

Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions

Ian Krajbich^a and Antonio Rangel^{a,b,1}

Departments of ^aHumanities and Social Sciences and ^bComputation and Neural Systems, California Institute of Technology, Pasadena, CA 91125

Edited by James L. McClelland, Stanford University, Stanford, CA, and approved July 6, 2011 (received for review January 24, 2011)

How do we make decisions when confronted with several alternatives (e.g., on a supermarket shelf)? Previous work has shown that accumulator models, such as the drift-diffusion model, can provide accurate descriptions of the psychometric data for binary value-based choices, and that the choice process is guided by visual attention. However, the computational processes used to make choices in more complicated situations involving three or more options are unknown. We propose a model of trinary value-based choice that generalizes what is known about binary choice, and test it using an eye-tracking experiment. We find that the model provides a quantitatively accurate description of the relationship between choice, reaction time, and visual fixation data using the same parameters that were estimated in previous work on binary choice. Our findings suggest that the brain uses similar computational processes to make binary and trinary choices.

A basic goal of decision neuroscience is to characterize the computational processes used by individuals to make different types of decisions, as well as the neurobiological substrates of such computations (1–5). A significant amount of effort has been devoted to characterizing these processes in the realm of perceptual decision making involving two-alternative forced choices (2, 6–8). However, many important decisions do not fit this framework: they involve choices among multiple alternatives ($n > 2$) associated with different reward values (e.g., which food to select from a buffet table). Here we investigate these types of decisions.

The standard drift-diffusion model (DDM), as well as closely related versions, such as the leaky competitive accumulator (LCA) model (3, 4, 9), have been highly successful in providing quantitative explanations of the psychometrics, chronometrics, and neurometrics of binary perceptual choice (2, 10–16), and more recently in binary value-based choice (17–20). These models assume that decisions are made by accumulating stochastic information over time until the net evidence in favor of one option exceeds a pre-specified threshold. The size of the threshold can be chosen to optimally balance the benefit of accumulating more information with the cost of taking more time to reach a decision (21). Consider, for example, the canonical dot-motion task that has been widely used to study perceptual decision making. Here the stimulus itself is stochastic and each instant is thought to provide noisy but informative evidence for the net direction of movement in the display. Thus, as the individual accumulates more evidence, his knowledge about the true net direction of movement increases (2).

The DDM has also been shown to provide highly accurate descriptions of accuracy and response times in domains such as memory retrieval and decision-making, where the stimuli are not explicitly stochastic (17, 19, 20, 22–28); this suggests that these decisions might be made using a similar process of random information accumulation and integration. To see why, consider the case of binary value-based choice. Here the individual needs to compare the value of the two items and select the one with the highest value. If the value signals at any particular instant are independent stochastic draws from a common distribution (e.g., a Gaussian distribution with a mean equal to the item's true value), then it is beneficial to integrate the signals over time to identify the best option. In fact, it is easily shown that in the binary versions of these tasks, the DDM implements an optimal sequential proba-

bility ratio test (SPRT) (2, 3, 21, 29). Unfortunately, despite the success of the basic binary DDM, it has two important limitations.

First, there is an ongoing debate about the right way to generalize the model to multialternative choice ($n > 2$). The question is difficult because a large number of possible extensions with significantly different behavioral and neurobiological properties are possible (5, 27, 30–34). Furthermore, the optimal statistical test associated with the case of multiple alternatives is unknown, and only approximations have been proposed (32, 33). Out of the alternatives that have been proposed, the multihypothesis sequential probability ratio test (MSPRT) is a particularly appealing one, because it reduces to the SPRT in the case of two options, and is asymptotically optimal in the sense that for sufficiently low error rates it minimizes expected decision time (5, 25, 33, 35).

Second, the classic DDM and leaky integrator models have ignored the role that visual attention plays in the choice process. This gap is an important limitation because both casual observation and previous research suggest that visual fixations play a role in the decision-making process (18, 36, 37).

In a previous study we investigated the role of visual fixations on binary value-based choice (18). In that study, we found that a simple extension of DDM (6, 14, 15, 22–25, 38–41), in which fixations modulate the value integration process, provides a remarkably good quantitative account of the relationship between fixations, reaction times, and choices. The key idea of that model is that the slope of integration in the DDM is biased in favor of the item being fixated on at any particular time. Interestingly, although this leads to sizable choice biases in favor of options that are looked at more, we also found that the fixation process was only slightly influenced by relative value.

