

Visual fixations and the computation and comparison of value in simple choice

Ian Krajbich¹, Carrie Armel² & Antonio Rangel^{1,3}

Most organisms facing a choice between multiple stimuli will look repeatedly at them, presumably implementing a comparison process between the items' values. Little is known about the nature of the comparison process in value-based decision-making or about the role of visual fixations in this process. We created a computational model of value-based binary choice in which fixations guide the comparison process and tested it on humans using eye-tracking. We found that the model can quantitatively explain complex relationships between fixation patterns and choices, as well as several fixation-driven decision biases.

There is a growing consensus in behavioral neuroscience that the brain makes simple choices by first assigning a value to all of the options under consideration and then comparing them^{1–3}. This has motivated an interest in the computational properties of value comparison processes, and in understanding the extent to which they can generate reward-maximizing choices.

Although many popular models of value-based choice implicitly assume that the comparison process involves a trivial instantaneous maximization problem^{4,5}, casual observation suggests that the underlying processes are more sophisticated and that visual fixations are likely to be involved. Consider, for example, a typical buyer at the grocery store choosing between two snacks. Instead of approaching the shelf and immediately selecting his preferred option, the individual's gaze shifts repeatedly between the items until one of them is eventually selected.

We propose a model of how simple value-based binary choices are made and of the role of visual fixations in the comparison of values. The model makes stark qualitative and quantitative predictions about the relationship between fixation patterns and choices, which we test using eye-tracking (Fig. 1a). Subjects are shown high-resolution pictures of two food items and are free to look at them as much as they want before indicating their choice with a button press.

The theory developed here builds on the framework of drift-diffusion models (DDM) of binary response selection^{6–19}, and especially on applications of these models to the realm of perceptual decision making^{18,20–31}, where they have provided accurate descriptions of the psychometric data and important insights into the activity of the lateral intraparietal area (LIP). These models assume that stochastic evidence for one response (compared to the other) is accumulated over time until the integrated evidence passes a decision-threshold and a choice is made. The level of the threshold is set to balance the benefit of accumulating more information with the cost of taking more time to reach a decision.

There are two key differences between our work and the previous studies on perceptual discrimination. First, in both tasks subjects must

determine the value of two potential responses, but in perceptual discrimination tasks subjects typically see a single stochastic stimulus, whereas in our task subjects see two non-stochastic pictures of food items. Second, fixations are not involved in the standard perceptual discrimination task because subjects maintain central fixation at all times, whereas here the fixations are crucial for the decisions. The key idea of our model is that fixations affect the DDM value comparison process by introducing a temporary drift bias toward the fixated item. This drift bias in turn leads to a choice bias for items that are fixated on more.

RESULTS

Computational model

Following the literature on DDMs, our model assumes that the brain computes a relative decision value (RDV) that evolves over time as a Markov Gaussian process until a choice is made (Fig. 1b). The RDV starts each trial at 0 and continually evolves over time at one of two possible rates (depending on which item is fixated), and a choice is made when it reaches a barrier at either +1 or –1. If the RDV reaches the +1 threshold the left item is chosen and if it reaches the –1 threshold the right item is chosen.

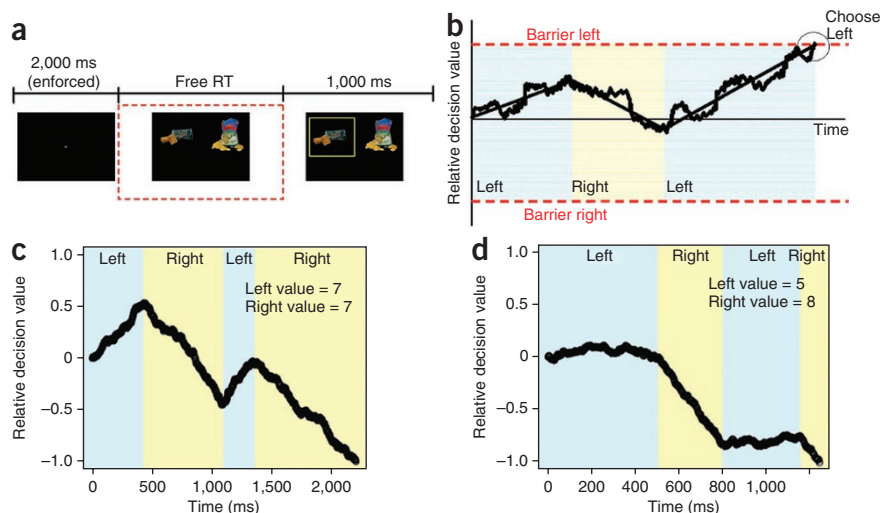
The key difference between our model and the standard drift diffusion model is that in our model the slope with which the RDV signal evolves at any particular instant depends on the fixation location. In particular, the slope is proportional to the weighted difference between the values of the fixated and unfixated items. The weight discounts the value of the unfixated item relative to the fixated one. When the subject is looking at the left item the RDV changes according to $V_t = V_{t-1} + d(r_{\text{left}} - \theta r_{\text{right}}) + \varepsilon_t$ and when he looks at the right item, it changes according to $V_t = V_{t-1} + d(r_{\text{right}} - \theta r_{\text{left}}) + \varepsilon_t$ where V_t is the value of the RDV at time t , r_{left} and r_{right} denote the values of the two options, d is a constant that controls the speed of integration (in units of ms^{-1}), θ between 0 and 1 is a parameter that reflects the bias toward the fixated option, and ε_t is white Gaussian noise with variance σ^2 (randomly sampled once every millisecond).

¹Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, California, USA. ²Precourt Institute for Energy Efficiency, Stanford University, Palo Alto, California, USA. ³Computational and Neural Systems, California Institute of Technology, Pasadena, California, USA. Correspondence should be addressed to A.R. (rangal@hss.caltech.edu).

Received 23 February; accepted 17 August; published online 12 September 2010; corrected after print 10 February 2011; doi:10.1038/nn.2635

Figure 1 Experiment and model.

