



Original Articles

An attentional drift diffusion model over binary-attribute choice



Geoffrey Fisher

Cornell University, Cornell SC Johnson College of Business and Dyson School of Applied Economics and Management, 340C Warren Hall, Ithaca, NY 14853, United States

ARTICLE INFO

Article history:

Received 6 May 2016

Revised 5 June 2017

Accepted 7 June 2017

Keywords:

Attention

Drift diffusion model

Sequential sampling models

Preferences

Multi-attribute choice

ABSTRACT

In order to make good decisions, individuals need to identify and properly integrate information about various attributes associated with a choice. Since choices are often complex and made rapidly, they are typically affected by contextual variables that are thought to influence how much attention is paid to different attributes. I propose a modification of the attentional drift-diffusion model, the binary-attribute attentional drift diffusion model (baDDM), which describes the choice process over simple binary-attribute choices and how it is affected by fluctuations in visual attention. Using an eye-tracking experiment, I find the baDDM makes accurate quantitative predictions about several key variables including choices, reaction times, and how these variables are correlated with attention to two attributes in an accept-reject decision. Furthermore, I estimate an attribute-based fixation bias that suggests attention to an attribute increases its subjective weight by 5%, while the unattended attribute's weight is decreased by 10%.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Except for very simple and familiar choices, most decisions require the identification and weighting of multiple attributes. Examples include choosing between two meals that differ in their taste, nutrition, and costs, or choosing between slot machines that differ in the likelihood and size of the potential rewards. Given their pervasiveness, understanding the algorithms that we use to make choices over alternatives with several attributes, and how they are affected by contextual variables, is a central question in psychology, economics, and neuroscience (Bussemeyer & Johnson, 2004; Fehr & Rangel, 2011; Glimcher & Fehr, 2014; Mas-Colell, Whinston, & Green, 1995).

While much evidence suggests we differentially weight attributes in decision-making, the extent to which these weights are influenced by attention has not been resolved. For instance, suppose a restaurant menu contains a daily special of steak with a side of green beans, and that a consumer enjoys steak, but dislikes green beans. Is the probability that the consumer orders the steak influenced by contextual variables (e.g., how the menu is presented) that change the relative attention paid to the steak and the green beans at the time of choice? Are there models capable of providing a quantitative explanation of these effects? These questions are important because, as hinted in the example, many choices require weighting attributes properly, which might be

impaired in the presence of the attentional effects hypothesized here.

This paper proposes and tests a modification of the attentional drift diffusion model (Krajbich, Armel, & Rangel, 2010; Krajbich, Lu, Camerer, & Rangel, 2012; Krajbich & Rangel, 2011) related to these effects, which I call the binary-attribute attentional drift diffusion model (baDDM). The model details the choice process by modeling how attention to two attributes, at the level of random eye fixations between those attributes, alters individual choices in an accept or reject decision.

The model builds on several main literatures. First, previous work has shown sequential sampling models of decision-making, such as the Drift-Diffusion model (Ratcliff, 1978; Ratcliff, Chierian, & Segraves, 2003; Ratcliff & Smith, 2004; Ratcliff, Smith, Brown, & McKoon, 2016), leaky-accumulator model (Usher & McClelland, 2001), Decision Field Theory (DFT) (Bussemeyer & Diederich, 2002; Bussemeyer & Townsend, 1992, 1993; Diederich, 1997; Roe, Bussemeyer, & Townsend, 2001), and the attention drift diffusion model (aDDM) (Fehr & Rangel, 2011; Krajbich & Rangel, 2011; Krajbich et al., 2010, 2012) provide accurate quantitative accounts of how choice probabilities and response times vary with properties of the choice options. Within the literature, there are varying classes of sequential sampling models but many assume that choices are made using a relative value signal that is dynamically computed by integrating an instantaneous noisy measure of the desirability of options. Once the accumulated relative value signal becomes sufficiently strong in favor of one of two options, a choice is made. Furthermore, a growing body of evidence from

E-mail address: gwf25@cornell.edu

neuroscience has found that the implementation of certain sequential integrator models is biologically plausible (Britten, Shadlen, Newsome, & Movshon, 1992; Gold & Shadlen, 2007; Hare, Schultz, Camerer, O'Doherty, & Rangel, 2011; Heekeren, Marrett, Bandettini, & Ungerleider, 2008; Rangel & Clithero, 2013).

Two broad classes of sequential sampling models are particularly related to this paper. The first concerns sequential sampling models that detail how multi-attribute decisions are made (Bhatia, 2013; Trueblood, Brown, & Heathcote, 2014; Tsetsos, Chater, & Usher, 2012; Usher & McClelland, 2004; Wollschläger & Diederich, 2012). A subset of these models treats choices as the accumulation of noisy evidence over time, although not all models of this class utilize momentary random fluctuations in preferences. Additionally, in some cases these multi-attribute models are able to incorporate attention effects. For instance, work in DFT has modeled attentional changes to attributes by appealing to a dynamic attention function that weights information over time, and Wollschläger and Diederich (2012) take a similar approach in their setting. Trueblood et al. (2014) use explicit attention weights that vary depending on how easily attribute values can be discriminated in their model of multi-attribute choice. In their model, attention weights are not meant to quantify the observed distribution of attention throughout a decision, but instead seek to capture the general trend that similar attributes receive more attention than vastly different ones. Additionally, Bhatia (2013) introduced a connectionist network that allowed more accessible attributes to be more likely to influence preferences. In the model, preferences are determined by weighting sums of attribute values where attributes with larger amounts also receive larger weights in decision-making. Although the models referenced above have differences in the how preferences and choices are formed, all are focused on detailing how the quantitative relationship between various attributes and their values impact decision-making.

A second class of relevant work consists of multi-stage sequential sampling models, some of which also model multi-attribute decision-making (Diederich, 1995, 1997; Diederich, 2015; Diederich & Oswald, 2014; Diederich & Oswald, 2016; Holmes, Trueblood, & Heathcote, 2016; Ratcliff, 1980). Multi-stage models explicitly represent evidence for different processing stages of a decision rather than combining all information into one source of evidence, which had previously described the majority of sequential sampling models found in the literature. This multi-stage approach began by allowing for varying drift rates in a drift diffusion model (Ratcliff, 1980) and has progressed to modeling the switching of attention between options and attributes throughout the course of a decision. Related to this paper, previous work in the aDDM has allowed the drift rate to vary depending on which of several options is currently attended (Krajbich & Rangel, 2011; Krajbich et al., 2010, 2012), though this model has only been extended to choice over a small number of options. Nevertheless, explicitly relating fixations to information accumulation and drift rate changes allows a natural extension to better understanding how we make decisions with more than one attribute, which is related to the model presented here. Relatedly, Diederich and Oswald (2016) propose a sampling model for multi-attribute choice that allows a separate sampling process for each attribute and for attention to switch between different attributes throughout the decision. They use numerical calculations of their model to demonstrate that the order in which attributes are processed can influence choices, but do not analyze empirical data. Although their model would need to be further specified in order to easily adapt to various choice environments and their model did not utilize fixation data, their work takes an important step in detailing how attentional distributions to attributes at the time of choice can influence decisions.

Additionally, this paper adds to a large literature that uses process-tracing methods to understand the decision process (Camerer & Johnson, 2004; Glöckner & Herbold, 2011; Johnson, Schulte-Mecklenbeck, & Willemsen, 2008; Russo & Doshier, 1983; Russo & Rosen, 1975; Willemsen, Böckenholt, & Johnson, 2011; Horstmann, Ahlgrimm, & Glöckner, 2009; Orquin & Mueller Loose, 2013; Towal, Mormann, & Koch, 2013). While much of this work makes use of eye tracking, others test process-based models by tracking mouse movements on a computer screen. Largely, previous work using these methods has broadly confirmed many predictions consistent with decisions being made by different classes of sequential integrator models. Relatedly, a portion of this work has focused on how well alternative models, such as heuristic models of choice, can explain behavior (Payne, Bettman, & Johnson, 1992). While certain heuristics can lead to particular attentional patterns (Day, 2010; Day, Lin, Huang, & Chuang, 2009; Renkewitz & Jahn, 2012), there is currently little evidence to suggest that the particular heuristic used can determine attentional deployment or that the underlying heuristic can be inferred from the distribution of attention (Orquin & Mueller Loose, 2013; Knoepfle, Yao-yi Wang, & Camerer, 2009; Reutskaja, Nagel, Camerer, & Rangel, 2011).