Here we propose a computational model of trinary value-based choice that addresses the two shortcomings listed above. The model is a natural generalization of our previous work. The key difference is that now we assume that there are three competing relative value signals, one for each option, with a rate of integration that depends on the fixation location. The model makes stark predictions about the relationship between the fixation patterns, reaction times, and choices. We test these predictions using an eye-tracking experiment in which subjects are shown high-resolution pictures of three food items and, after looking at them freely, indicate their choice through a button press (Fig. 1A).

The theory and experiments allow us to address two important questions about how the choice process changes when we go from the binary to the multialternative case. First, are the same computational processes at work in making binary and trinary choices? Second, are subjects able to identify and rule out especially bad alternatives early in the choice process by not looking at them? This

Author contributions: I.K. and A.R. designed research; I.K. and A.R. performed research; I.K. contributed new reagents/analytic tools; I.K. analyzed data; and I.K. and A.R. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: rangel@hss.caltech.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1101328108/-DCSupplemental.

second question is important because, given that fixations bias choices, a fixation process that directs relatively more attention to the best alternatives can significantly improve the quality of choices.

Results

Computational Model. To better understand the model for the trinary case it is useful to briefly review the case of binary choice developed in our previous study (18). That model assumed that the brain computes a relative decision value (RDV) signal that evolves over time as a Markov Gaussian process until a choice is made. The RDV starts each trial at zero, and every 1-ms step it exhibits a mean change that is proportional to the difference in value between the fixated and unfixated items, with a bias toward the fixated one. A left choice is made when the RDV crosses the barrier at +1; a right choice is made when it crosses -1. The model also assumed that fixations were exogenous to the integration process.

A useful observation to see the link with the trinary case is that the binary model described above can also be written as consisting of two competing parallel processes: one that computes the RDV in favor of the left option, and one that computes the RDV in favor of the right option. To be equivalent, the two must be co-dependent in the sense that at all times the two RDVs add up to zero, and a choice is made as soon as one crosses a +1 threshold.

The model for trinary choice that we propose here is a straightforward extension of the two parallel processes version of the binary choice model. Evidence in favor of each item is accumulated at different rates depending on the item's value and whether it is being fixated on. For example, when the subject is looking at the left item, the rates of evidence accumulation for the three items are given by

$$E_t^{left} = E_{t-1}^{left} + d^* r^{left} + \epsilon_t^{left},$$

$$E_t^{center} = E_{t-1}^{center} + \theta^* d^* r^{center} + \epsilon_t^{center},$$

$$E_t^{right} = E_{t-1}^{right} + \theta^* d^* r^{right} + \epsilon_t^{right},$$

where E_t denotes the evidence accumulated in favor of the option as of time t , r denotes the underlying value of each item (measured independently in the experiment), d is a constant controlling the speed of integration (in units of ms^{-1}), θ between 0 and 1 is a parameter reflecting the bias against the unfixated

options, and ϵ is white Gaussian noise with variance σ^2 (sampled independently every ms). The RDV signals for each option can then be defined as follows:

$$V_t^{left} = E_t^{left} - \max(E_t^{center}, E_t^{right}),$$

$$V_t^{center} = E_t^{center} - \max(E_t^{left}, E_t^{right}),$$

$$V_t^{right} = E_t^{right} - \max(E_t^{left}, E_t^{center}).$$

Just as before, a choice is made as soon as one of these RDVs crosses a barrier with a constant value of +1. The model has three free parameters: d , θ , and σ^2 ; these are the same free parameters as in the binary case.

Fig. 1B describes a simulated run of one decision trial for illustration. The RDV generally increases for the fixated item and decreases for the unfixated items, but the rate of change depends on the values of the three items. In the model there is only ever one positive RDV at a time, which represents the difference in evidence between the items with the most and next-most evidence. Therefore, the model falls into the class of “best vs. next” models that are known to implement the asymptotically optimal MSPRT in the absence of attention biases (5, 33, 35). We emphasize, however, that once visual attention biases are present, neither the two-item version of the DDM, nor the multi-item version presented here, is optimal. We chose the MSPRT framework to model the multi-item case because it is the most natural way to generalize DDM, and because, as we will see, it provides a remarkably accurate description of the data.

The model assumes that fixations are produced by a stochastic process that is exogenous to the path of the RDV. This does not rule out the possibility that the fixation process might be affected by the latent value of the stimuli. In fact, as is described in detail in *Methods* and *SI Methods*, in all of our data analyses we assume that fixation locations and lengths are drawn from the correct empirical distribution (as measured in the experiment), and that the integration process proceeds until the end of the fixation, in which case another fixation is drawn, or until a barrier is crossed, in which case the trial ends. By using this approach, our model and experimental tests focus on understanding the comparator process, while taking the fixation process as given (*SI Methods* and Fig. S1).