(a) Choice trial. Subjects are forced to fixate at the center of the screen for 2 s. They are then presented with images of two items and given as much time as they want to make their choice. After a selection is made a yellow box highlights the chosen item for 1 s. RT, reaction time. (b) Model. A relative decision value (RDV) evolves over time with a slope that is biased toward the item that is being fixated. The slope dictates the average rate of change of the RDV, but there is also an error term drawn from a Gaussian distribution. When the RDV hits the barrier a choice is made for the corresponding item. The shaded vertical regions represent the item being fixated. (c,d) Simulated runs of the model using $d = 0.005$, $\sigma = 0.05$ and $\theta = 0.6$, to give a better intuition for the decision process.



With respect to the fixation process, the model assumes that the first fixation goes to the left item with probability p (independent of the values of the options), that fixations alternate between the two items until a barrier is crossed, and that fixations have a maximum duration given by a fixed distribution that depends on the difficulty of the choice, as measured by $r_{\text{best}} - r_{\text{worst}}$. A fixation terminates if either its maximum duration is reached, or the RDV terminates the choice process by crossing a barrier.

In simulated runs of the model, the RDV generally moved toward the fixated item, but the slope depended on the values of the two items (Fig. 1c,d). For example, the RDV signal integrated toward the left barrier when the subject fixated on the left item, even though it had a lower value than the right item (Fig. 1d). This suggests that visual fixations are important for the integration process.

Hypotheses and model fitting

We carried out a simple eye-tracking experiment to investigate the extent to which the drift diffusion model outlined here captures key patterns of the relationship between the fixation and choice data. We were particularly interested in distinguishing between three alternative models: model 1, the regular DDM given by the case of $\theta = 1$, model 2, a DDM with full fixation bias given by the case of $\theta = 0$, and model 3, a DDM with partial fixation bias given by the case $0 < \theta < 1$. The experiment consisted of two stages. In the first stage subjects rated how much they would like to eat 70 food items at the end of the experiment (scale -10 to 10). The liking ratings provide an independent measure of the value of individual items. In the second stage subjects made 100 choices between pairs of neutral or appetitive foods. Afterwards they ate the item chosen in a randomly selected trial. We measured eye-movements at 50 Hz.

We fitted the model to the even-numbered trials of the group data using maximum likelihood estimation (MLE) on the observed choices and reaction times. The best fitting model had parameters $d = 0.0002$ ms⁻¹, $\theta = 0.3$ and $\sigma = 0.02$, with a log-likelihood value of -3,704.

Table 1 Summary of goodness-of-fit statistics

Figure	2a	2b	2c	4b	4c	5a, left	5a, right	5b	5c	5d	5e
$\theta = 0.3$	0.92*	0.1*	0.39*	0.997*	0.824	0.96*	0.96*	0.75*	0.0062*	0.21	0.0016
$\theta = 0$	10 ⁻⁵	0.01	10 ⁻⁵	0.01	0.96*	0.83	0.19	0.0002	10 ⁻¹³	0.76*	0.1*
$\theta = 1$	10 ⁻¹⁶	0.0007	10 ⁻¹⁵	10 ⁻¹⁶	0.04	10 ⁻¹⁶	0.0001	10 ⁻¹³	10 ⁻¹¹	0.0009	10 ⁻⁹

Each number is the P -value from the goodness-of-fit test of that model to the data. Note that the intermediate model ($\theta = 0.3$) fits better than the other two models in most cases except **Figures 4c** and **5d,e** (*).

We used the same procedure to fit models with $\theta = 1$ and $\theta = 0$. In both cases the best-fitting models had parameters $d = 0.0002$ ms⁻¹ and $\sigma = 0.02$. We then used the likelihood ratio statistic to test the hypothesis that θ was significantly less than 1 (log-likelihood = -3,708, $P = 0.008$) and significantly larger than 0 (log-likelihood = -3,710, $P = 0.0005$). This provides support for model 3 over the standard and full fixation bias drift-diffusion models (see also **Table 1** and **Supplementary Figs. 1–10**, which compare the fits of the three models). We also carried out a restricted fit of the model to individual subject data. The mean (s.d.) estimated θ value from individual model fits was 0.52 (0.3), and 35/39 subjects had an estimate of this parameter less than 1 (**Supplementary Figs. 11** and **12**).

To investigate the ability of the model to predict the data quantitatively, we then simulated the model 1,000 times for each pair of liking ratings, using the estimated maximum likelihood parameters, and by sampling fixation lengths from the empirical fixation data (taking into account that fixation durations are related to decision difficulty, as described below). We assumed that the location of the first fixation is chosen probabilistically to match the empirical data (look left first with probability 74%). The results of the simulations are described below. Note that in all comparisons of the model with the data, we present only the odd-numbered trials, as the model was fitted to the even-numbered trials.

Basic psychometrics

The model predicted the choice and reaction time curves quite well in the odd-numbered trials. The choice data (χ^2 goodness-of-fit statistic = 4.47, $P = 0.92$) indicated that choices were a logistic function of the value differences (**Fig. 2**), which means that the best option was selected only 78% of the time. Note that the amount of noise in the choice process is controlled by the Gaussian noise in the integration process and by the random fixation durations. The $\theta = 0$ model was comparably poor fitting.




Figure 2 Basic psychometrics. (a) Psychometric choice curve. P, probability. (b) Reaction times as a function of the difference in liking ratings between the best and worst items, which is a measure of difficulty. (c) Number of fixations in a trial as a function of choice difficulty. The black dashed lines indicate the simulated data using the MLE parameters. Subject data includes only odd-numbered trials. In **a**, the gray dash-dotted line indicates the simulated data for the $\theta = 0$ model. Bars denote s.e., clustered by subject. Tests are based on a paired two-sided *t*-test.

We also examined reaction times and number of fixations (goodness-of-fit: $P = 0.10$ and 0.39 , respectively; **Fig. 2**). Both measures correlated with difficulty (mixed effects regression estimate -211 ms per rating, $P = 10^{-11}$ and -0.171 fixations per rating, $P = 10^{-15}$, respectively), a property of drift-diffusion models that also extends to our model.