The model proposed in this paper expands on the work above in a number of ways. First, it extends the previous theory and applications of the aDDM. Formerly, the aDDM has been used to estimate how attention biases the drift rate depending on which of several choice options is currently fixated. This operationalizes by applying a fixation bias parameter to the unattended option so that its value is discounted in the evidence accumulation process. The baDDM described here extends this model to cover a simple binary-attribute choice environment in which an individual accepts or rejects a consumption option. Critically, the model and experimental design allow for separate estimation for the degree to which the weight of the attended attribute is increased as well as the degree to which the weight of the unattended attribute is decreased. Estimating multiple fixation bias parameters that describe how attribute weights change over the course of a decision may yield new insights compared to modeling a single fixation bias. Furthermore, since the baDDM investigates an accept-reject choice with two attributes, the results can help us understand the additional tasks that models such as these are able to accurately capture, but also to what extent they can fail. By pushing these limits, we may ultimately be able to design rigorous models that more accurately capture human behavior across a variety of contexts.

Second, the work here extends previous multi-attribute and multi-stage sequential sampling models by collecting and incorporating physiological data on attention, as measured by fixations, throughout the duration of a choice with two attributes. I estimate the model using choices, response times, and fixation data and test how well the model can explain observed patterns that subjects display. Although several previous models of multi-attribute choice are able to incorporate attention effects to varying degrees (e.g., Bhatia, 2013; Trueblood et al., 2014; Wollschläger & Diederich, 2012) they do not explicitly allow fixation information at the time of choice and many do not test their predictions in out of sample data. Despite focusing on a simplified version of multi-attribute choice, which this paper refers to as binary-attribute choice, the setting here can help understand how fixation data can be fit to novel tasks and can ultimately better inform, design, and test models that are grounded in more traditional multi-attribute choice settings. Furthermore, similar to previous work in the aDDM and other multi-stage sequential sampling models, the baDDM also allows for varying drift rates and permits those drift rates to vary as a function of the currently attended information. Although the baDDM is highly related to Diederich

and Oswald (2016), the estimation and tests of the highly specified model found here can deepen our understanding of simple binary-attribute choice by computationally analyzing the extent to which tractable sequential sampling models that incorporate fixation information can fit empirical data.

To test the model, I conduct a laboratory experiment where participants make decisions over whether to consume pairs of foods, or bundles, while I record their eye movements between two attributes. A subset of the data estimates parameters from the model while the remaining data is used to test the model's predictions. Critically, the results provide a quantitative estimate for how attending to particular attributes of a choice can alter the weights those attended features receive when computing value: I find subjects overweight the currently attended attribute and underweight the unattended attribute.

2. Methods

2.1. Subjects

Forty-six subjects recruited from the California Institute of Technology community participated in the experiment (63% male; mean age = 26.2). No participant who completed the entire experiment was excluded from analysis, although the experiment was stopped partway through for six subjects who either did not exhibit enough variance in their behavior to generate enough experimental trials to complete the experiment, failed an eye tracking calibration, or decided to leave partway through the experiment. Based on previous work, I planned to collect at least 4000 trials for model fitting, which required at least 40 subjects. All subjects had normal or corrected-to-normal vision with the use of either contact lenses or glasses. Participants were paid a \$5 show-up fee and received an additional \$25 upon successful completion of the experiment. The local Institutional Review Board approved the study.

2.2. Task

Subjects were asked to fast for four hours prior to the start of the experiment. Compliance was verified through self-report upon subject arrival, and was required for participation.

The experiment consisted of three tasks. Subjects were informed of this at the outset, but the tasks were only described to them just before each took place.

In Rating Task 1, subjects performed a liking-rating task over individual snack food items shown in a computer screen, one at a time. The image size was 300 × 300, with a screen resolution of 1280 × 1024. Subjects were asked to enter a liking rating for each food using an integer scale (−3 to 3, framing: “How much would you enjoy that particular food at the end of today’s experiment?”). The ratings were entered using the bottom row of the keyboard. Subjects could take as long as desired to enter each rating. Thirty unique foods were rated and each food was shown twice to each subject in random order.

The foods were selected based on previous studies (Plassmann, O’Doherty, & Rangel, 2007, 2010) and contained eighteen foods that were consistently rated as appetitive by previous subjects, and twelve foods that were consistently rated as aversive.

For every subject, I averaged the two ratings provided for each food in order to create subject-specific food classes. Snacks with a positive average rating were labeled as “appetitive,” snacks with a negative average rating were classified as “aversive,” and foods with a zero average rating were omitted from the remaining tasks. On average, subjects had 17 appetitive foods and 11 aversive foods. See Table 1 for details.

Table 1

Food stimuli used in the experiment. Each stimulus contains the mean rating across subjects, with standard deviation below in parentheses.

Food	Rating	Food	Rating
Ghirardelli Milk Chocolate	2.24 (1.02)	Chocolate Pudding	0.60 (2.10)
KitKat Candy Bar	2.16 (1.14)	Almond Joy Candy Bar	0.37 (2.17)
Twix Candy Bar	1.95 (1.28)	Tootsie Rolls	0.23 (1.55)
Milano Cookies	1.91 (1.43)	Canned Tuna	0.09 (2.06)
Crunch Bar	1.86 (1.23)	Canned White Meat Chicken	−0.70 (1.81)
Peanut M&M’s	1.76 (1.10)	Canned Sweet Peas	−0.84 (1.76)
Reese’s Peanut Butter Cups	1.70 (1.70)	Canned Garbanzo Beans	−0.88 (1.66)
MilkyWay Candy Bar	1.68 (1.24)	Canned Vienna Sausage	−1.18 (1.87)
Snickers Candy Bar	1.65 (1.35)	Canned Sardines	−1.24 (1.87)
Oreos	1.57 (1.30)	Canned Artichoke Hearts	−1.33 (1.84)
3 Musketeers Candy Bar	1.45 (1.52)	Canned Deviled Ham Spread	−1.38 (1.65)
Doritos Cool Ranch Chips	1.43 (1.44)	Soy Sauce	−1.45 (1.49)
Nature Valley Granola Bar	1.12 (1.33)	Canned Spinach	−1.45 (1.61)
Butterfinger Candy	0.95 (1.69)	Pureed Carrots	−1.68 (1.38)
Flamin’ Hot Cheetos	0.72 (2.12)	Pureed Green Beans	−1.84 (1.38)

In Rating Task 2, subjects saw bundles of two foods on the screen and had to provide liking ratings over the bundles, using the same integer scale (−3 to 3). In particular, subjects were asked to rate “How much would you enjoy taking at least three bites from both of the foods shown on the screen?” Every bundle contained one appetitive food and one aversive food from the previous round, and subjects could take as long as needed to enter their ratings. As shown in Fig. 1, one of the items was shown in the left and the other on the right, with their location randomized every trial. The number of trials in this task varied across subjects as they were asked to rate every potential bundle made of one appetitive and one aversive food.

Finally, subjects participated in a Choice Task (Fig. 1). Every trial they were shown one of the bundles from Task 2, and they had to decide (yes/no) if they wanted to take at least three bites from both of the foods at the end of the experiment. Choices were indicated using a keyboard button press using the subject’s dominant hand with the index and middle fingers. Subjects could take as long as desired to make each choice. The choice task consisted of 200 trials, selected at random from the set of all possible bundles described above. Subjects were instructed that at the end of the experiment they would need to remain in the lab for an additional twenty minutes. During this time, one of the two hundred trials was randomly selected and their choice in that trial was implemented. This procedure encouraged subjects to give incentive compatible responses.

I focus on this choice task because it is an extremely simple setting in which binary-attribute choice can be studied. Here, the choice objects are the bundles. Each bundle consisted of one appetitive food and one aversive food. Hence, the bundles contain two attributes: an appetitive and an aversive stimulus. The liking




Fig. 1. Experimental design. Subjects participated in a task where they made decisions over whether they were willing to take at least three bites from each of the two foods on the screen at the end of the experiment. One food was previously rated as appetitive while the other was previously rated as aversive. Participants' eye movements were recorded as they made these choices. The timing of each screen is depicted at the bottom of the figure. Each subject saw a fixation cross for 500 ms, then had as long as they liked to enter a rating or make a choice. They then saw feedback for 2000 ms and moved to the next trial. Eye tracking was only used in the choice task.

ratings provide a measure of the attribute values for each bundle. Although this set up differs from the prototypical multi-attribute choice setting, with several options each of which contain several attributes, it can serve as a first step to verify whether this model can be applied to a multi-attribute setting.

Importantly, in order to study the role of relative attention to the two attributes, I monitored fixations during the choice task with eye tracking. A desktop mounted SR Research Eyelink 1000 eye tracker recorded eye movements throughout the Choice Task at 500 Hz. The eye tracker was calibrated immediately after reading the instructions for the section. Those subjects who were unable to pass an initial 13-point calibration eye tracking exercise were excluded from participating in the choice task. In order to measure calibration drift over time, every fifty trials of the experiment subjects were asked to fixate to a center fixation point and press a button. The experiment would only continue once their fixation was within a certain region of the calibration point. All subjects passed these three additional calibration tests.

The SR Research Eyelink software determined the duration and location of all fixations in the choice task. I defined a region of interest (ROI) around each food image that consisted of the 300×300 image.