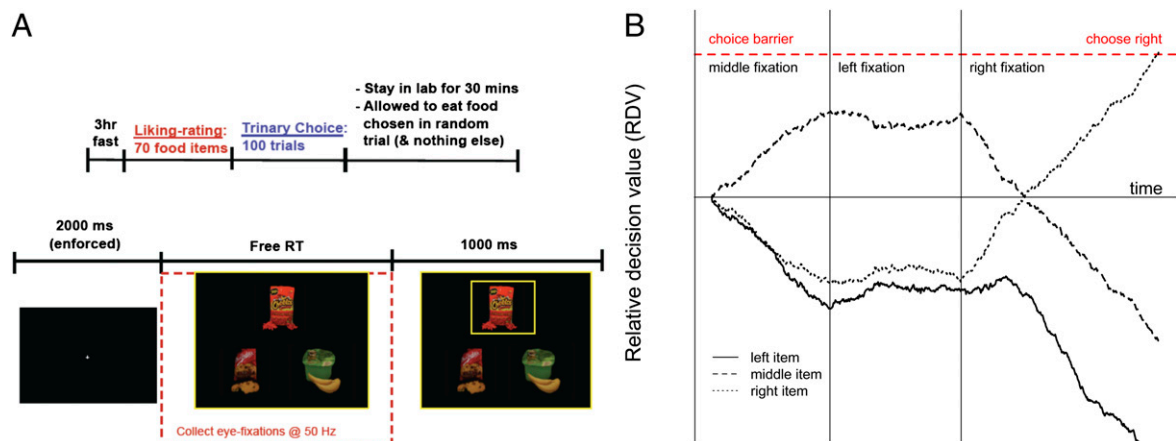


Fig. 1. (A) Task. Subjects were forced to fixate at the center of the screen for 2 s before the trial commenced, and were then presented with images of three food items and given as much time as needed to make a choice. After a selection was made, a yellow box highlighted the chosen item for 1 s. Fixations were recorded at 50 Hz. (B) Model. An RDV is computed for each item based on the evidence accumulated for that item compared with the highest accumulated evidence for the other items. The average rate of evidence accumulation is higher for an item when it is fixated on than when it is not. In addition to the average accumulation rate, there is also noise drawn from a Gaussian distribution. When one of the RDVs hits its barrier at +1, then that item is chosen. In this particular simulation of the model we assumed $r^{left} = 3$, $r^{center} = 5$, and $r^{right} = 7$, and it was simulated with an exaggerated signal-to-noise ratio (larger d and smaller σ) for illustrative purposes.

However, future work must explore the basis of the fixation process in more detail.

Hypotheses, Simulation, and Estimation. We hypothesized that our generalized model would be able to provide an accurate quantitative description of the fixation and choice data using the same parameters estimated in our previous binary experiment, which provides a stringent test of the hypothesis that similar computational processes are at work in binary and trinary value-based choice. These parameters are $\theta = 0.3$, $d = 0.0002 \text{ ms}^{-1}$, and $\sigma = 0.014$ (SI Results).

To see whether these parameters of the model were able to predict the data quantitatively, we simulated the model 50 times for each possible configuration of liking ratings (which provide independent measures of the value of each stimulus). The results of these simulations are described below, and are depicted in red in the appropriate figures.

Unless otherwise noted, throughout the following results, goodness-of-fit P values are based on two-sided t tests of the regression parameters against zero (SI Methods), and P values for trends in the subject data are based on two-sided t tests of the mixed-effects regression parameters against zero (SI Methods).

Basic Psychometrics. The simulated model accounted for the choice and reaction time curves well. The data show that choices were a logistic function of the difference between an item's value and the average value of the two other items (χ^2 goodness of fit, $P = 0.74$; Fig. 2A). Here and below, the items' values are measured using the independent liking ratings. Fig. 2B also shows that choice probabilities are affected by the value of the lowest-rated item, which indicates that all items, even the worst one, are taken into consideration during the choice process (χ^2 goodness of fit, $P = 0.64$; Fig. 2B). Fig. 2C shows that the mean reaction times, as a function of difficulty, were also accurately predicted by the model (goodness-of-fit slope: $P = 0.31$, intercept: $P = 0.13$) and that they increased with difficulty, as measured by the

difference in ratings between the best item and the average of the other two ($-145 \text{ ms per rating}$, $P < 0.001$).

We also looked at the match between the predicted and actual number of fixations as a function of difficulty (Fig. 2D). Although the slope of the fixation curve was not significantly different from the mixed-effects estimates from the data (goodness-of-fit slope: $P = 0.17$, intercept: $P = 10^{-8}$), the model systematically predicted 0.6 excess fixations. This mismatch is an unavoidable consequence of the procedures used to carry out our simulations, and does not reflect an inherent limitation of the model (SI Results).