Properties of the visual search process

The model makes strong assumptions about the nature of the fixation process (**Fig. 3**). First, the probability that the first fixation was to the best item was not significantly different from 0.5 and was unaffected by the difference in ratings (**Fig. 3a**) (mixed effects regression estimates: intercept = 0.518 , $P = 0.31$; slope = -0.0009 per rating, $P = 0.88$). Second, the middle fixation durations were independent of the value of the fixated items (mixed effects regression estimate: 6.4 ms per rating, $P = 0.21$; **Fig. 3b**). Third, the middle fixation duration depended slightly on the difference in value between the fixated and nonfixated items (mixed effects regression estimate: 11.4 ms per rating, $P = 0.0052$; **Fig. 3c**), but depended more on the difficulty of the decision (mixed effects regression estimate -33.8 ms per rating, $P = 10^{-5}$; **Fig. 3d**).

In the estimation and simulation procedures, we took into account the dependency of middle fixation durations on value (**Fig. 3d** and **Supplementary Figs. 13–15**). The fixation distributions were best approximated by log-normal distributions (**Supplementary Table 1** and **Supplementary Figs. 16–22**).

Core model predictions

The model makes several strong predictions about the relationship between visual fixations, choices and reaction times and we tested

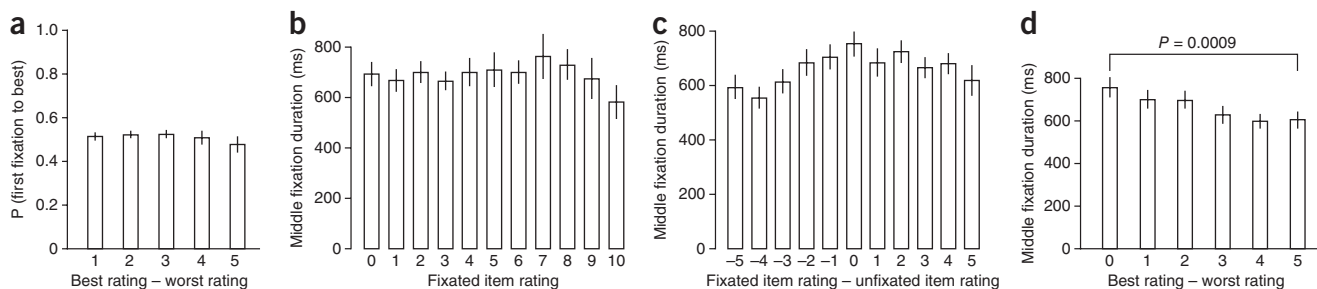


Figure 3 Fixation properties. (a) Probability that the first fixation is to the best item. In all cases they are not significantly different from 50%. (b) Middle fixation duration as a function of the liking rating of the fixated item. (c) Middle fixation duration as a function of the difference in liking ratings between the fixated and unfixated items. (d) Middle fixation duration as a function of the difference in liking ratings between the best- and worst-rated items. Bars denote s.e., clustered by subject. Tests are based on a paired two-sided *t*-test.

these predictions using the eye-tracking data. First, consistent with the data, the model predicts that final fixations should be shorter than middle fixations, as fixations are interrupted when a barrier is crossed ($P = 0.0002$; **Fig. 4a**). First fixations were also shorter than middle ones ($P = 10^{-14}$) which was not predicted *ex ante* by the theory, but which we incorporated *ex post* into the computational model's estimation and simulation procedures.

Second, the model correctly predicts that subjects will generally choose the item they looked at last, unless that item is much worse than the other one (χ^2 goodness-of-fit statistic = 1.96 , $P = 0.997$; **Fig. 4b**). To see

why, recall that the RDV climbs toward the barrier of the fixated item unless the fixated item is worse enough than the other item that the drift rate becomes negative. The $\theta = 0$ model cannot account for this pattern (**Fig. 4b**).


Third, the model predicts that the longer you have looked at item A during a trial, the longer you will have to look at item B before choosing it over item A. The intuition is simple: on average, the longer one looks at item A the farther the RDV gets from item B's barrier, and thus the farther it will have to travel back to hit that threshold. This was approximately the case (mixed effects regression coefficient = -0.08 , $P = 0.11$, median-split test $P = 0.03$; **Fig. 4c**). The estimated model is consistent with this pattern (goodness-of-fit: $P = 0.82$), whereas the $\theta = 1$ model is not.

Choice biases

The model also predicts that when $\theta < 1$ the decision processes should show several choice biases. First, it predicts a last-fixation bias: subjects should be more likely to choose an item (for a given rating difference) if their last fixation is to that item as opposed to the other one. This is a direct implication of the fact that the value of the unfixated item is discounted. As predicted, there was a sizable bias in both the simulated and the subject data (logit mixed effect regression: $P = 10^{-16}$, χ^2 goodness-of-fit statistic = 3.64 , $P = 0.96$ for last fixation left and = 3.58 , $P = 0.96$ for last fixation right; **Fig. 5a**). In contrast, the $\theta = 1$ model predicts that the last fixation will have no effect.

Second, the model predicts that there should be a choice bias that depends on the total amount of time spent looking at one item versus the other: controlling for value differences, the probability of choosing an item should increase with the excess time for which it is fixated. A mixed effects logit regression shows that this was the case

Figure 4 Basic model predictions. **(a)** Fixation duration by type. Middle fixations indicate any fixations that were not the first or last fixations of the trial. **(b)** Probability that the last fixation is to the chosen item as a function of the difference in liking ratings between the fixated and unfixated items in that last fixation. **(c)** Amount of time spent looking more at Item B before the last fixation (to Item A), as a function of the duration of that last fixation. The black dashed line indicates the simulated data using the MLE parameters. Subject data includes only odd-numbered trials. In **b**, the gray dash-dotted line indicates the simulated data for the $\theta = 0$ model, and the vertical dotted lines indicate the points at which the simulation curves cross the horizontal line at chance. In **c**, the gray dash-dotted line indicates the simulated data for the $\theta = 1$ model. Bars denote s.e., clustered by subject.



($P = 10^{-8}$). This prediction follows from the fact that the RDV always evolves more toward an item's barrier when it is being looked at than when it is not. Fixation duration and order are independent of an item's value, highlighting this bias in both the data and the simulations (χ^2 goodness-of-fit statistic = 3.43, $P = 0.75$; Fig. 5b). We further tested this hypothesis by correcting for the difference in liking ratings. For each trial we took the actual choice (1 or 0) and subtracted the average probability that left was chosen in all trials with that difference in liking ratings. These 'corrected' choice probabilities were plotted as a function of the fixation time advantage for the left item (goodness-of-fit: $P = 0.0062$; Fig. 5c). This eliminates any possible influence of the measured liking ratings on the fixation durations and shows that there is a substantial effect of total fixation time on choice. In contrast, the $\theta = 1$ model predicts that exposure time will have no effect on choice.

