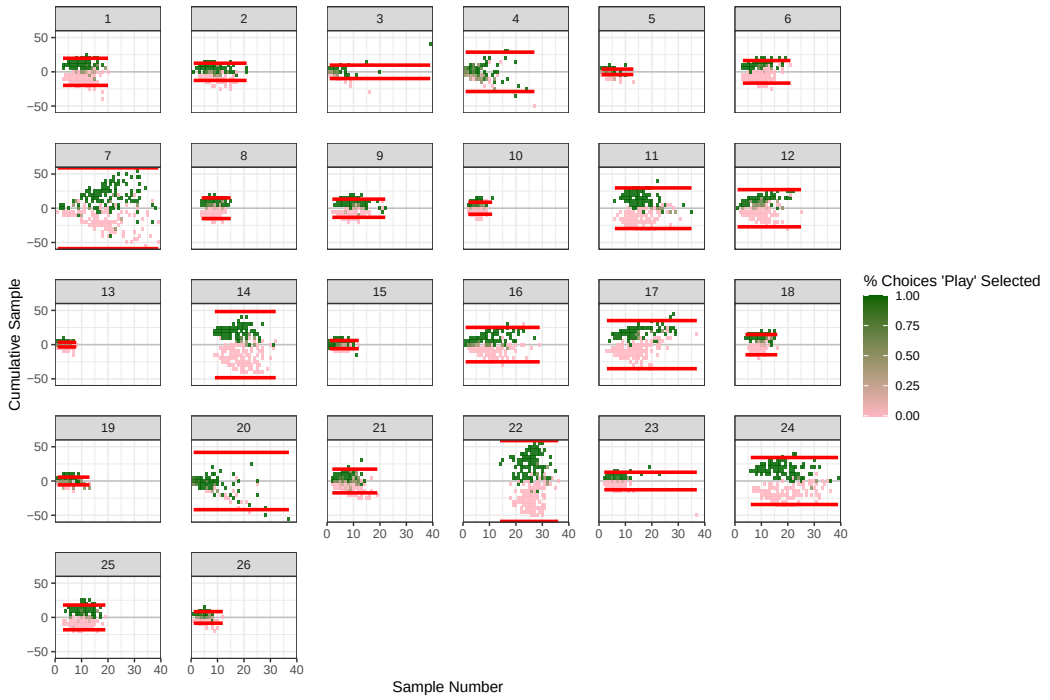

Figure C.3: Policy over Cumulative Sample for Every Participant. We define a state by the combination of cumulative sample and sample number. For every state, we plot the proportion of trials in which Play was selected, conditional on a choice being made in this state. Red lines plot the estimated decision boundaries from Theorem 1.