Core Model Predictions and Choice Biases. The model with $\theta < 1$ makes several strong predictions about the relationship between fixations, choices, and reaction times that we tested using the eye-tracking data.

First, the model qualitatively predicts that final fixations will be biased toward the chosen item, unless the final fixated item is much worse than the others (Fig. 3A; χ^2 goodness of fit, $P = 0.02$). The intuition for the effect is that the evidence for the fixated item accumulates faster than the evidence for unfixated items, unless r_{fixated} is smaller than $\theta^* r_{\text{unfixated}}$.

Second, the model correctly predicts that the duration of the final fixation should be correlated with the excess amount of time that has been spent looking at the nonchosen items before that fixation (Fig. 3B; -0.26 ms/ms , $P = 0.0012$, goodness-of-fit slope: $P = 0.06$, intercept: $P = 0.04$). The intuition for this effect comes from the fact that, due to the attention bias, more time spent looking at the other items generally means more accumulated evidence in favor of those items, and thus a bigger gap in evidence against the chosen item must be overcome during the last fixation.

Third, the model correctly predicts that the probability of choosing an item should increase with the total excess time spent looking at that item (Fig. 3C; χ^2 goodness of fit, $P = 0.56$; mixed-effects logit, $P < 10^{-14}$). Here we define excess time as the total amount of time spent looking at the left item minus the sum total time spent looking at the other two items. One possible concern is that this effect could be a direct result of the fixation process if subjects tend to look more at higher-valued items, because higher-valued items are usually chosen; in our previous work (18) we showed that this was not the case in binary choice. To address this concern here, in Fig. 3D we replicate Fig. 3C, but using simulations of the model with $\theta = 1$, where the fixation location does not affect the drift rates of the relative value integrators. Because the model is simulated using the empirical fixation pattern, this figure tells us how much of the effect is due to a general tendency for subjects to look more at higher-valued items. In fact, only a small fraction of the effect can be explained by the influence of value on fixations, which suggests that most of the excess fixation time effect is due to the bias in the drift.

Fourth, the model predicts that, other things being equal, longer fixations to a particular item should be associated with a higher probability of choosing that item. Fig. 3E shows that this pattern also matches the data well using the first fixations (χ^2 goodness of fit, $P = 0.56$; mixed-effects logit, $P < 10^{-5}$).

Fifth, the model predicts that final fixations should be shorter than other nonfinal "middle" fixations, because crossing a barrier interrupts final fixations. Fig. 3F shows that this was the case (mean difference = 59 ms , $P = 0.032$).

Fixation Process. As in the case of binary choice, we found that the location of first fixations was random with respect to value ($P = 0.23$ based on a mixed-effects logit of item rating on first fixation probability), and that first and second fixations were shorter than later ones (mean difference = 161 ms , $P = 10^{-7}$ and mean difference = 122 ms , $P = 10^{-8}$, respectively; Fig. 3F). We found no effect of item value or choice difficulty (as measured by the difference in value between the best item and the average of the other two items) on middle fixation duration (-6.6 ms/rating , $P = 0.35$ and -7.5 ms/rating , $P = 0.21$, respectively). If instead we define choice difficulty as the difference in value between the best and middle item then we do see a significant but small effect

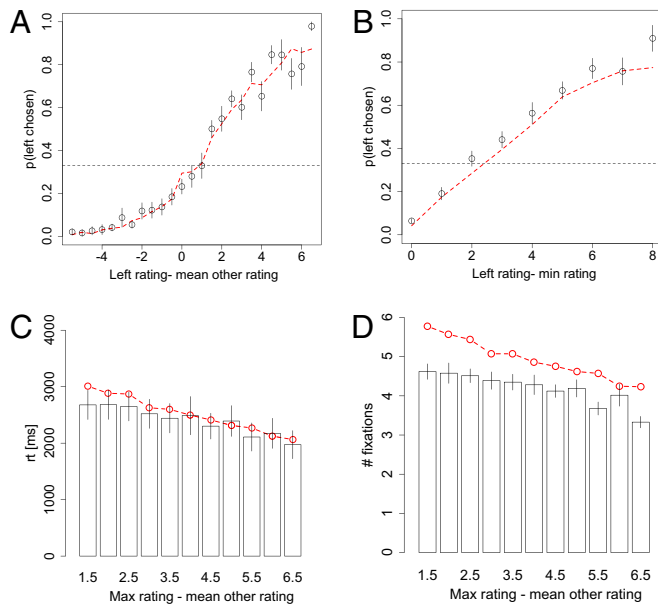


Fig. 2. Basic psychometrics. (A) Psychometric choice curve for the left item relative to the average rating of the other two items. (B) Psychometric choice curve for the left item relative to the worst-rated item. (C) Reaction time and (D) number of fixations as a function of the difference in liking ratings between the maximum rated item and the average of the other two items. In all panels, the axes were cut short whenever less than half of the subjects had any observations at those values. The red dashed lines indicate the simulated model. Bars denote SEs.