Although it is well known that fixations and attention can be dissociated (Egly, Driver, & Rafal, 1994; Posner, Nissen, & Ogden, 1977), for the purposes of this experiment fixations appear to provide a reasonable measure of attention at any instance during the choice process.

3. Results

3.1. Model

The experiment was designed to test the ability of the baDDM to account for the relationship between fixations, choices, and reaction times in a simple binary-attribute choice setting. To see why, I begin by describing the model and its properties.

The model assumes that the value of a bundle, denoted by V_B , is given by a linear combination of the values of the appetitive food (V_P) and the aversive foods (V_N); i.e.,

$$V_B = \beta_0 + \beta_P V_P + \beta_N V_N$$

Note that the rating tasks provide a measure of each of these values, which allows me to test the general validity of this assumption. To do so, for every subject I estimated a linear regression of the bundle ratings on the ratings of the appetitive and the aversive foods, and found that the data approximates the assumption rea-

sonably well (mean $\beta_0 = 0.59$, $SD = 2.17$; mean $\beta_P = 0.61$, $SD = 0.40$; mean $\beta_N = 0.99$, $SD = 0.74$; mean $R^2 = 0.32$, $SD = 0.18$). I further tested for an interaction effect by including an additional regressor, $V_P * V_N$, in the above linear combination. After estimating this regression for every subject, I found the mean coefficient on the interaction was 0.14 ($SD = 0.31$); however, the mean difference in R^2 before and after adding this term was only 0.007 (minimum = 0.00, maximum = 0.03).

As depicted in Fig. 2, the baDDM assumes that decisions are made by integrating a relative decision value (RDV) signal over time until enough evidence is accumulated in favor of one of the two options: choice = "yes" or choice = "No." In particular, the subjects choose "yes" if the barrier crossed is at $B = +1$, and choose "no" if the barrier crossed is at $B = -1$. The model also predicts reaction times, since choice time equals the time the barrier is crossed.

A key property of the model is that both the bundle properties and attention are allowed to influence the evolution of the RDV signal, and thus how choices are made. In particular, the model assumes that there is a fixation bias, so that attending to a particular attribute increases the weight that attribute is assigned in the integration process. Specifically, when looking at the appetitive attribute the RDV signal evolves according to

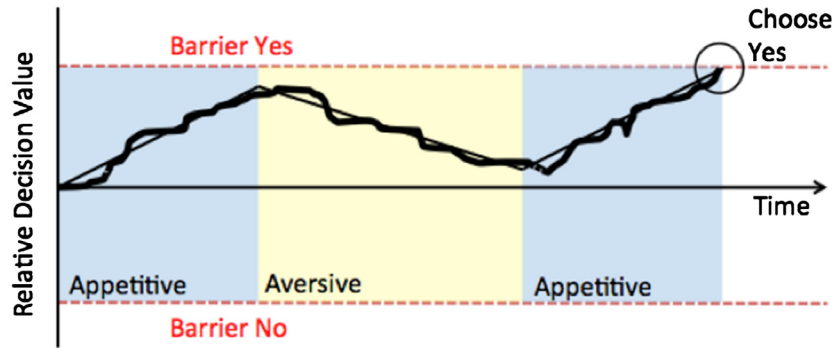
$$RDV_t = RDV_{t-1} + d(\beta_0 + \delta\beta_P V_P + \theta\beta_N V_N) + \varepsilon_t$$

and when looking at the aversive attribute, it evolves according to

$$RDV_t = RDV_{t-1} + d(\beta_0 + \theta\beta_P V_P + \delta\beta_N V_N) + \varepsilon_t.$$

Here, RDV_t indicates the value of the RDV signal at time t , d is a constant that controls the speed of integration (in units ms^{-1}), δ is a parameter that can take values greater than or equal to 1 and reflects the fixation bias towards the currently fixated attribute, θ is a constant between 0 and 1 that reflects a fixation bias to the currently non-fixated attribute, and ε_t is i.i.d. white Gaussian noise with variance σ^2 that reflects the stochastic nature of the process.

Importantly, the model assumes that the fixation process between the two attributes is independent of each individual attribute's value, or of the location of the positive and negative items. In particular, the first fixation is assumed to go to the left attribute with a constant probability p . Fixations then alternate between the two foods until a barrier is crossed. At the beginning of each fixation, a maximum fixation length is drawn from a distribution that depends on the type of attribute (appetitive or aversive), and whether the fixation is a first fixation or a later one. The fixation is then allowed to run its course unless a barrier is crossed before it terminates, which ends the choice process.



$$\text{Look at Appetitive Attribute: } RDV_t = RDV_{t-1} + d(\delta\beta_p R_p + \theta\beta_v R_v + \beta_0) + \epsilon_t$$

$$\text{Look at Aversive Attribute: } RDV_t = RDV_{t-1} + d(\theta\beta_p R_p + \delta\beta_v R_v + \beta_0) + \epsilon_t$$

Fig. 2. Depiction of the baDDM. A relative decision value (RDV) signal evolves over time. Its slope is biased towards the fixated item, but random noise is added to the RDV at every millisecond. When the RDV hits a barrier, a decision is made. The shaded vertical regions represent what item is currently fixated. In this example, three fixations are made (appetitive, aversive, appetitive) and the individual chose “yes.” The equations below the image describe how the RDV is integrated over time. The blue δ parameter describes an increase in weight that the attended item receives, while the red θ parameter describes a decrease in weight that the unattended item receives. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Several properties of the model are worth highlighting. First, the model includes as a special case a DDM without an attentional bias, which arises when $\delta = \theta = 1$. This model is almost identical to the “standard” DDM that has been widely used in previous literature to study binary choices in a large number of domains, including simple choices (Ratcliff, 1978; Milosavljevic, Koch, & Rangel, 2011; Ratcliff & Smith, 2004).¹

Second, the model exhibits a fixation bias when $\delta > 1$ or $\theta < 1$. In that case, an exogenous relative increase in attention to the appetitive food biases choices towards consuming the bundle, while the opposite is true for an exogenous decrease. Fig. 2 provides an intuition for why this is the case. Consider a case in which $\beta_p = \beta_N = 1$, $V_p = -V_N$, and $\beta_0 = 0$. Here, in the absence of an attentional bias (i.e., when $\delta = \theta = 1$), the slope of the RDV is always zero, and the choice is determined simply by the realization of noise. In contrast, when $\theta < 1 < \delta$ the slope of the RDV signal is positive when looking at the appetitive attribute, and negative otherwise. As a result, the probability of choosing “yes” depends on the relative allocation of attention.

Third, the model has four free parameters ($d, \delta, \theta, \sigma$) that can be fit using the choice, fixation, and reaction time data. The model has a fifth parameter, given by the height of the barrier, which is assumed to be fixed at ± 1 . Fixing the barrier separation comes without loss of generality since multiplying the barriers, slope (d), and noise by a fixed constant has no effect on the data that the model generates (Ratcliff & McKoon, 2008; Ratcliff et al., 2016).

Fourth, and somewhat more technical, the model allows for an asymmetric bias on the attended and unattended attributes (as opposed to requiring that $\delta - 1 = 1 - \theta$). This asymmetry can be identified from the data as long as β_0 is non-zero.²

¹ The “standard” DDM referenced does not include variability in drift rate or a starting point; however, it does include variability in non-decision time as later described in the model fitting section.

² Evidence suggests that the asymmetry here is well identified. First, given that the mean estimated β_0 across subjects is 0.59 with a standard deviation of 2.17, a t -test across subjects reveals marginally significant evidence that β_0 is non-zero ($t(45) = 1.85$; $p = 0.071$). Second, given that the estimates of β_0 arise from subject-specific regressions of the data from the bundle liking task on the data from the individual food liking task and that subjects faced different numbers of trials in these tasks depending on their subjective ratings, we can analyze the significance of β_0 across subjects in pooled data. To do this, we ran a linear mixed-effects regression with random slopes and intercepts for each subject where the dependent variable, bundle liking rating, was regressed on a constant, value of the appetitive food, and value of the aversive food using all trials from the bundle rating task. The estimate of the constant, β_0 , was 0.57 with a standard error of 0.22, resulting in a p -value of 0.013.

3.2. Model fitting

I fitted the model using MLE on the pooled group data. As described above, the baDDM has four free parameters: the constant determining the speed of integration d , the positive fixation bias δ , the negative fixation bias θ , and the noise parameter σ . I fitted these parameters at the group level by pooling the data from all subjects into a single data set. The parameters were fitted to maximize the maximum likelihood of the observed choices and reaction times. Importantly, the model was fitted using only even trials, and the odd trials were reserved for out-of-sample comparisons, as described later. I fitted the model at the group level to both choice and reaction time data for all 46 subjects by pooling all even numbered trials into a single data set. The model requires a large amount of data to estimate the parameters accurately, and fitting them at the individual level would result in highly noisy estimates. A similar variant and emphasis of this group estimation procedure has been used in all previous work that estimates parameters from attentional drift diffusion models (Krajbich & Rangel, 2011; Krajbich et al., 2010, 2012).