In a related result, the duration of the first fixation was correlated with choice (χ^2 goodness-of-fit statistic = 4.55, $P = 0.21$, mixed effects regression: $P = 0.028$; Fig. 5d). This effect is still there after correcting for the difference in liking ratings (goodness-of-fit: $P = 0.0016$; Fig. 5e). As before, this relationship is not predicted by the $\theta = 1$ model.

Third, the model correctly predicts that any left-looking biases should translate into left-choice biases (Fig. 5f). The more likely a

subject was to look left first, the more likely he was to choose items on the left, with a correlation of 0.38 ($P = 0.017$) and a Spearman's non-parametric correlation of 0.49 ($P = 0.002$).

Alternative models

The results above showed that our variant of the drift-diffusion model can account for a wide range of correlations between the pattern of fixations, reaction times and choices. A natural question to ask is whether alternative models can provide a similar account of the data. Although a comprehensive estimation and comparison of alternative models is beyond the scope of this study, here we present exploratory results regarding the ability of three natural alternatives to account for the key patterns in the data (Online Methods).

The first alternative model is a DDM in which the RDV signal evolution is independent of the fixations (so that $\theta = 1$). Instead, the model reverses the direction of causality by assuming that fixation lengths are affected by the local value of the RDV signal. In particular, fixations to an item are longer the more the RDV favors that item. This alternative model is interesting because it produces a correlation between fixation lengths and choices without giving a causal role to fixations.

We investigated the qualitative properties of this model by assuming a simple and concrete functional form for the impact of the RDV

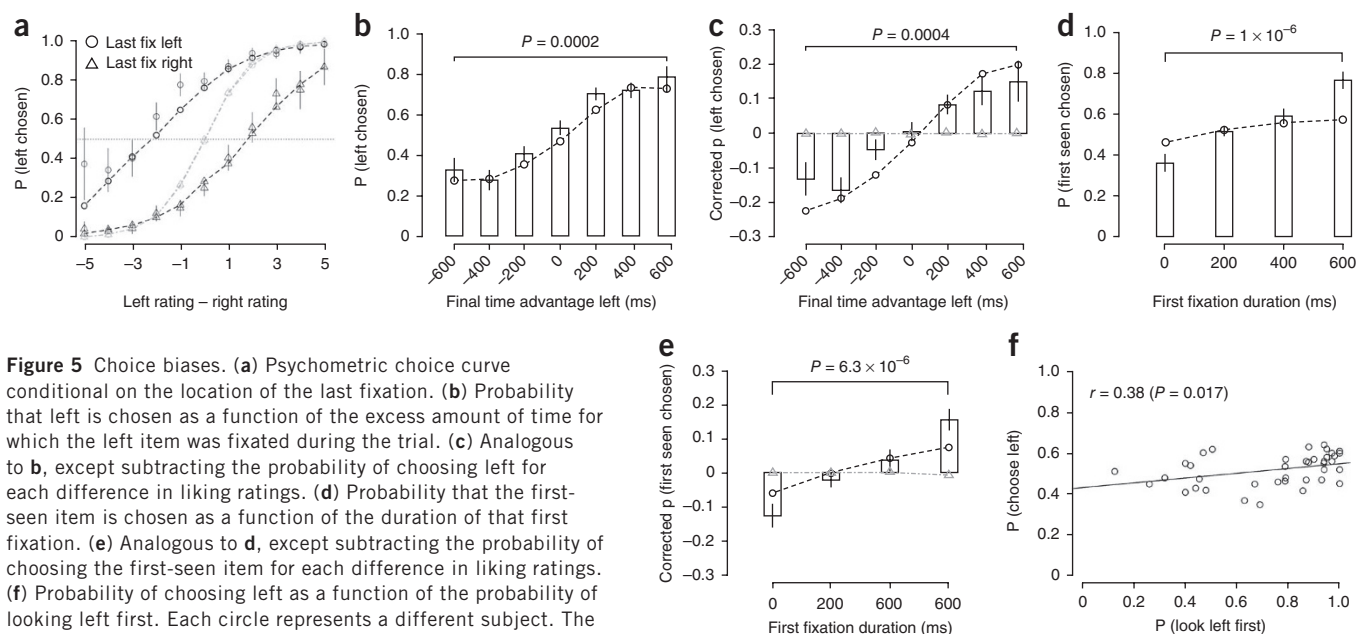


Figure 5 Choice biases. **(a)** Psychometric choice curve conditional on the location of the last fixation. **(b)** Probability that left is chosen as a function of the excess amount of time for which the left item was fixated during the trial. **(c)** Analogous to **b**, except subtracting the probability of choosing left for each difference in liking ratings. **(d)** Probability that the first-seen item is chosen as a function of the duration of that first fixation. **(e)** Analogous to **d**, except subtracting the probability of choosing the first-seen item for each difference in liking ratings. **(f)** Probability of choosing left as a function of the probability of looking left first. Each circle represents a different subject. The black dashed line indicates the simulated data using the MLE parameters. The gray dash-dotted line indicates the simulated data for the $\theta = 1$ model. Subject data includes only odd-numbered trials. Bars denote s.e., clustered by subject. Tests are based on a paired two-sided t -test, except in **f**, where we used standard two-sided t -tests.

