Title

Mechanisms underlying the influence of saliency on value-based decisions.

Journal

Journal of Vision, 13(12)

Authors

Mihalas, Stefan

Niebur, Ernst

Stuphorn, Veit

et al.

Publication Date

2013-10-28

DOI

10.1167/13.12.18

Peer reviewed

Mechanisms underlying the influence of saliency on value-based decisions

Xiaomo Chen
Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine and Zanvyl Krieger Mind/Brain Institute, Baltimore, MD, USA
Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD, USA

Stefan Mihalas
Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine and Zanvyl Krieger Mind/Brain Institute, Baltimore, MD, USA

Ernst Niebur
Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine and Zanvyl Krieger Mind/Brain Institute, Baltimore, MD, USA

Veit Stuphorn
Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine and Zanvyl Krieger Mind/Brain Institute, Baltimore, MD, USA
Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD, USA

Objects in the environment differ in their low-level perceptual properties (e.g., how easily a fruit can be recognized) as well as in their subjective value (how tasty it is). We studied the influence of visual salience on value-based decisions using a two alternative forced choice task, in which human subjects rapidly chose items from a visual display. All targets were equally easy to detect. Nevertheless, both value and salience strongly affected choices made and reaction times. We analyzed the neuronal mechanisms underlying these behavioral effects using stochastic accumulator models, allowing us to characterize not only the averages of reaction times but their full distributions. Independent models without interaction between the possible choices failed to reproduce the observed choice behavior, while models with mutual inhibition between alternative choices produced much better results. Mutual inhibition thus is an important feature of the decision mechanism. Value influenced the amount of accumulation in all models. In contrast, increased salience could either lead to an earlier start (onset model) or to a higher rate (speed model) of accumulation. Both models explained the data from the choice trials equally well. However, salience also affected reaction times in no-choice trials in which only one item was present, as well as error trials. Only the onset model could explain the observed reaction time distributions of error trials and no-choice trials. In contrast, the speed model could not, irrespective of whether the rate increase resulted from more frequent accumulated quanta or from larger quanta. Visual salience thus likely provides an advantage in the onset, not in the processing speed, of value-based decision making.

Introduction

Value-based decision making is the selection of an action among several alternatives based on the subjective value of their outcomes. While ideally this choice should be independent of irrelevant target properties, it is well known that low-level physical properties can profoundly influence decision making. During free viewing of natural scenes and video sequences, saccades are drawn to more salient parts of an image (Berg, Boehnke, Marino, Munoz, & Itti, 2009; Parkhurst, Law, & Niebur, 2002; Parkhurst &

Citation: Chen, X., Mihalas, S., Niebur, E., & Stuphorn, V. (2013). Mechanisms underlying the influence of saliency on value-based decisions. *Journal of Vision*, *13*(12):18, 1–23, http://www.journalofvision.org/content/13/12/18, doi:10.1167/13.12.18.

doi: 10.1167/13.12.18 Received February 27, 2013; published October 28, 2013 ISSN 1534-7362

Niebur, 2003), observers find these parts more interesting (Masciocchi, Mihalas, Parkhurst, & Niebur, 2009), and high-salience targets are detected faster and more accurately (Egeth & Yantis, 1997; Wolfe, 1998). The question thus arises whether visual salience influences not only simple perceptual but also value-based decisions.

A number of recent studies demonstrated that both visual salience and subjective value can affect decision making (Markowitz, Shewcraft, Wong, & Pesaran, 2011; Navalpakkam, Koch, Rangel, & Perona, 2010; Schutz, Trommershauser, & Gegenfurtner, 2012). Nevertheless, the mechanisms underlying this behavioral phenomenon might differ depending on the specific influence of salience in the task. Sensory stimuli can vary in many different feature dimensions and any of these feature domains can influence the overall salience of the target. However, value information is typically carried only by some of the features of a visual target. It is therefore of importance whether salience is manipulated on the same or a different feature dimension as value.

In a situation in which salience is manipulated on a different feature dimension than the one indicating value, the main effect of salience manipulations will be to influence the overall contrast of the target relative to the background and other targets. In other words, low salience will lead to a lower probability that the target will be detected to be present. However, once detected a low-salience target will provide as much information about its value as a high-salience target. Such salience manipulations were often employed in previous research, either by modulating detectability of targets (Markowitz et al., 2011) or of distractors (Navalpakkam et al., 2010). This generates a "neglect" situation, in which high-value targets can be overlooked if they are of lower salience than the background or alternatives.

The situation is different when the salience of the feature dimension is manipulated that carries value information, but other feature dimensions of the target are still highly salient. In this situation, the perception of the value information is selectively influenced by the salience manipulation, while all the other perceptual dimensions are the same. The influence of the salience manipulation on choice behavior is therefore not simply to make the agent unaware of a low-salience target, but rather to create targets whose value is harder to perceive. There are fewer studies of this type (Schutz et al., 2012) and as a result we know much less about the influence of salience on value-related information.