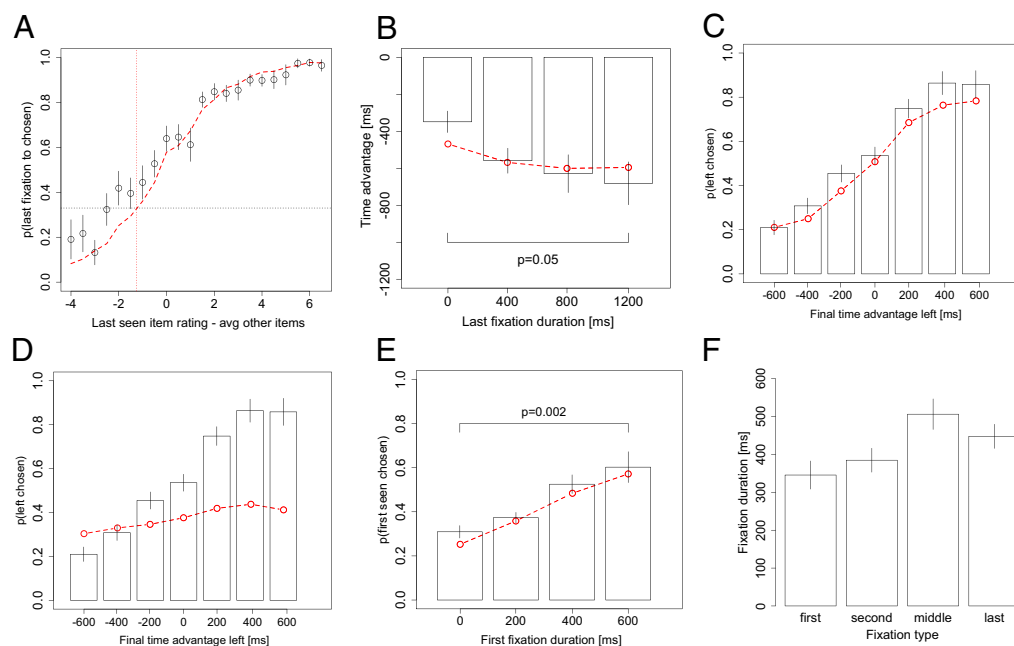


Fig. 3. Model prediction and choice biases. (A) Probability that the last fixation is to the chosen item as a function of the difference in liking ratings between the last seen item and the average of the other two items. (B) Difference in time spent looking at the unchosen items compared with the chosen item, before the last fixation, as a function of the duration of that last fixation. (C) Probability that left is chosen as a function of the excess amount of time that the left item was fixated on during the trial and (D) the same but for a model with $\theta = 1$ in which there is no fixation bias. (E) Probability that the first-seen item is chosen as a function of the duration of that first fixation. The values are rounded to the nearest 200 ms, so the bars correspond to 0–100 ms, 100–300 ms, 300–500 ms, and 500–700 ms. (F) Fixation duration by type. Middle fixations are any fixations that are not the first, second, or last fixations of the trial. Bars denote SEs. The red dashed lines indicate the simulated model. *P* values are based on paired two-sided *t* tests.

on middle fixation duration (-13.8 ms/rating, $P = 0.018$). We also found that middle fixation duration depended on the difference in ratings between the fixated item and the average of the other two items (18.5 ms/rating, $P = 0.0005$).

Fig. 4 *A* and *B* provides evidence that the underlying item values, as measured by the liking ratings, affected the fixation process late in the trial. The figure panels show the distribution of fixations to the best, middle, and worst items as a function of fixation number from the beginning of the trial and the end of the trial, respectively. The fixation location was mostly independent of value for the first three fixations but favored the best item later on. Also, the last fixation was usually to the best item, which is a natural consequence of the DDM with a visual fixation bias (Fig. 3*A*; see *SI Results* for an alternative explanation).