The MLE procedure was conducted as follows. First, I simulated the model 4000 times for each combination of the model parameters in the grid described below, and for each of the seven possible bundle ratings (ranging from -3 to 3). The simulations were carried out using 1 ms time steps.

In each simulation, individual liking ratings for both the appetitive and aversive foods were drawn from the empirical distribution of liking ratings conditional on the rating of the bundle. Furthermore, once a pair of liking ratings was drawn, I chose subject-estimated regression weights (β_0, β_p , and β_N) associated with the randomly selected simulated liking ratings. For instance, if the drawn liking rating for the appetitive item and the aversive item was drawn and belonged to subject i , then subject i 's regression weights were used throughout the simulated trial.

In each simulation, I randomly sampled fixations lengths from the empirical distribution for the group, conditional on whether it was a first or middle fixation in the trial, and whether the fixation was to either the positive or negative food. I also assumed that subjects looked first to the left item 68% of the time, which is the frequency observed in the data. As in the observed data, I assume fixations alternate between the foods. Although the model assumes that fixations between the two items on the screen occur

instantaneously, in practice there are observed saccade length transitions in each trial. To take this into account, in every simulated trial I randomly sampled from the empirical distribution of transition times, and add that sampled transition time to the simulated total fixation time. To clarify, in every trial I have a response time (the length of time from stimulus onset to response) and total time spent fixating to both items. I define the transition time as the difference between these two variables, and the sum of the transition time and total simulated fixation time represents the simulated response time in a trial.

Second, I used the simulations to compute the likelihood of each observation, for each vector of parameters, as follows. Reaction time was discretized into bins of 100 ms, from 0 to 7400 ms, with an additional bin representing a trial that took longer than 7400 ms. Choice data is automatically discretized into yes/no bins. I then used the simulation results, conditional on the bundle rating, to compute the frequency with which responses followed into each time-choice bin.

Third, I used the data from the previous step to compute the log-likelihood of the data for each vector of parameters, and carried out a grid search to identify the vector of parameters with the largest maximum likelihood. I then performed the Nelder-Mead optimization algorithm on the grid search algorithm's identified maximum.

To reduce computational costs, this maximization was done in two steps. In step one I first did a coarse search over the following parameter space:

d in {0.0001, 0.0005, 0.001, 0.0015, 0.002, 0.0025}
 σ in {0.005, 0.01, 0.02, 0.025}
 δ in {0.85, 0.90, 0.95, 1, 1.05, 1.1, 1.15}
 θ in {0.85, 0.90, 0.95, 1, 1.05, 1.1, 1.15}

which identified ($d = 0.0015$, $\sigma = 0.02$, $\delta = 0.95$, $\theta = 0.90$) as the parameters that maximized the log-likelihood. In step two I did a finer search around this vector using the grid:

d in {0.001, 0.0011, 0.0012, 0.0013, 0.0014, 0.0015, 0.0016, 0.0017}
 σ in {0.015, 0.0175, 0.02, 0.0225}
 δ in {0.85, 0.875, 0.9, 0.925, 0.95, 0.975, 1, 1.025, 1.05, 1.075, 1.1, 1.125, 1.15}
 θ in {0.85, 0.875, 0.9, 0.925, 0.95, 0.975, 1, 1.025, 1.05, 1.075, 1.1, 1.125, 1.15}

The best fitting parameters where $d = 0.0013$, $\sigma = 0.02$, $\delta = 1.05$, and $\theta = 0.90$ (log-likelihood = $-18,016$; AIC = 36,040). Finally, I ran a Nelder-Mead optimization algorithm with starting vector ($d = 0.0013$, $\sigma = 0.02$, $\delta = 1.05$, $\theta = 0.90$). The algorithm identified ($d = 0.0013$, $\sigma = 0.0200$, $\delta = 1.0500$, $\theta = 0.9000$) as the maximum after rounding to four decimal places.

In order to test for the presence of a fixation bias, I also fitted a model with the restriction $\delta = \theta = 1$. The best restricted model also had $d = 0.0013$ and $\sigma = 0.02$ (log-likelihood = $-18,052$; AIC = 36,108). A likelihood ratio test statistic from a chi-square test of nested models provided support in favor of the unrestricted model with a small but significant fixation bias ($p < 0.001$; $df = 2$), and the AIC of the unrestricted model is lower than that of the restricted (unrestricted AIC = 36,040; restricted AIC = 36,108). Furthermore, the restricted model is less than 0.001 times as probable as the unrestricted model to minimize the information loss, providing essentially no evidence in favor of the restricted model.

In order to test for the asymmetry of the fixation bias, I also fitted a model with the restriction $\delta - 1 = 1 - \theta$. The best restricted model had parameters $d = 0.0012$, $\sigma = 0.0225$, $\delta = 1.025$, and $\theta = 0.975$ (log-likelihood value = $-18,039$; AIC = 36,084). A likelihood ratio test statistic from a chi-square test of nested models provided support in favor of the model with the asymmetric fixation bias ($p < 0.001$; $df = 1$) and the AIC of the unrestricted model is lower than that of the restricted (unrestricted AIC = 36,040; restricted AIC = 36,084). Furthermore, the restricted model is less than 0.001 times as probable as the unrestricted model to minimize the information loss, providing essentially no evidence in favor of the restricted model.

Together, these results are consistent with the existence of small and asymmetric fixation bias. It is worth emphasizing, however, that the size of these fixation biases are significantly smaller than those that have been found in previous studies of simple choice (Krajbich & Rangel, 2011; Krajbich et al., 2010, 2012).

The fit of the above three models is further explored in Fig. 3 which depicts differences in the log-likelihood for the unrestricted model compared to the two restricted versions. A positive difference at a bundle value in the figure indicates the unrestricted asymmetric bias model fits better than a restricted version. While the symmetric bias model fits better at the bundle value of -3 and the no bias model fits better at the bundle value of -1 , the asymmetric fixation bias model fits the data better at all other bundle values.

3.3. Basic psychometrics

Fig. 4 compares the basic psychometric properties of the data with the predictions generated by the best fitting model. In this figure, and the following ones, black denotes data and red denotes out of sample predictions. Both data and predictions are shown only for odd trials, to insure that the comparison is out-of-sample.

Predictions were made by simulating the best-fitting model 4000 times for each bundle liking rating, and sampling fixation lengths from the empirical distribution of observed fixations, conditioning only on whether a fixation was to an appetitive or aversive attribute, and whether the fixation was a first a later one. Many of the tests here are strongly inspired by previous tests of the aDDM (Krajbich & Rangel, 2011; Krajbich et al., 2010, 2012).

Fig. 4A depicts the psychometric choice curve. It shows that the probability of choosing yes is a logistic function of the bundle value which matches well the predictions of the best fitting model.

Fig. 4B depicts the reaction time curve, which exhibits the typical inverted-U pattern of reaction time when plotted against the liking rating of the bundle, so that more difficult choices take longer. The data also matches the predictions of the best fitting model.

Finally, Fig. 4C depicts the fixation curve, which shows that the number of fixations that it takes to make a choice increases with the difficulty of the choice. Although both the data and predictions exhibit the same general pattern, the model over predicts the impact of choice difficulty on the number of fixations, as well as the average number of fixations (data: coefficient on difficulty = -0.24 , mean = 2.79; model: coefficient on difficulty = -0.42 , mean = 3.14). Part of the mismatch between actual and predicted fixations has to do with technical limitations of the fitting and prediction procedure, namely that fixation were sampled from non-final fixations meaning that more fixations in the simulated data than in the odd numbered data will be required to achieve the same reaction times.

3.4. Properties of the fixation process

As described above, the basic baDDM assumes that fixations are independent of the value of the foods. Here I test if the pattern of observed fixations is consistent with this assumption.

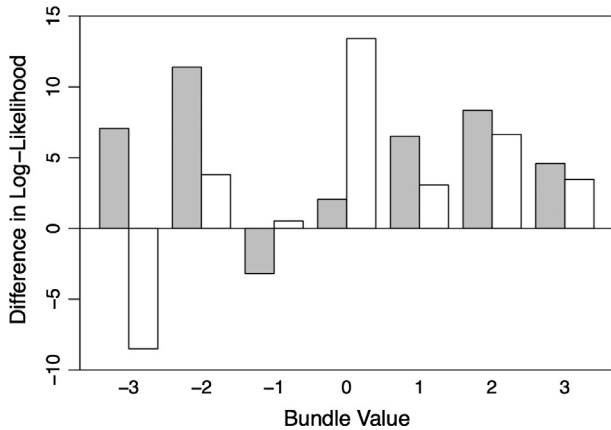


Fig. 3. Model fits. For each bundle value on the horizontal axis, gray bars represent the difference in log-likelihood between the best fitting unrestricted model and the best fitting restricted model with no fixation bias. White bars represent the difference in log-likelihood between the best fitting unrestricted model and the best fitting restricted model with a symmetric fixation bias. Positive differences indicate the unrestricted model fits better at a given bundle value.