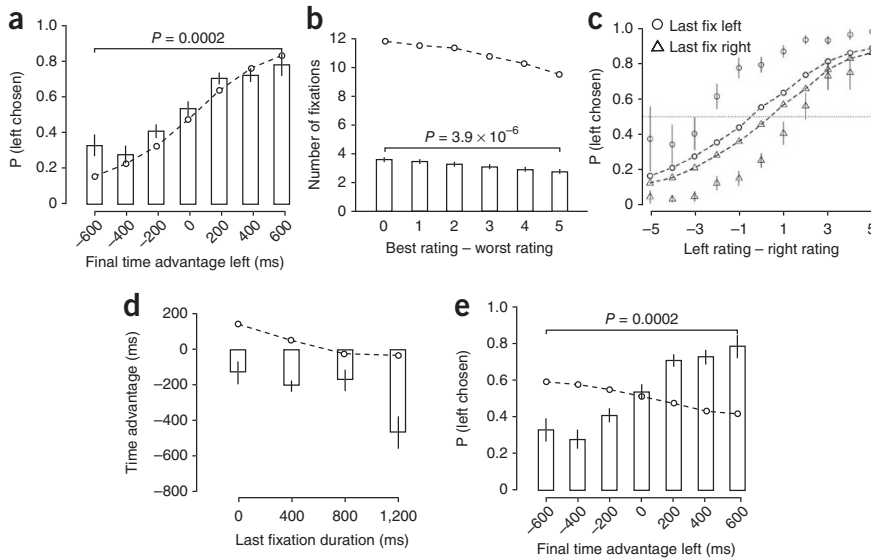


Figure 6 Alternative models. (a,b) Replications of **Figures 5b** and **2c** for alternative model 1. (c,d) Replications of **Figures 5a** and **4c** for alternative model 2. (e) Replication of **Figure 5b** for alternative model 3. The black dashed lines indicate the simulated data using the alternative models. Subject data includes only odd-numbered trials. Bars denote s.e., clustered by subject. Tests are based on paired two-sided *t*-tests.

on the probability of terminating fixations with one free parameter P^* (Online Methods). All other parameters were assumed to be the same as those that we obtained in the best-fitting version of our model with $\theta = 1$. The free parameter P^* was chosen to recreate the effect of net total fixation time on choice probability. The model could do this quite well (**Fig. 6a**). Unfortunately, however, the value of P^* needed to match this psychometric curve also generated too many fixations (**Fig. 6b**). This problem could not be overcome by adjusting two other key parameters of the model: the integration slope d , or the Gaussian noise σ (**Supplementary Figs. 23** and **24**).

The second alternative model assumes that the RDV signal is modulated by fixations, as in our main model, but differs on how a decision is triggered. It assumes that the reaction time is determined exogenously by a separate unmodeled process, and that the subject chooses the option with the best RDV at that time. We investigated the qualitative properties of this model by simulating it using the best-fitting parameters from the $\theta = 0.3$ version of our model, and randomly sampling decision times from the actual empirical distribution.

Owing to its similarity to our model, it is not surprising that the model fits most of the psychometric data well (**Supplementary Fig. 25**). However, this set of parameters had problems accounting for the effect of the last fixation on choice, as well as the relationship between the last fixation duration and the relative fixation advantage (**Fig. 6c,d**). Varying other key parameters of the model did not improve the ability of the model to overcome this limitation (**Supplementary Figs. 26–28**). A comparison of the forces at work in both models provides an intuitive idea of why this is the case. In the barrierless model, there is no reason to expect that the final fixation duration would depend on the current fixation advantage for the other item, as the fixation is terminated exogenously. Furthermore, as the reaction time is determined exogenously, instead of being determined by crossing a barrier, the integration bias at work in our model is more limited here.

The third alternative model is similar to ours, except that the fixations change the locations of the choice barriers rather than the

drift rate. We simulated this model using the best-fitting parameters from the $\theta = 1$ model, and assuming that the magnitude of the barrier for the fixated item is lowered from 1 to 0.8. The value of 0.8 was chosen so that the model would fit the basic psychometric data (**Supplementary Fig. 29**). However, although the model approximated the basic psychometric data fairly well, it could not reproduce the effect of total fixation duration on choice (**Fig. 6e**), even when we varied the magnitude of the barrier drop (**Supplementary Fig. 30**). Note that in this model there are two competing forces at work. First, the longer an item is seen, the more time that item has a lower barrier and the more likely it is to be chosen. Second, the decrease in the barrier height for the fixated item makes the final fixations shorter than they would have otherwise been. The simulations show that the second effect is dominant.

DISCUSSION

The results presented here provide insight into the nature of the computational and psychological processes that guide simple choices. In particular, we found that a simple extension of the DDM in which fixations are involved in the value integration process could provide a remarkably good quantitative account of various relationships between the fixation and choice data, as well as of several sizable choice biases.

An important question raised by our results is whether the visual fixation process has a causal effect on the value comparison process. Several pieces of evidence suggest that this might be the case. First, our model assumes a causal effect and fits important moments of the data notably well. Second, we have shown that one simple alternative DDM in which values affect fixations, but in which the opposite is not true, does not seem to be able to simultaneously account for all the trends in the data. Third, consistent with the findings reported here, related studies have shown that it is possible to bias choices by manipulating relative fixation durations exogenously^{32,33}. However, it is important to emphasize that by itself, the evidence provided here is not sufficient to establish a causal effect of fixations on choices.

Our model does not rule out the possibility that values have some effect on the pattern of fixations. In fact, as assumed in the model's estimation and simulation procedures, fixation durations increased with the difficulty of the choice (**Fig. 3d**). However, our results suggest that even if these feedback results are present, random variation in fixation durations can affect the choice process directly. In this study we have treated these feedback effects from values to fixations as exogenous. Understanding the computational properties of these effects is an important open question for future research in this area.

Our exploratory robustness analysis casts some doubt on the ability of three natural alternative models of how fixations might interact with the choice process. Each of the models we tried fit many of the psychometric trends, but also had trouble with other important aspects of the data. However, these analyses were purely exploratory, and a systematic estimation and model comparison of alternative models is an important topic for future research.

The theory developed here builds on the framework of DDM of binary response selection^{6–19}, and especially on applications of these

models to the realm of perceptual decision making^{18,20–31}. These models assume that noisy information is integrated over time up to a choice threshold. However, the nature of the noise in our task is quite different. In the standard dot motion task the stimulus itself is stochastic, whereas in our task the stimuli are static and the noise is generated internally. In particular, we hypothesize that the brain assigns value to the stimuli by sequentially and stochastically extracting features of the stimuli, retrieving the learnt values for such features, and then integrating those values. In the absence of fixations the computational problem has similar properties, and under appropriate assumptions, it can be shown that the DDM implements an optimal decision-making process that amounts to a sequential-likelihood ratio test^{18,23,24,34–36}. However, the model that we investigated here does not seem to have that property because of the integration bias for the fixated item. An important question for future research is to determine the extent to which the model approximates an optimal Bayesian decision-making problem in which fixations are determined endogenously.