In this study, we designed therefore a two alternative forced choice task, in which items in a visual display were endowed with different values (rewards). Human participants rapidly chose between items, attempting to maximize the reward amount. We manipulated visual salience and value of the targets simultaneously and independently across trials, while keeping the detectability of all targets constant. That is, we designed visual stimuli for which one visual feature (luminance contrast with the background) was large enough to ensure that their location could be detected rapidly. Another visual feature (color saturation) was manipulated so that the visual feature carrying value information (color hue) was more or less perceptually salient. Therefore, the manipulation of salience influences the perception of value. The salience is with respect to behaviorally relevant information (i.e., target value), but not with respect to the ability of the subject to localize the targets on the screen. In this way we could manipulate the visual salience of value information without directly affecting motor processes used to report the choice. We also mixed no-choice trials in with the choice trials, in which subjects could only select the single target on the screen. The no-choice trials were controls to test the effect of visual salience and value on behavior without any interference by the choice process.

We found that both value and salience have strong effects on the decision process by themselves, as has congruency between value and salience. Specifically, reaction times were significantly correlated with both value and salience of the chosen target and with value difference and salience difference between the chosen target and the non-chosen target. In addition, the error rate, defined as the rate of choosing the lower valued target, was significantly higher in incongruent trials than in any other type of trials.

After characterizing behavior in a descriptive regression model, we analyzed the neuronal mechanisms underlying our behavioral data using a series of four stochastic accumulator models based on different functional assumptions (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Cisek, Puskas, & El-Murr, 2009; Hanks, Mazurek, Kiani, Hopp, & Shadlen, 2011; Krajbich & Rangel, 2011; Purcell et al., 2010). All models consist of two accumulators, each adding up value information in support of one of the two possible choices. For the *independent model*, we assumed that the two accumulators did not interact, while the other accumulator models implemented mutual inhibition between them. The *speed model* assumed that salience influenced the rate with which value information was accumulated by modulating the strength of the incoming value information. The *onset model* was motivated by the observation that salience can reduce visual processing time (Ratcliff & Smith, 2011; White & Munoz, 2011) and assumed an earlier accumulation onset time rather than an increased accumulation rate. The full model combined both ways for salience to affect the accumulation process. Comparison between model predictions and behavior suggested that mutual