As shown in Table 2, subjects exhibited a left-first bias: they looked at the left attribute before the right 64% of the time ($t(45) = 4.43, p < 0.001$). However, the location of first fixation was not significantly different for appetitive and aversive foods (Table 2, $t(45) = 1.80, p = 0.08$). In the pooled even and odd-numbered data, subjects spent on average 37.9% (SD = 8.0%) of each trial not fixating to an ROI. This time includes latency, time of saccades, as well as fixations outside the ROI. A mixed-effects regression with combined even and odd-numbered data of all time spent not fixating to an ROI on the trial number and constant revealed a small but significant increase in non-ROI fixation time over time (slope = 0.0006, $p < 0.001$).

As shown in Fig. 5, fixations to aversive foods were about 58 ms longer on average than fixations to appetitive items, both for first, middle, and last fixations (first fixation: $t(45) = 4.68, p < 0.001$; middle fixation: $t(45) = 3.20, p = 0.003$; last fixation: $t(45) = 3.89, p < 0.001$). This is consistent with the assumptions of the model listed above, since fixations to different attribute types might follow a different process (e.g., they might have a different processing latency). The key assumption of the model, however, is that fixation duration is not dependent on the value of the fixated or unfixated attributes (controlling for their attribute type; i.e., appetitive or aversive). I tested this assumption by examining how the duration of different types of fixations, either first or middle fixations, was affected by the value of the attended and unattended attribute. I ignore final fixations in this analysis since their duration is endogenous to the choice process.

Table 2

Depicts the mean percent of first fixations to each item by spatial features (Panel A) and attribute features (Panel B). Standard deviation given below in parentheses.

Percent of First Fixations to Each Item		
A. Spatial		
Percentage	Left 64.4 (22.1)	Right 35.6 (22.1)
B. Attribute of Interest		
Percentage	Appetitive 48.2 (6.6)	Aversive 51.8 (6.6)

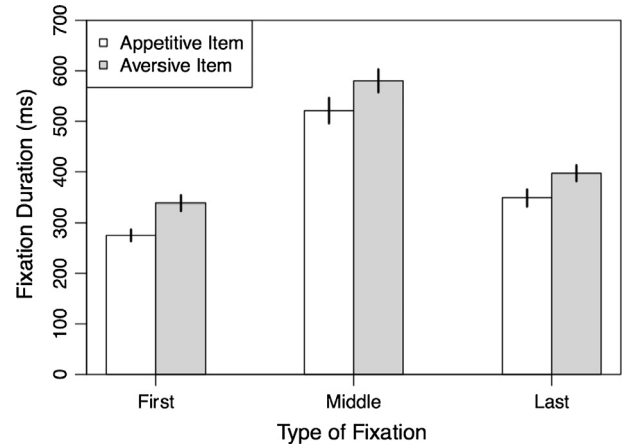


Fig. 5. Fixation Durations. The mean fixation duration for first, middle, and last fixations to appetitive and aversive items. Standard errors are clustered by subject.

The duration of the first fixation was not significantly related to the value of the attended item (mixed-effects regression of first fixation length on an indicator for whether the item is appetitive, the weighted value of the item and the weighted value of the unattended item: beta of indicator = -37.92 , t -statistic: $-1.84, p = 0.07$; beta for attended value = -1.90 , t -statistic = $-0.29, p = 0.77$; beta for unattended value = -5.38 , t -statistic = $1.26, p = 0.21$). Furthermore, the duration of the first fixation was not related to the value of the bundle (mixed-effects regression of first fixation length on an indicator for whether the item is appetitive and the value of the bundle: beta of indicator = -64.44 , t -statistic = $-5.31, p < 0.001$; beta for value of bundle = -0.35 , t -statistic = $-0.11, p = 0.91$) or the difficulty of the choice (mixed-effects regression of first fixation length on an indicator for whether the item is appetitive and the absolute value of the

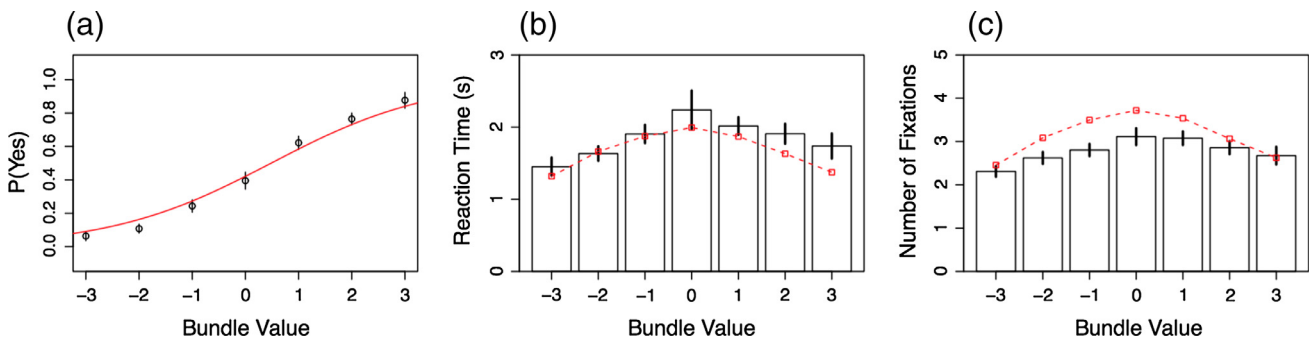


Fig. 4. Basic psychometrics. (a) Psychometric choice curve as a function of the bundle liking rating. (b) Reaction times as a function of the bundle liking rating. (c) The number of fixations in a trial as a function of the bundle liking rating. The red lines indicate the model's predictions. Subject data is shown for the odd-numbered trials, with standard errors clustered by subject. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

bundle: beta of indicator = -64.44 , t -statistic = -5.26 , $p < 0.001$; beta for absolute value of bundle = -1.49 , t -statistic = -0.32 , $p = 0.75$). Clearly, the duration of the first fixation was not dependent on the strength of value within the appetitive or aversive categories.

For middle fixations, I found no significant relationship with the attended value and a significant but quantitatively small effect of the unattended value (analogous mixed-effects regression: beta of indicator = -24.85 , t -statistic = -0.80 , $p = 0.43$; beta for attended value = 6.51 , t -statistic = 0.92 , $p = 0.36$; beta for unattended value = 17.13 , t -statistic = 2.32 , $p = 0.02$). Importantly, the size of this effect is relatively small as a change in value of the unattended attribute of 2.5, the maximum possible change, would only alter a middle fixation duration by 43 ms, on average. Furthermore, the duration of the middle fixation was not related to the value of the bundle (analogous mixed-effects regression: beta for indicator = -62.04 , t -statistic = -3.82 , $p < 0.001$; beta for value of bundle = -5.55 , t -statistic = -0.91 , $p = 0.37$), but was slightly related to the difficulty of the choice (analogous mixed-effects regression: beta for indicator = -61.50 , t -statistic = -3.78 , $p < 0.001$; beta for absolute value of bundle = -21.61 , t -statistic = -2.25 , $p = 0.03$). Again, even though I found a significant effect here, the effect was quite small in size as a move from the most difficult choice, with bundle value 0, to the simplest, with an absolute bundle value of 3, only corresponds to a fixation duration change of 65 ms.

Together, the results in this section suggest that the properties of the observed fixation process are largely consistent with the assumption that fixations are independent of the value of specific attributes (e.g., example, changing a mildly appetitive item for a highly appetitive one). In addition, there was considerable variation in the duration of fixations across trials, which makes it possible for the existence of random fluctuations in attention to influence choices.

3.5. Model predictions

The baDDM makes additional predictions about the pattern of the fixations, and their relationship to choices, which I test here.

First, the model predicts that final fixation durations should be shorter than middle fixations, since final fixations are terminated prematurely when a barrier is crossed. As shown in Fig. 5, this also holds in the data (mean last = 376 ms; mean middle = 550 ms; $t(45) = 11.42$, $p < 0.001$). Relatedly, the model also predicts that final fixations are shorter than penultimate fixations (mean last = 376 ms; mean penultimate = 500; $t(45) = 9.00$, $p < 0.001$). Interestingly, I also found that first fixations were shorter than middle fixations (mean first = 309 ms; mean middle = 550 ms; $t(45) = 15.20$, $p < 0.01$), consistent with previous work finding decision makers parse their decisions into several attention related tasks (Glaholt & Reingold, 2011) and that shorter fixations occur early on in a choice (Glöckner & Herbold, 2011; Krajbich & Rangel, 2011; Krajbich et al., 2010, 2012). Note that although the model made no ex-ante prediction about the relationship between these two types of fixations, this pattern was incorporated in the prediction exercise since fixation durations were conditioned on whether a fixation was first or later.