Our model is also related to the models of decision field theory^{13,37–40}, which also consider sequential integration models in the spirit of the DDM in which fixations matter. There are several differences between these studies and our study. First, decision field theory assumes that items are multidimensional and that fixations matter by focusing the integration of value to a subset of dimensions. In contrast, we focused on choices among unidimensional stimuli where fixations matter because they bias the integration of value in favor of one of the items. Second, the predictions of decision field theory regarding the impact of fixations on choice have not been tested directly using eye-tracking.

One essential question is how the brain implements this model of decision-making. One brain region that is likely to be important is the medial orbital frontal cortex. A number of studies have shown that the medial orbital frontal cortex encodes value signals at the time of choice^{41–47}, and these are the likely inputs to the comparator process studied here. We conjecture that fixations affect this process by amplifying the relative value signal for the fixated item in the medial orbital frontal cortex.

Our results have important implications for the quality of choice processes and decision-making in general. First, as fixations might be affected by visual features of the items that are uncorrelated with value, such as their visual saliency⁴⁸ or their location, the model predicts that such irrelevant factors could affect choices. Second, the model predicts that systematic biases in fixations could lead to deficits in decision-making. Extensions of this framework might help us to understand why individuals with autism who generally avoid eye contact show deficits in social decision making⁴⁹. Finally, the model explains how cultural norms (for example, reading left to right) can interact with comparator processes to produce cultural choice biases. These biases help to explain, for example, why shelf and computer screen space on the top-left is more valuable than other positions.

METHODS

Methods and any associated references are available in the online version of the paper at <http://www.nature.com/natureneuroscience/>.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank E. Johnson, P. Bossaerts and C. Koch for comments and J. Pulst-Korenberg for help with data collection. This work received financial support from the Moore Foundation.

AUTHOR CONTRIBUTIONS

A.R. and C.A. devised the experiment. I.K. programmed and conducted the experiment, performed the analyses and co-wrote the manuscript. A.R. designed the model, co-wrote the manuscript and supervised the project.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience/>.

Reprints and permissions information is available online at <http://www.nature.com/reprintsandpermissions/>.

- Rangel, A., Camerer, C. & Montague, P.R. A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* **9**, 545–556 (2008).
- Wallis, J.D. Orbitofrontal cortex and its contribution to decision-making. *Annu. Rev. Neurosci.* **30**, 31–56 (2007).
- Padoa-Schioppa, C. & Assad, J.A. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).
- Kahneman, D. & Tversky, A. Prospect theory: an analysis of decision under risk. *Econometrica* **4**, 263–291 (1979).
- Mas-Colell, A., Whinston, M. & Green, J. *Microeconomic Theory* (Oxford University Press, New York, 1995).
- Luce, R.D. *Response Times: Their Role in Inferring Elementary Mental Organization* (Oxford University Press, Oxford, 1986).
- Stone, M. Models for choice-reaction time. *Psychometrika* **25**, 251–260 (1960).
- Ratcliff, R. A theory of memory retrieval. *Psychol. Rev.* **85**, 59–108 (1978).
- Ratcliff, R., Cherian, A. & Segraves, M. A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J. Neurophysiol.* **90**, 1392–1407 (2003).
- Ratcliff, R. & Smith, P. A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* **111**, 333–367 (2004).
- Laming, D. A critical comparison of two random-walk models for choice reaction time. *Acta Psychol. (Amst.)* **43**, 431–453 (1979).
- Link, S.W. *The Wave Theory of Difference and Similarity* (Lawrence Erlbaum, Hillsdale, New Jersey, 1992).
- Usher, M. & McClelland, J. The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592 (2001).
- Smith, P. Psychophysically-principled models of visual simple reaction time. *Psychol. Rev.* **102**, 567–593 (1995).
- Smith, P. Stochastic dynamic models of response time and accuracy: a foundational primer. *J. Math. Psychol.* **44**, 408–463 (2000).
- Ditterich, J. Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw.* **19**, 981–1012 (2006).
- Bogacz, R. Optimal decision-making theories: linking neurobiology with behaviour. *Trends Cogn. Sci.* **11**, 118–125 (2007).
- Gold, J.I. & Shadlen, M.N. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* **5**, 10–16 (2001).
- Gold, J.I. & Shadlen, M.N. Banburisms and the brain: decoding the relationship between sensory stimuli, decisions and reward. *Neuron* **36**, 299–308 (2002).
- Churchland, A.K., Kiani, R. & Shadlen, M.N. Decision-making with multiple alternatives. *Nat. Neurosci.* **11**, 693–702 (2008).
- Ditterich, J. Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw.* **19**, 981–1012 (2006).
- Ditterich, J., Mazurek, M.E. & Shadlen, M.N. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat. Neurosci.* **6**, 891–898 (2003).
- Gold, J.I. & Shadlen, M.N. Banburisms and the brain: decoding the relationship between sensory stimuli, decisions and reward. *Neuron* **36**, 299–308 (2002).
- Gold, J.I. & Shadlen, M.N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
- Hanks, T.D., Ditterich, J. & Shadlen, M.N. Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat. Neurosci.* **9**, 682–689 (2006).
- Huk, A.C. & Shadlen, M.N. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* **25**, 10420–10436 (2005).
- Mazurek, M.E., Roitman, J.D., Ditterich, J. & Shadlen, M.N. A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13**, 1257–1269 (2003).
- Palmer, J., Huk, A.C. & Shadlen, M.N. The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* **5**, 376–404 (2005).
- Roitman, J.D. & Shadlen, M.N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489 (2002).
- Shadlen, M.N. & Newsome, W.T. Motion perception: seeing and deciding. *Proc. Natl. Acad. Sci. USA* **93**, 628–633 (1996).
- Shadlen, M.N. & Newsome, W.T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936 (2001).
- Armell, K.C., Beaumel, A. & Rangel, A. Biasing simple choices by manipulating relative visual attention. *Judgm. Decis. Mak.* **3**, 396–403 (2008).
- Shimojo, S., Simion, C., Shimojo, E. & Sheier, C. Gaze bias both reflects and influences preference. *Nat. Neurosci.* **6**, 1317–1322 (2003).