Fig. 4 *C* and *D* provides further insight into how the process of visual search changes during the course of a trial. The figure panels show the distribution of fixations to the left, center, and right items as a function of fixation number from the beginning of the trial and the end of the trial, respectively. The transitions were clearly dominated by location early in the trial but not at the end of the process. Specifically, 46% of first fixations were to the center item, suggesting that diagonal eye-movements from the central fixation cross may be less natural than horizontal or vertical movements. This is further supported by the fact that only 16% of second fixations are to the center item, because those fixations must be the result of diagonal fixation transitions.

Finally, Fig. 4*E* shows a histogram of the number of fixations that subjects made within a trial. Overall, on 83% of the trials, subjects made their decisions only after fixating on all three items at least once.

Discussion

The results presented here suggest that the same computational processes are at work in simple binary and trinary value-based choices. In both cases, a simple extension of the DDM in which fixations bias the slope of the value integration process provides

good quantitative accounts of the relationship between choice, reaction times, and fixation data. Remarkably, we were able to predict the data for the trinary case with high quantitative accuracy using parameters estimated from the binary case, demonstrating that the model is robust.

Our results also show that subjects tend to take all items into consideration when making a choice. In a vast majority of the trials, the subjects looked at all three items before making a choice, and though subjects were more likely to fixate on the best item as the trial progressed, they still continued to look at the worst item at a rate no less than 20%. These patterns are interesting for several reasons. First, they show that the fixation process is not fully independent from the valuation process, and contains an element of choice that needs to be explained in further work. Second, because items that are fixated on more are more likely to be chosen, this contributes to the quality of the choice process, although the effect is minor given the results in Fig. S1 in which the fixations were treated as fully random (as opposed to being described by the actual observed process; *SI Discussion*). However, subjects typically do not rule out the worst alternative, which degrades the quality of the choice process and leads to an overall worse percentage of choosing the best option (72% in the trinary case compared with 78% in the binary case) and a 6% chance of choosing the worst option.

The influence of value on the fixation process was taken into account in our numerical simulations, although the results do not change if we instead assume that fixations are random and independent of value or location after the first fixation (*SI Discussion* and Fig. S1). Thus, this aspect of the fixation process did not interfere with our ability to test the model of the value comparison process. Nevertheless, a critical goal for future research is to understand how the fixation process takes value into account, and how, if at all, it interacts with the integration process. However, in our previous binary choice work (18) we investigated an alternative class of models where there is no fixation bias in the DDM, but the RDV drives the termination of

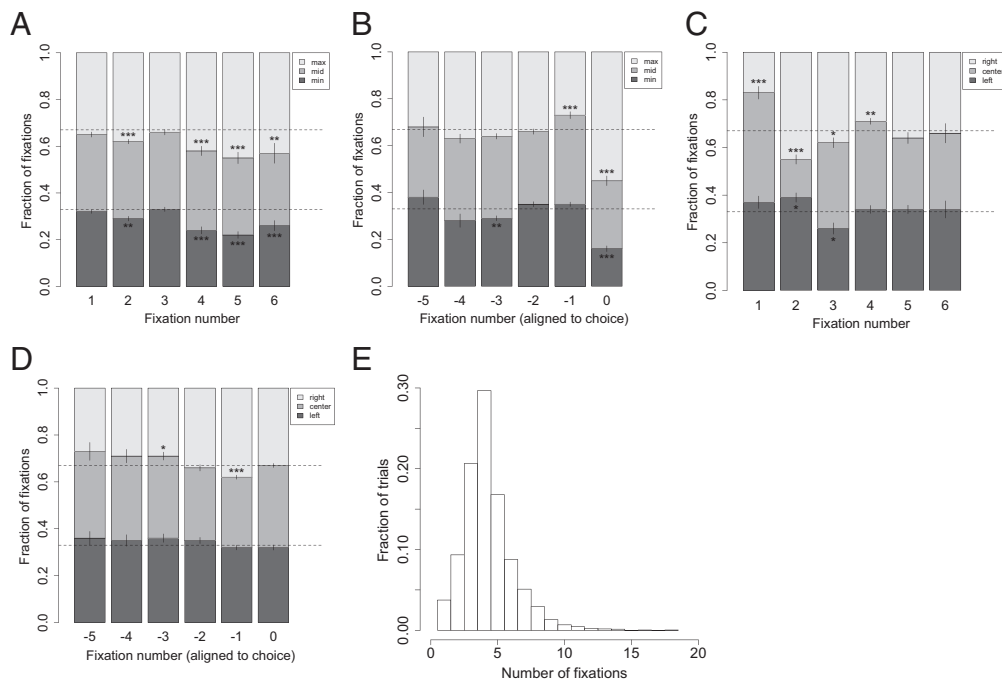


Fig. 4. Fixation properties. (A) Fraction of fixations to items with the maximum, middle, and minimum ratings as a function of fixation number from the beginning of the trial, and (B) from the end of the trial (0 indicates the final fixation). (C) Fraction of fixations to the left, center, and right items as a function of fixation number from the beginning of the trial and (D) from the end of the trial (0 indicates the final fixation). (E) Histogram of the number of fixations that subjects made within a trial. Bars denote SEs. Tests are based on paired two-sided *t* tests against 33.33% for the two extremes in each panel (maximum and minimum, or left and right). **P* = 0.05; ***P* = 0.0042 (Bonferroni corrected for the 12 tests in each panel), and ****P* = 0.001 (Bonferroni corrected for all 48 tests in the figure).