Finally, the model predicts a strong relationship between the fixation-averaged value at the start of the final fixation and the duration of the final fixation, conditional on the choice made. Specifically, the model predicts that conditional on a “no” choice, the duration of the final fixation should increase with the variable

$$F_P(\beta_0 + \delta\beta_P V_P + \theta\beta_N V_N) + (1 - F_P)(\beta_0 + \theta\beta_P V_P + \delta\beta_N V_N),$$

where F_P denotes the fraction of the trial spent attending to the appetitive item (as of the beginning of the last fixation). Essentially,

this variable measures the average slope of the RDV signal during the initial phase of the choice, given the realization of fixations up to that point. The intuition for this relationship illustrates the key forces at work in the baDDM, and are best seen using a hypothetical case in which $\beta_P = \beta_N = 1$, $V_P = -V_N$, and $\beta_0 = 0$. In this case, when $\theta < 1 < \delta$ the slope of the RDV is positive when fixating to the appetitive attribute, and negative otherwise. As a result, the larger F_P , the farther the RDV signal is likely to be from the “choose no” barrier at the beginning of the last fixation. Thus, the process needs to cover more distance during the last fixation to reach the “no” barrier, leading to a longer last fixation.

To test this prediction, I estimated a mixed-effects regression of final fixation duration on the final fixation average value variable, for trials in which the subjects choose “No.” Consistent with the prediction, I found a significant effect between the two (slope = 33.87 , t -statistic = 4.22 , $p < 0.001$). The model makes an analogous prediction for trials in which the subject chooses yes, albeit with the opposite sign, which was also present on the data (slope = -24.17 , t -statistic = -2.43 , $p = 0.02$).

These results demonstrate that several key predictions regarding the pattern of fixations and their relationship to choices hold in the data.

3.6. Choice biases

When $\theta < 1 < \delta$, the baDDM predicts a number of attentional driven biases that the $\theta = \delta = 1$ model does not predict. This provides an additional set of model tests, which I carried out below.

First, the model predicts that, controlling for bundle values, the probability of choosing “yes” increases with the relative attention to the appetitive attribute interacted with its subjective value, and decreases with additional time spent attending to the aversive attribute interacted with its subjective value with no main effects. To test for this effect in the data, I ran a mixed-effects logistic regression of choice on bundle rating, the weighted value of the appetitive food ($=\beta_P V_P$) interacted with its relative fixation time, and the weighted value of the aversive food ($=\beta_N V_N$) interacted with its relative fixation time. Consistent with the predictions, I found a negligible bias (constant = -0.20 , $p = 0.43$), a significant increase in the probability of choosing “yes” with bundle value (slope = 1.14 , $p < 0.001$) and the value of the appetitive food weighted by its share of relative attention (slope = 0.37 , $p < 0.04$), and a significant decrease in the probability of saying yes with the value of the aversive item weighted by its relative attention (slope = -0.56 , $p = 0.001$). Similar effects were found in the simulated data (constant = 0.07 , $p = 0.004$; slope bundle rating = 0.63 , $p < 0.001$, slope weighted value appetitive item interacted with relative fixation time = 0.17 , $p < 0.001$; slope weighted value aversive item interacted with relative fixation time = -0.50 , $p < 0.001$).

It is worth emphasizing that, despite the small effect of the fixation bias coefficients (θ and δ), the resulting choice biases need not be small. To see why, consider an example in which the bundle liking rating is 0, the appetitive attribute has a rating of 2, the aversive attribute has a rating of -2 , and the value weights take the mean value over all subjects. The estimate predicts that when an individual spends 10% of the trial attending to the appetitive attribute, there is only a 23.9% chance of agreeing to consume the bundle; however, if that individual instead spends 90% of the trial attending to the appetitive attribute, the probability of responding “yes” increases to 52.4%. In contrast, in a model without a fixation bias ($\theta = \delta = 1$), this change in the fixation pattern has no effect on the choices.

Second, the model predicts that, controlling for bundle value, the probability of choosing “yes” depends on the relative amount of time spent attending to the appetitive item. As shown in Fig. 6A, this was true in both the data and the model predictions,

and the size of the two effects was remarkably similar. This test was conducted as follows. For every trial, I computed a corrected choice measure by subtracting the observed choice (yes = 1, no = 0) from the average frequency with which the bundle was chosen for all trials with that bundle rating. I then estimated a linear regression of the corrected choice probabilities on the relative time advantage to the appetitive item and found the predicted effect in both the simulated data (slope = 0.06; $p < 0.001$) and the data (mixed-effects regression slope = 0.04; t -statistic = 2.26, $p = 0.03$).

Third, the model predicts that the longer the first fixation is to the appetitive item, interacted with its weighted value, the more likely the subject is to choose “yes.” To see why, note that longer first fixations move the RDV signal towards the “yes” barrier, which all else equal biases choices towards “yes.” To test this, I estimated a mixed-effects logistic regression of choice on the duration of the first fixation, conditioning on a first fixation to the positive attribute, and found the predicted effect in both the data and the simulations (simulated model: slope = 0.0008, $p < 0.001$; data: slope = 0.0017, $p < 0.001$). I also found the analogous effect for first fixations to the aversive item (simulated model: slope = -0.0008 , $p < 0.001$; data: slope = -0.0017 , $p < 0.001$).

In a related result, any initial biases in the first attended attribute should translate into choice biases. Fig. 6B shows that this is the case: a linear regression of the probability that a subject looks first at the appetitive food first on the subjects' average probability of choosing shows a significant positive relationship (slope = 0.09; $p = 0.04$).

Fourth, the model predicts a relationship between the identity of the last fixation and choice. In particular, conditional on value of the bundle, it predicts that the probability of choosing “yes” is larger when the last fixation is to the appetitive item. The intuition for this prediction stems from the fact that the slope of the RDV signal is more likely to be positive, and thus climbing towards the “yes” barrier, during fixations to the appetitive item.

I estimated a logistic regression of choice on a constant, the bundle liking rating, an indicator variable for when the last fixation is to the appetitive item, and the interaction of the bundle liking rating with the indicator variable, in the simulations and the actual data. In the simulations, I found a significant effect of the identity of the final fixation (beta = 0.69, $p < 0.001$), but not of the interaction term (beta = -0.00 , $p = 0.89$). A similar pattern was found in the data (last fixation bias to positive indicator: beta = 0.21; $p = 0.06$; interaction term: beta = 0.03; $p = 0.60$). As shown in Fig. 6C, although these biases are small, they follow a quantitatively similar pattern in the model predictions and data.

4. Discussion

The results described here suggest the baDDM provides a quantitatively accurate description of the choice process in a simple binary-attribute accept-reject environment. Specifically, the model quantitatively describes the relationship between choices, response time, and the correlation of these variables with attentional deployment as measured by fixations. The data suggest that individuals increase the weight of an attended attribute by 5% and decrease the weight of an unattended attribute by 10%. Consistent with this estimation, a number of attention-based choice biases found in the data support the model's predictions.

Although the estimated model parameters suggest a fixation bias alters the decision process, it has a relatively small effect size. Notably, previous findings have estimated that only 30% of an unattended option's value is accounted for during the choice process (Krajbich et al., 2010). There are several possibilities that can explain this finding. First, as the baDDM is one of the first

applications to model the fixation bias in attributes rather than options, it is possible the fixation bias over attributes is simply smaller than over options. Second, the bundle task forces subjects to accept or reject an outcome. If they were instead making choices over two or more bundles, consistent with a more prototypical multi-attribute setting, a larger bias might be present. Understanding how the size of the bias changes with the task is an important step for future work.

One question about the model concerns the direction of causality between fixations and choice. Namely, while the model assumes that fixations bias the value estimation process, another possibility is that the value of the attributes directly affects the fixation process. Although the data here suggest such an explanation is not responsible for driving the observed results, the best way to address this question is through follow-up work that provides a causal test of this theory. Several related papers address this issue in the contexts of simple food choice and risky decision-making (Armel, Beaumel, & Rangel, 2008; Kim, Seligman, & Kable, 2012). Furthermore, additional work finds that the vmPFC encodes attention-modulated relative value signals, suggesting neurobiological evidence that fixations alter the choice process (Lim, O'Doherty, & Rangel, 2011). While this literature speculates that there is a causal role from fixations to choice, one cannot rule out the possibility that causality also works in the other direction.

The paper tests a number of main properties of the baDDM, which largely fall into four main categories. First, the model with the best fitting asymmetric fixation bias appears to make accurate out of sample quantitative predictions regarding the relationship between the value of the bundle and choices, response time, and number of fixations between the two attributes. Specifically, choices increase as the value of the bundle increases and fixations and response time decrease as the absolute value of the bundle increases. Second, although the model makes fairly strong assumptions by assuming independence between the value of the attributes and fixation durations, the data largely support this assumption in the simple binary-attribute task analyzed here. Third, there is evidence that supports several predictions about patterns and durations of fixations and their relationship to choices. Notably, the paper replicates previous work that has found middle fixation durations are longer than final fixation durations. Fourth, given that the best fitting model's estimated parameters included fixation biases that increased the weight of the attended attribute and decreased the weight of the unattended attribute, a number of attentional driven choice biases were predicted. These predictions were largely supported in the data and, critically, a model without fixation biases would be unable to predict these relationships which provides for a strong test of the model.