34. Bogacz, R. Optimal decision-making theories: linking neurobiology with behaviour. *Trends Cogn. Sci.* **11**, 118–125 (2007).
35. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J.D. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).
36. Reddi, B.A.J. & Carpenter, R.H.S. The influence of urgency on decision time. *Nat. Neurosci.* **3**, 827–830 (2000).
37. Roe, R.M., Busemeyer, J. & Townsend, J.T. Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol. Rev.* **108**, 370–392 (2001).
38. Busemeyer, J. & Townsend, J.T. Decision field theory: a dynamic-cognitive approach to decision-making in an uncertain environment. *Psychol. Rev.* **100**, 432–459 (1993).
39. Busemeyer, J. & Johnson, J.G. Computational models of decision making. in *Blackwell Handbook of Judgment and Decision Making* (eds. Koehler, D. & Narvey, N.) 133–154 (Blackwell, Oxford, 2004).
40. McClelland, J. & Rumelhart, D. An interactive activation model of context effects in letter perception. I. An account of basic findings. *Psychol. Rev.* **88**, 375–407 (1981).
41. Plassmann, H., O'Doherty, J. & Rangel, A. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J. Neurosci.* **27**, 9984–9988 (2007).
42. Hare, T.A., O'Doherty, J., Camerer, C.F., Schultz, W. & Rangel, A. Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *J. Neurosci.* **28**, 5623–5630 (2008).
43. Erk, S., Spitzer, M., Wunderlich, A., Galley, L. & Walter, H. Cultural objects modulate reward circuitry. *Neuroreport* **13**, 2499–2503 (2002).
44. Arana, F.S. *et al.* Dissociable contributions of the human amygdala and orbitofrontal cortex to incentive motivation and goal selection. *J. Neurosci.* **23**, 9632–9638 (2003).
45. Paulus, M.P. & Frank, L.R. Ventromedial prefrontal cortex activation is critical for preference judgments. *Neuroreport* **14**, 1311–1315 (2003).
46. Valentin, V.V., Dickinson, A. & O'Doherty, J. Determining the neural substrates of goal-directed learning in the human brain. *J. Neurosci.* **27**, 4019–4026 (2007).
47. Hare, T.A., Camerer, C.F. & Rangel, A. Self-control in decision-making involves modulation of the vmPFC valuation system. *Science* **324**, 646–648 (2009).
48. Itti, L. & Koch, C. Computational modelling of visual attention. *Nat. Rev. Neurosci.* **2**, 194–203 (2001).
49. Spezio, M.L., Adolphs, R., Hurley, R.S.E. & Piven, J. Abnormal use of facial information in high-functioning autism. *J. Autism Dev. Disord.* **37**, 929–939 (2007).

ONLINE METHODS

Subjects. The experiment consisted of 39 Caltech students. Only subjects who self-reported regularly eating the snack foods (for example, potato chips and candy bars) used in the experiment and not being on a diet were allowed to participate. These steps were taken to ensure that the food items we used would be motivationally relevant. This would not have been the case if the subjects did not like junk food. Subjects were paid a \$20 show-up fee, in addition to receiving one food item. Caltech's Human Subjects Internal Review Board approved the experiment. Written informed consent was obtained from all participants.

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 the items were available. Furthermore, subjects briefly saw all the items at this point so that they could effectively use the rating scale.

In the choice phase, subjects made their choices by pressing the left or right arrow keys on the keyboard. The choice screen had a free response time. Food items that received a negative rating in the rating phase of the experiment 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 chosen pseudo-randomly according to the following rules: (i) no item was used in more than 6 trials; (ii) the difference in liking ratings between the two items was constrained to be 5 or less; (iii) if at some point in the experiment (i) and (ii) could no longer both be satisfied, then the difference in allowable liking ratings was expanded to 7, but these trials occurred for only 5 subjects and so were discarded from the analyses. The spatial location of the items was randomized.

After subjects indicated their choice, a yellow box was drawn around the chosen item (with the other item still on-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. Choice trials with no item fixations for more than 40 ms at the beginning or end of the trial were excluded from analysis. The mean (s.e.m.) number of trials dropped per subject was 2.8 ± 1.5 . For all measurements following the first item fixation and preceding the last item fixation of the trial, blank fixations were dealt with according to the following rules.

If the blank fixations were recorded between fixations on the same item, then those blank fixations were changed to that item. So, for example, a fixation pattern of 'Left', 'Blank', 'Left' would become 'Left', 'Left', 'Left'. The assumption here is that the eye-tracker simply lost the subject's eyes during this time. The alternate hypothesis is that the subject looked away from the item without looking at the other item, but we consider this to be an unlikely scenario.

If blank fixations were recorded between fixations on different items, then those blank fixations were recorded as non-item fixations and discarded from further analysis. The assumption here is that the subject took time to shift his gaze from one item to the other, and during that time was not fixating on either item.

Group model fitting. The computational model was fit to the choice and reaction time data from the even numbered trials of the pooled data from the 39 subjects. The MLE procedure was implemented as follows. First, we set apart the even trials of the data to estimate the model. Then for each set of parameters and pair of liking ratings in the data we ran 1,000 simulations of the model. In the simulations

we randomly sampled fixation times from the empirical distribution conditional on the measure of choice difficulty given by $r_{\text{best}} - r_{\text{worst}}$. First fixations were sampled separately from the rest. We also used the empirical fact that subjects looked left first 74% of the time and that the first fixations were independent of value. Finally, the simulations assume instantaneous transitions between fixations, while in the data there are often delays between fixations. To compensate, we calculated the total amount of 'transition' time in each trial, randomly sampled from the empirical distribution of those 'transition' times, and added them to the simulated reaction times.