fixations. Though we could not rule out this entire class of models, we did show that a simple version failed to accurately account for basic properties of the last fixations, or the relationship between time exposure and choice probability. Furthermore, this class of models cannot explain the cross-subject correlation between first fixation location and the probability of choosing the first seen item. Therefore, it seems that this class of models is difficult, if not impossible, to fit to our data. Furthermore, another previous study by our group (36) manipulated fixation times exogenously and found that the manipulation affected choice probabilities. This evidence is consistent with our model, where fixations drive the drift process, and not vice versa.

One prominent alternative generalization of the DDM is the SPRT model, where the RDV compares the item with the most evidence with the average evidence for the other two items (5, 30, 33). Though this “best vs. average” model is not asymptotically optimal like the “best vs. next” MSPRT used here, it has been shown to do as well as the MSPRT in some perceptual decision-making datasets. We carried out a comparison of the two frameworks by also simulating the “best vs. average” extension of the model with the same parameters and fixation process used here (SI Discussion). As shown in Fig. S2, we found similar qualitative trends, but a worse overall fit than the “best vs. next” model. Of course, this analysis does not rule out the possibility that a different set of parameters might improve the fit of this model, but that would ignore one important feature of the model proposed here, which is that the same set of parameters fit the binary and trinary datasets with high quantitative accuracy.

Even in the binary choice case, there are many variations to the DDM that have been proposed in the literature (9, 17, 27, 42). This gives rise to a large number of possible model permutations, and thus to a difficult problem of model comparison (5, 27). From this viewpoint, it is remarkable that a simple extension of the binary case without the additional mechanisms was able to account for the experimental data with good quantitative accuracy using the previously identified parameters. Further-

more, in Krajbich et al. (18) we ran explicit model comparisons between some of these variations and our model, and found that they did not perform as well.

Finally, a remarkable aspect of our results is that the same computational processes, with the same parameters, seem to be at work in binary and trinary value-based choice. A basic question for future work is whether there is a size of the choice set at which these processes break down, because it seems unlikely that the brain would be able to track large numbers of alternatives. We hypothesize that the fixation-biased MSPRT process will hold up to a handful of items, but that it will eventually break down, to be substituted by a more complicated and yet-to-be-discovered algorithm.

Methods

Subjects. Thirty California Institute of Technology students participated in the experiment. Only subjects who self-reported regularly eating the snacks foods (e.g., potato chips and candy bars) used in the experiment and not on a diet were allowed to participate. Subjects were paid a \$20 show-up fee, in addition to receiving one food item. California Institute of Technology’s Human Subjects Internal Review Board approved the experiment.

Task. Subjects were asked to refrain from eating for 3 h before the start of the experiment. After the experiment, they were required to stay in the room with the experimenter for 30 min while eating the food item that they chose in a randomly selected trial (see below). Subjects were not allowed to eat anything else during this time.

In the initial rating phase, subjects entered liking ratings for 70 different foods using an on-screen slider bar (“How much would you like to eat this at the end of the experiment?”; scale –10 to 10). The initial location of the slider was randomized to reduce anchoring effects. This rating screen had a free response time. The food was kept in the room with the subjects during the experimental session to assure them that all of the items were available. Furthermore, subjects briefly saw all of the items at this point so that they could effectively use the rating scale.

In the choice phase, subjects made their choices using the keyboard. The choice screen had a free response time. Food items that received a negative

rating were excluded from the choice phase. We did not tell subjects about this feature of the experiment because doing so could have changed their incentives during the rating phase.

The items shown in each trial were randomly chosen. In all trials the three items were displayed in a triangular formation with the left and right items at the same vertical position, and the center item at the opposite vertical position. In half of the trials the center item was on the top half of the screen, and in the other half it was on the bottom half of the screen. We do not distinguish between these two types of trials in the paper.

Subjects indicated their choice by pressing the left, down, or right arrow keys for the left, center, and right items, respectively. After subjects indicated their choice, a yellow box was drawn around the chosen item (with the other item still on the screen) and displayed for 1 s, followed by a fixation screen, before the beginning of the next trial.