Naturally, one criterion when evaluating the baDDM concerns its comparison with existing evidence accumulation models in the literature. Rather than explicitly modeling and predicting the distribution of attention, as some previous work as done, the baDDM uses the observed distributions of attention as an input and therefore must make strong assumptions concerning independence. Given this, the model is able to make a sizable number of predictions regarding the relationship between visual attention and decision-making. Although certain competing models may share a subset of these predictions, others remain silent on particular correlations between attention and choice.

One example of this can be seen in analyzing how the order in which attributes are attended can influence decisions. The baDDM makes a number of predictions between these variables, including that choices are biased towards the first and last fixated attributes. While some previous multi-stage models make predictions indicating the order in which attributes are processed matters (Diederich & Oswald, 2014, 2016), a significant portion of the sequential sampling literature fails to make such predictions.

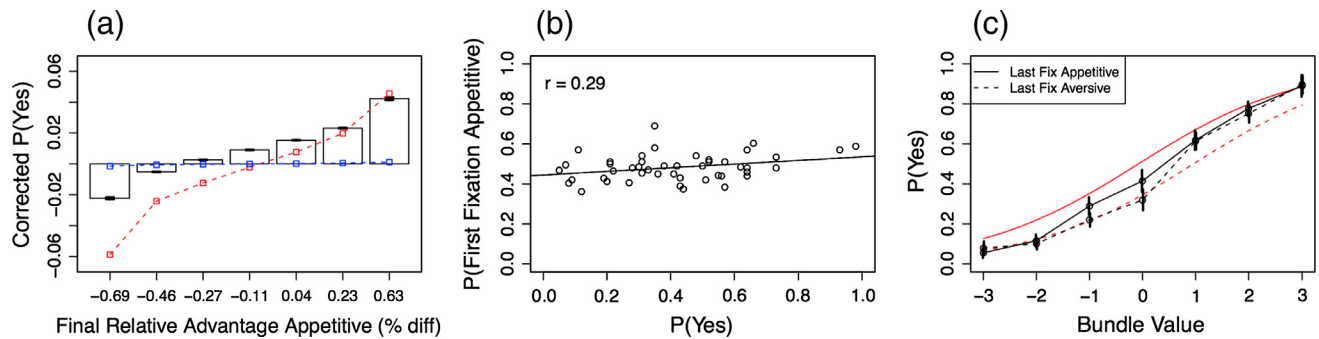


Fig. 6. Choice Biases. (a) Corrected probability of agreeing to eat the bundle as a function of the relative time advantage looking at the appetitive item. Bins depict the odd-numbered trials, the red dotted line is the model simulation, and the blue dotted line is the no bias model simulation. To compute the bins, the data was split into seven equal bins and the median of each is reported on the horizontal axis. (b) Probability of looking at the appetitive item as a function of agreeing to consume the bundle. Each circle represents a different subject. (c) The probability of agreeing to consume the bundle as a function of the bundle value and whether the last fixation was to the appetitive or aversive attribute. The red solid line indicates the model's prediction when the last fixation was to the appetitive item and the red dashed line indicates the model's prediction when the last fixation was to the aversive item. Black dots with a solid connecting line indicates the odd-numbered data when last fixation was to the appetitive item and black dots with a dashed connecting line indicates the data when the last fixation was to the aversive item. Standard errors clustered by subject. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Furthermore, the baDDM is mathematically well specified enough to test how well its predictions hold in empirical data, rather than solely numerical simulations of decisions. However, as the model here was applied to a specific experimental task it remains unclear how well it would perform in settings that more closely match the paradigms described in other models, such as a choice set that consists of at least two options with at least two attributes each. One example can be seen in a comparison to [Diederich and Oswald \(2014\)](#) who predict that if two attributes both favor some alternative, and the first attended attribute offers more evidence in favor of choosing an alternative than the second attribute, then regardless of how attention is deployed one should observe shorter response times for the more frequently chosen alternative. Alternatively, if the processing order was reversed then their model predicts a faster response time for the less frequently chosen option. Unfortunately, testing this and related hypothesis in the data collected here is not possible as by definition of the task here one attribute favors choosing “yes” while the other favors choosing “No.” Future work that models a more traditional multi-attribute choice setting and records fixation data will be better set up to directly address these and other related model comparisons.

Indeed, many real-world multi-attribute decisions are more complex than the binary-attribute bundle environment explored here. Much previous work in multi-attribute choice has explored two option choices with at least two attributes each and has found that attention typically increases with attribute importance ([Fiedler & Glöckner, 2012](#); [Glöckner, Fiedler, Hochman, Ayal, & Hilbig, 2012](#); [Glöckner & Herbold, 2011](#); [Hristova & Grinberg, 2008](#)). The fixation process detailed here is consistent with this underlying finding in that aversive attributes are more highly weighted than appetitive ones in choices and fixation durations to aversive attributes are longer than those to appetitive ones. The assumption made in this paper is that within an attribute, the value or level of the attribute does not influence attention. Overall, the data from the bundle task is largely consistent with this assumption; however, some previous work has found a u-shaped relationship between attribute levels and fixations, where high and low importance attribute levels receive more attention ([Meißner, Musalem, & Huber, 2016](#); [Mueller Loose & Orquin, 2012](#); [Orquin & Mueller Loose, 2013](#); [Sütterlin, Brunner, & Opwis, 2008](#)). Importantly, this u-shaped relationship has not been found in certain other tasks, such as risky gambles ([Fiedler & Glöckner, 2012](#); [Glöckner & Herbold, 2011](#)). However, unlike the current paper this previous work did not involve aversive stimuli, as all

risky choices were located in the gain and not loss domains. It remains an open question as to which tasks generate relationships between attribute levels and fixations although the data here suggest that an accept-reject task with an appetitive and aversive attributes may not generate a u-shaped relationship between attribute levels and fixation durations. However, given the rating elicitation used here, it might be particularly difficult to find a relationship in the collected data. In the task here, individuals rated the attributes over a relatively small range of values and it is possible that a u-shaped relationship could have been observed if subjects were allowed to be more discriminatory in their ratings.

The simple experimental design in this paper was chosen for two critical reasons. First, I wanted to understand how the baDDM could be applied to the simplest possible scenario with at least two attributes, an important feature when understanding whether and how these types of models may break down as the task becomes increasingly complex. Future work should expand this model to test the relative weight between attribute and option fixation biases. Second, as the number of attributes and options on a screen grows, the model must be able to account for the fixation process between all visible options and attributes. This is a complex task that may be a rich area for future work, but we currently lack a systematic understanding of how to model more than two fixations. For this reason, having a constant reference option of nothing off-screen was convenient. Indeed, a number of commonly viewed multi-attribute choice environments, such as choice under risk, are often elicited using similar techniques that utilize an off-screen reference option and ask subjects whether they are willing to accept a mixed-valence screen option.

Finally, it is worth noting that I find much evidence that has a qualitative flavor of loss aversion. Specifically, the duration of a fixation to an aversive attribute is on average longer than the duration to an appetitive one, individuals weight aversive attributes more heavily than appetitive attributes in choices, and aversive attributes are attended to for a longer period of time throughout the choice process. This differential attribute weighting is consistent with the literature on loss aversion ([Kahneman & Tversky, 1984](#)) while the differences in fixations to attributes are consistent with previous process tracking studies of loss aversion ([Willemsen et al., 2011](#)). This difference in fixation duration to the appetitive and aversive attributes can be further explained given that the amount of time one attends to a feature appears to influence its weight ([Fiske, 1980](#); [Schkade & Johnson, 1989](#); [Wedell & Senter, 1997](#); [Willemsen et al., 2011](#)).

Acknowledgments

I thank Eric Johnson and Jay Russo. I am especially grateful to Antonio Rangel for his comments, encouragement, and support of the project. Jennifer Cheng and Emily Tubb provided excellent research assistance. I thank the Foundation for Food and Agriculture Research and the USDA National Institute of Food and Agriculture, Hatch project 1010381 for financial support.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cognition.2017.06.007>.