Second, we computed the probability of each data point for each set of parameters as follows. The empirical spread of reaction times ranged from 525 ms to 25 s so in the fitting procedure we discarded any simulation trials below 500 ms or above 25 s. The rest of the reaction times were separated into bins from 500–6,000 ms, each one spanning 100 ms, except for the first bin, which went from 500–1,000 ms, and the final bin, which went from 7,000–25,000 ms. For each combination of liking ratings, we then split the data into the trials where Left was chosen and where Right was chosen, and then for each group we counted the number of data trials in each reaction time bin. For the simulations, we similarly calculated the probability that a simulation trial occurs in each reaction time bin, conditional on Left or Right being chosen.

Third, we computed the set of parameters that maximized the log-likelihood of the data by taking the logarithms of each of these probabilities, multiplying by the number of data trials in each bin, and summing them up. The resulting number is used to assess how well the model fit the data, with less negative numbers indicating better fits.

In the simulations, we vary σ as a function of the slope d , rather than absolutely. Therefore, we let $\sigma = d^* \mu$ and perform a grid search over values of d , μ and θ . The search for the maximum likelihood parameters was carried out in three steps. First we did a coarse grid search with d in {0.0001, 0.00015, 0.0002, 0.00025}, μ in {80, 100, 120, 140} and θ in {0.3, 0.5, 0.7, 0.9}. Second, we used the results from the first search to define a finer search with d in {0.000175, 0.0002, 0.000225}, μ in {90, 100, 110} and θ in {0.2, 0.3, 0.4}.

Group likelihood ratio tests. We tested whether θ was significantly different from 1 and 0 by performing likelihood ratio tests. These tests use the results from the MLE described above, as well as those from another MLE model in which θ was fixed to 0 or 1. This procedure was carried out exactly the same as before, using only the even-numbered trials and starting with a coarse search with d in {0.0001, 0.00015, 0.0002, 0.00025} and μ in {80, 100, 120, 140}, followed by a finer search with d in {0.000175, 0.0002, 0.000225} and μ in {90, 100, 110}. The best fitting set of parameters in both cases was $d = 0.0002 \text{ ms}^{-1}$ and $\mu = 100$, with a log-likelihood value of -3,708 for $\theta = 1$ and -3,710 for $\theta = 0$. From these log-likelihood values we calculated the likelihood ratio statistic, which for the case of $\theta = 1$ is given by $\text{LR} = 2(x(\theta = 0.3) - x(\theta = 1))$. Here, x is the log-likelihood value for each set of parameters. This test statistic is distributed as $\chi^2(1)$.

Group simulations. We carried out 1,000 simulations for every combination of values in the dataset using the maximum likelihood parameter estimates and sampling fixation durations from the odd-numbered trials.

Individual model fitting. An important concern with the group fits above is that they do not provide a good description of the underlying distribution of parameters in the subject population. We investigated this issue with two further analyses. First (**Supplementary Fig. 11**), we set $d = 0.0002 \text{ ms}^{-1}$ and $\mu = 100$ from the group level analysis and performed an MLE grid search over θ in {0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1} using all trials and 1,000 simulations for every combination of values.

Second (**Supplementary Fig. 12**), we calculated the average difference in left-choice probability between last-fixation-left trials and last-fixation-right trials, for each subject (curves in **Fig. 5a**). Subjects with $\theta = 1$ should show no difference between these two curves, whereas subjects with $\theta = 0.3$ and $\theta = 0$ should show differences of 0.47 and 0.58, respectively (assuming $d = 0.0002 \text{ ms}^{-1}$ and $\mu = 100$).

Goodness-of-fit calculations. For **Figures 2b,c, 4c** and **5c,e** we could not compute χ^2 goodness-of-fit statistics because the dependent variables are not binary. R^2 statistics were also uninformative because of the high variability in



average fixation duration from subject to subject. Therefore, we devised a different goodness-of-fit statistic, which works as follows. For each value of the independent variable we 'correct' the dependent variable by subtracting the average simulated value from each subject's average value. We then run a weighted least-squares (WLS) regression with the 'corrected' dependent variable. The weights in the regression were equal to the inverse of the variance. Note that if the simulations fit the data well, then the 'corrected' data should be a flat line at 0. On the other hand, if the simulation fits poorly, then the WLS coefficient should be nonzero. So, for goodness-of-fits, we report the P values for the coefficients of these WLS regressions. If the P values are less than 0.05 then we reject that the model fits the data.

Fitting the fixation distributions. To determine the best-fitting distributions for the first and middle fixation durations, we used a log-likelihood method to fit several different types of distribution to all the trials, as well as dividing trials by the absolute difference in the liking ratings. **Supplementary Table 1** summarizes the best-fitting parameters from log-normal distributions (which were consistently the best or near-best distribution) and the log-likelihoods for the different distributions. **Supplementary Figures 16–22** show the log-normal fits to the data.

Mixed effect regressions. All mixed effect regressions had random effects for subject-specific constants and slopes. Only one regression, with three independent variables, was performed for **Figure 3** (and **Supplementary Figs. 13–15**).

Alternative model simulations. The three alternative models (**Fig. 6** and **Supplementary Figs. 23–30**) were each simulated using 500 runs for every combination of values in the dataset, and fixation durations and reaction times were sampled (where appropriate) from the odd-numbered trials.

For the first alternative model, the probability of the fixation ending at each time point was given by $p = \frac{P^*}{3 - k}$, where k is the magnitude of the distance

between the RDV and the choice barrier for the currently fixated item. A value of $P^* = 0.02$ was used as the benchmark value for this model. Two additional models were simulated for **Supplementary Figure 24** with $P^* = 0.002$ and 0.0005.

The second alternative model is described in the text. For the third alternative model, when a fixation was made to an item, the magnitude of that item's choice barrier was lowered to 0.8. These values were chosen to roughly fit the patterns seen in **Figure 2**. Two additional models were simulated for **Supplementary Figure 30** with the barriers lowered to 0.5 and 0.2.

Erratum: Visual fixations and the computation and comparison of value in simple choice

Ian Krajbich, Carrie Armel & Antonio Rangel

Nat. Neurosci. 13, 1292–1298 (2010); published online 12 September 2010; corrected after print 10 February 2011

In the version of this article initially published, there were symbols dropped from the equations in the second paragraph of the results section. The term θ_{right} should have been θr_{right} in the first equation and the term θ_{left} should have been θr_{left} in the second equation. The error has been corrected in the HTML and PDF versions of the article.