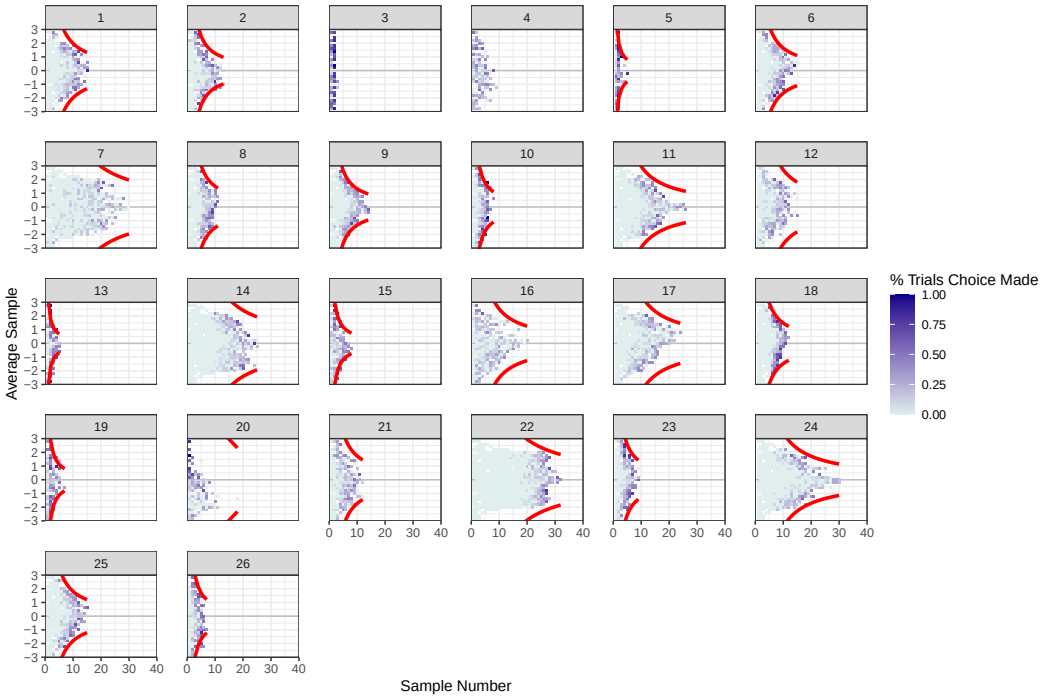

Figure C.4: Termination Rule over Average Sample for Every Participant. Define a state by the combination of average sample and sample number. For every state, we plot the proportion of trials in which a choice was made in this state, conditional on having reached this state. If less than 5 trials reached a state, that state was excluded. Red lines plot the estimated decision boundaries from Theorem 1.

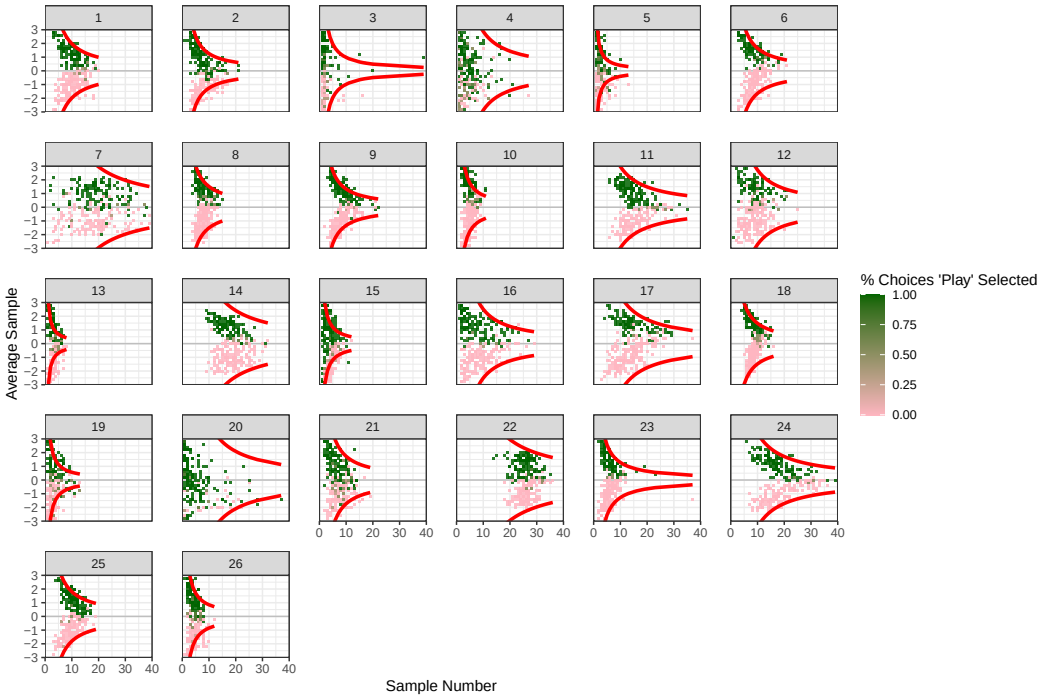

Figure C.5: Policy over Average Sample for Every Participant. Define a state by the combination of average sample and sample number. For every state, we plot the proportion of trials in which Play was selected, conditional on a choice being made in this state. Red lines plot the estimated decision boundaries from Theorem 1.

C.2 Tables

Row	Dept. Var.	Indept. Var.	Estimate	Standard Error	
1	Play	Intercept	-0.43	0.21	*
	(Logistic)	Machine Mean	1.96	0.19	*
2	Play	Intercept	-0.52	0.27	
	(Logistic)	Avg. Sample	3.43	0.38	*
3	# Samples	Intercept	6.90	1.59	*
	(Linear)	Abs. Machine Mean	-1.22	0.23	*
4	# Samples	Intercept	7.08	1.77	*
	(Linear)	Abs. Avg. Sample	-2.11	0.33	*

* indicates significance at the 95% confidence level.
"Avg.": Average.
"Abs.": Absolute.

Table C.1: Regressions associated with the psychometrics results in Fig. 3.2.