References

- Armel, C., Beuvel, A., & Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, 3(5), 396–403.
- Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, 120(3), 522–543.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12(12), 4745–4765.
- Busemeyer, J. R., & Diederich, A. (2002). Survey of decision field theory. *Mathematical Social Science*, 43(3), 345–370.
- Busemeyer, J. R., & Johnson, J. G. (2004). In D. J. Koehler & N. Narvey (Eds.), *Handbook of judgment and decision making* (pp. 133–154). New York: Blackwell Publishing Co.
- Busemeyer, J. R., & Townsend, J. T. (1992). Fundamental derivations from decision field theory. *Mathematical Social Sciences*, 23(3), 255–274.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432–459.
- Camerer, C., & Johnson, E. J. (2004). Thinking about attention in games: Backward and forward induction. In J. Carillo & I. Brocas (Eds.), *The psychology of economic decisions* (pp. 111–129). New York: Oxford University Press.
- Day, R. F. (2010). Examining the validity of the Needleman-Wunsch algorithm in identifying decision strategy with eye-movement data. *Decision Support Systems*, 49(4), 396–403.
- Day, R. F., Lin, C. H., Huang, W. H., & Chuang, S. H. (2009). Effects of music tempo and task difficulty on multi-attribute decision making: An eye-tracking approach. *Computers in Human Behavior*, 25(1), 130–143.
- Diederich, A. (1995). Intersensory facilitation of reaction time: Evaluation of counter and diffusion coactivation models. *Journal of Mathematical Psychology*, 39(2), 197–215.
- Diederich, A. (1997). Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology*, 41(3), 260–274.
- Diederich, A. (2015). A multistage attention-switching model account for payoff effects on perceptual decision tasks with manipulated processing order. *Decision*, 2(4).
- Diederich, A., & Oswald, P. (2016). Multi-stage sequential sampling models with finite or infinite time horizon and variable boundaries. *Journal of Mathematical Psychology*, 74(2016), 128–145.
- Diederich, A., & Oswald, P. (2014). Sequential sampling model for multiattribute choice alternatives with random attention time and processing order. *Frontiers in Human Neuroscience*, 8(697), 1–13.
- Egley, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, 123(2), 161–177.
- Fehr, E., & Rangel, A. (2011). Neuroeconomic foundations of economic choice: Recent advances. *Journal of Economic Perspectives*, 25(4).
- Fiedler, S., & Glöckner, A. (2012). The dynamics of decision making in risky choice: An eye-tracking analysis. *Frontiers in Psychology*, 3, 335.
- Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology*, 38(6), 889–906.
- Glaholt, M. G., & Reingold, E. M. (2011). Eye movement monitoring as a process tracing methodology in decision making research. *Journal of Neuroscience, Psychology, & Economics*, 4(2), 125–146.
- Glimcher, P. W., & Fehr, E. (2014). *Neuroeconomics: Decision making and the brain*. London: Elsevier.
- Glöckner, A., Fiedler, S., Hochman, G., Ayal, S., & Hilbig, B. (2012). Processing differences between descriptions and experience: A comparative analysis using eye-tracking and physiological measures. *Frontiers in Psychology*, 2, 1–15.
- Glöckner, A., & Herbold, A.-K. (2011). An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes. *Journal of Behavioral Decision Making*, 24(1), 71–98.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574.
- Hare, T. A., Schultz, W., Camerer, C. F., O'Doherty, J. P., & Rangel, A. (2011). Transformation of stimulus value signals into motor commands during simple choice. *PNAS*, 108(44), 18120–18125.
- Heekeren, H. R., Marrett, S., Bandettini, P. A., & Ungerleider, L. G. (2008). A general mechanism for perceptual decision-making in the human brain. *Nature*, 431, 859–862.
- Holmes, W. R., Trueblood, J. S., & Heathcote, A. (2016). A new framework for modeling decisions about changing information: The piecewise linear ballistic accumulator model. *Cognitive Psychology*, 85, 1–29.
- Horstmann, N., Ahlgrim, A., & Glöckner, A. (2009). How distinct are intuition and deliberation? An eye-tracking analysis of instruction-induced decision modes. *Judgment and Decision Making*, 4, 335–354.
- Hristova, E., & Grinberg, M. (2008). Disjunction effect in prisoner's dilemma: Evidences from an eye-tracking study. In *Proceedings of the 30th annual conference of the cognitive science society*.
- Johnson, E. J., Schulte-Mecklenbeck, M., & Willemsen, M. C. (2008). Process models deserve process data: Comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review*, 115(1), 263–273.
- Kahneman, D., & Tversky, A. (1984). Choices, values and frames. *American Psychologist*, 39, 341–350.
- Kim, B. E., Seligman, D., & Kable, J. W. (2012). Preference reversals in decision making under risk are accompanied by changes in attention to different attributes. *Frontiers in Neuroscience*, 6, 109.
- Knoepfle, D. T., Yao-yi Wang, J., & Camerer, C. F. (2009). Studying learning in games using eye-tracking. *Journal of the European Economic Association*, 7(2/3), 388–398.
- Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and comparison of value in simple choice. *Nature Neuroscience*, 13, 1292–1298.
- Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift diffusion model extends to simple purchasing decisions. *Frontiers in Cognitive Science*, 3, 193.
- Krajbich, I., & Rangel, A. (2011). A multi-alternative drift diffusion model predicts the relationship between visual fixations and choice in value-based decision. *Proceedings of the National Academy of Sciences*, 108, 13853–13857.
- Lim, S. L., O'Doherty, J., & Rangel, A. (2011). The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. *Journal of Neuroscience*, 31, 13214–13223.
- Mas-Colell, A., Whinston, M., & Green, J. (1995). *Microeconomic theory*. Cambridge: Cambridge University Press.
- Meißner, M., Musalem, A., & Huber, J. (2016). Eye tracking reveals processes that enable conjoint choices to become increasingly efficient with practice. *Journal of Marketing Research*, 53(1), 1–17.
- Milosavljevic, M., Koch, C., & Rangel, A. (2011). Consumers can make decisions in as little as a third of a second. *Judgment and Decision Making*, 6(6), 520–530.
- Mueller Loose, S., & Orquin, J. (2012). How stimuli presentation format affects visual attention and choice outcomes in choice experiments. In *Australian and New Zealand Marketing Academy Conference*.
- Orquin, J. L., & Mueller Loose, S. (2013). Attention and choice: A review on eye movements in decision making. *Acta Psychologica*, 144(1), 190–206.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology*, 43, 87–131.
- Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness-to-pay in simple economic transactions. *Journal of Neuroscience*, 27(37), 9984–9988.
- Plassmann, H., O'Doherty, J., & Rangel, A. (2010). Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making. *Journal of Neuroscience*, 30(32), 10799–10808.
- Posner, M. I., Nissen, M. J., & Ogden, W. C. (1977). *Attended and unattended processing modes: The role of set for spatial location. Modes of perceiving and processing information* (pp. 137–157). Hillsdale, N.J.: Erlbaum.
- Rangel, A., & Clithero, J. A. (2013). The computation of stimulus values in simple choice. In Paul Glimcher & Ernst Fehr (Eds.), *Neuroeconomics: Decision-making and the brain* (2nd ed., pp. 125–147).
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108.
- Ratcliff, R. (1980). A note on modeling accumulation of information when the rate of accumulation changes over time. *Journal of Mathematical Psychology*, 21(2), 178–184.
- Ratcliff, R., Cheria, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two choice decisions. *Journal of Neurophysiology*, 90, 1392–1407.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922.
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2), 333–367.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Science*, 20(4), 260–281.
- Renkewitz, F., & Jahn, G. (2012). Memory indexing: A novel method for tracing memory processes in complex cognitive tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(6), 1622–1639.
- Reutskaja, E., Nagel, R., Camerer, C. F., & Rangel, A. (2011). Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, 101(2), 900–926.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108(2), 370–392.

- Russo, J. E., & Doshier, B. A. (1983). Strategies for multiattribute binary choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(4), 676–696.
- Russo, J. E., & Rosen, L. D. (1975). An eye fixation analysis of multialternative choice. *Memory & Cognition*, 3(3), 267–276.
- Schkade, D., & Johnson, E. J. (1989). Cognitive processes in preference reversals. *Organizational Behavior and Human Decision Processes*, 44(2), 203–231.
- Sütterlin, B., Brunner, T. A., & Opwis, K. (2008). Eye-tracking the cancellation and focus model for preference judgments. *Journal of Experimental Social Psychology*, 44(3), 904–911.
- Towal, R. B., Mormann, M., & Koch, C. (2013). Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proceedings of the National Academy of Sciences*, 110(40), 3858–3867.
- Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, 121(2), 179–205.
- Tsetsos, K., Chater, N., & Usher, M. (2012). Saliency driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109(24), 9659–9664.
- Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, 108(3), 550–592.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, 111(3), 757–769.
- Wedell, D. H., & Senter, S. M. (1997). Looking and weighting in judgment and choice. *Organizational Behavior and Human Decision Processes*, 70(1), 41–64.
- Willemsen, M. C., Böckenholt, U., & Johnson, E. J. (2011). Choice by value encoding and value construction: Processes of loss aversion. *Journal of Experimental Psychology: General*, 140(3), 303–324.
- Wollschläger, L. M., & Diederich, A. (2012). The 2N-ary choice tree model for N-alternative preferential choice. *Frontiers in Psychology*, 3(189), 1–11.