Eye Tracking. Subjects' fixation patterns were recorded at 50 Hz using a Tobii desktop-mounted eye tracker. Before each choice trial, subjects were required to maintain a fixation at the center of the screen for 2 s before the items would appear, ensuring that subjects began every choice fixating on the same location.

Data Analysis. The eye-tracking data were processed using the same procedures used in our binary choice work (18), and is detailed in *SI Methods*.

- Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci* 9:545–556.
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Bogacz R (2007) Optimal decision-making theories: Linking neurobiology with behaviour. *Trends Cogn Sci* 11:118–125.
- Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27:161–168.
- Ditterich J (2010) A comparison between mechanisms of multi-alternative perceptual decision making: Ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Front Neurosci* 4:184.
- Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111:333–367.
- Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci* 9:467–479.
- Romo R, Salinas E (2003) Flutter discrimination: Neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4:203–218.
- Usher M, McClelland JL (2001) The time course of perceptual choice: The leaky, competing accumulator model. *Psychol Rev* 108:550–592.
- Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5:10–16.
- Ditterich J (2006) Evidence for time-variant decision making. *Eur J Neurosci* 24:3628–3641.
- Ditterich J, Mazurek ME, Shadlen MN (2003) Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat Neurosci* 6:891–898.
- Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat Neurosci* 9:682–689.
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cereb Cortex* 13:1257–1269.
- Shadlen MN, Newsome WT (1996) Motion perception: seeing and deciding. *Proc Natl Acad Sci USA* 93:628–633.
- Milosavljevic M, Malmaud J, Huth A, Koch C, Rangel A (2010) The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgm Decis Mak* 5:437–449.
- Krajibich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13:1292–1298.
- Philiastides MG, Biele G, Heekeren HR (2010) A mechanistic account of value computation in the human brain. *Proc Natl Acad Sci USA* 107:9430–9435.
- Basten U, Biele G, Heekeren HR, Fiebach CJ (2010) How the brain integrates costs and benefits during decision making. *Proc Natl Acad Sci USA* 107:21767–21772.
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113:700–765.
- Ratcliff R (1978) A theory of memory retrieval. *Psychol Rev* 85:59–108.
- Ratcliff R (2002) A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychon Bull Rev* 9:278–291.
- Ratcliff R, Rouder JN (2000) A diffusion model account of masking in two-choice letter identification. *J Exp Psychol Hum Percept Perform* 26:127–140.
- Ratcliff R, McKoon G (1997) A counter model for implicit priming in perceptual word identification. *Psychol Rev* 104:319–343.
- McKeeff TJ, Tong F (2007) The timing of perceptual decisions for ambiguous face stimuli in the human ventral visual cortex. *Cereb Cortex* 17:669–678.
- Leite FP, Ratcliff R (2010) Modeling reaction time and accuracy of multiple-alternative decisions. *Atten Percept Psychophys* 72:246–273.
- Domenech P, Dreher JC (2010) Decision threshold modulation in the human brain. *J Neurosci* 30:14305–14317.
- Wald A, Wolfowitz J (1948) Optimum character of the sequential probability ratio test. *Ann Math Stat* 19:326–339.
- Niwa M, Ditterich J (2008) Perceptual decisions between multiple directions of visual motion. *J Neurosci* 28:4435–4445.
- Bogacz R, Usher M, Zhang J, McClelland JL (2007) Extending a biologically inspired model of choice: Multi-alternatives, nonlinearity and value-based multidimensional choice. *Philos Trans R Soc Lond B Biol Sci* 362:1655–1670.
- Bogacz R, Gurney K (2007) The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput* 19:442–477.
- McMillen T, Holmes P (2006) The dynamics of choice among multiple alternatives. *J Math Psychol* 50:30–57.
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11:693–702.
- Dragalin VP, Tartakovsky AG, Veeravalli VV (1999) Multihypothesis sequential probability ratio tests—part I: Asymptotic optimality. *IEEE Trans Inf Theory* 45:2448–2461.
- Armel C, Rangel A (2008) Biasing simple choices by manipulating relative visual attention. *Judgm Decis Mak* 3:396–403.
- Shimojo S, Simion C, Shimojo E, Scheier C (2003) Gaze bias both reflects and influences preference. *Nat Neurosci* 6:1317–1322.
- Ratcliff R, McKoon G (2008) The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Comput* 20:873–922.
- Gold JI, Shadlen MN (2002) Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36:299–308.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936.
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489.
- Cisek P, Puskas GA, El-Murr S (2009) Decisions in changing conditions: The urgency-gating model. *J Neurosci* 29:11560–11571.