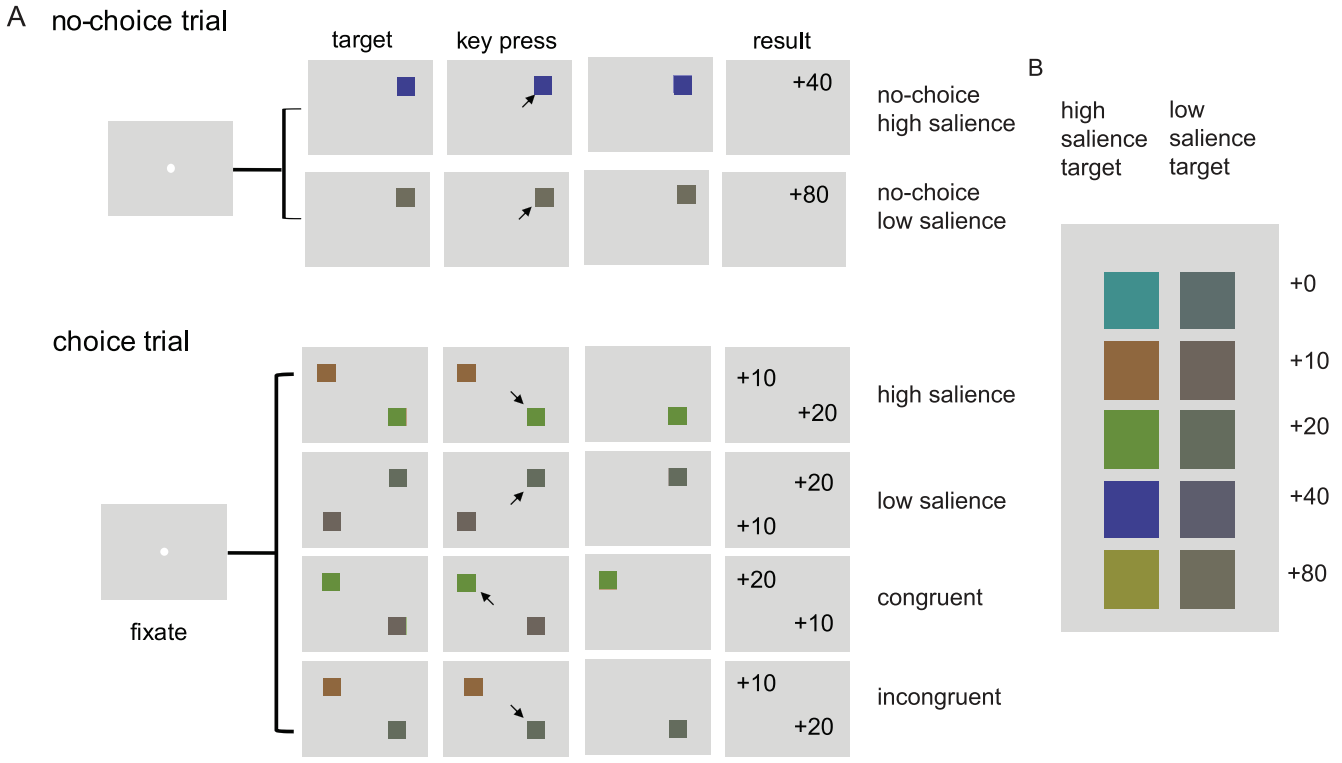

Figure 1. Behavioral paradigm. (A) Sequence of events during no-choice trials (top) and choice trials (bottom). The type of stimulus is listed to the right. The black arrow is not part of the stimuli; it symbolizes the participants' choices that then leads to the next stimulus shown. In all cases shown, the arrow corresponds to the optimal choice. (B) Visual cues used. High- and low-salience targets are left and right, respectively. Rows correspond to values, as shown to the right.

inhibition and salience-induced differences in accumulation onset time, but not accumulation rate, are necessary to explain the behavior of the human participants.

Methods

Subjects

Fifteen participants (age: 18–30, eight female), undergraduate and graduate students naïve to the purpose of the study, were recruited from the Johns Hopkins University community and participated in the experiment after providing informed consent. All participants reported normal or corrected-to-normal vision and no history of color blindness. Among these 15 participants, nine (age: 26–30, five females) participated in the two pilot experiments, nine (age: 18–30, five females) participated in the main experiment, and three participated in all three experiments. All procedures were approved by the Johns Hopkins University Homewood Institutional Review Board.

Pilot experiments

In the first pilot experiment, we determined saturation levels to be used in the main experiment based on simple color detection. Four targets with equal brightness appeared on the screen, of which one was colored and the others were gray. The participants were required to localize the colored target and to indicate its position by pressing the corresponding key on a keyboard. In order to encourage accuracy, we did not set a response deadline. We tested all five colors (cyan, brown, green, blue, yellow) used in the main experiment (Figure 1) in a range (1%–16%) of saturation levels. In the pilot experiment, the color of the targets in the detection experiment did not carry any value information and we did not observe systematic behavioral differences across colors (Figure 2). In contrast, mean reaction time varied systematically with the saturation level, which for our purposes served as a measure for the salience of a target. Low saturation levels were defined as in the 3%–5% range because the reaction times in response to these targets were around 50 ms longer than those in response to high-saturation targets, defined as 16% where performance plateaued. Accuracy was nearly perfect for all values above 6% (with one outlier). Thus, low saturation level was

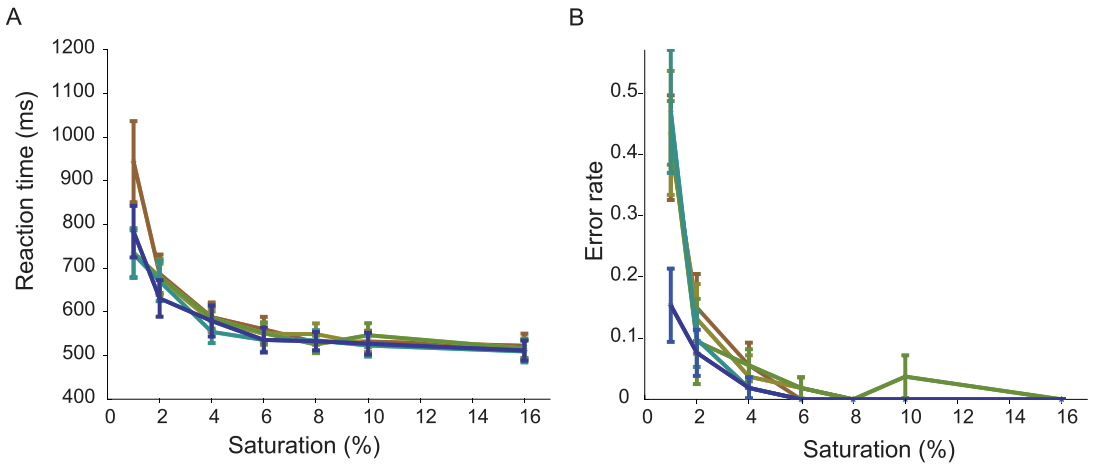

Figure 2. Effect of saturation level on reaction time and error rate for different color targets. The mean reaction times and error rates for nine subjects are plotted. The colors of the lines in the plots correspond to the colors of the targets used in the task. (A) Mean reaction times plotted as a function of the saturation level. Error bars represent standard error of the mean reaction time. (B) Error rates plotted as a function of the saturation level. Error bars represent standard error of error rate.

chosen for each color (5% for yellow, 4% for cyan, brown, and green, and 3% for blue), so that the salience of the color information was substantially reduced from high saturation level (16% for all colors), but still strong enough to be detectable.

In order to test whether the manipulation of the color saturation by itself had an effect on the speed with which the targets could be detected, we performed a second pilot experiment (the singleton task) with nine participants. In this task, 10 targets with different color or saturation as selected in the pilot experiment were used. There was no difference in value associated with the targets. In each trial, only one target appeared on the screen and the participants were asked to indicate its location by pressing the corresponding key. They were encouraged to do so as fast as possible while maintaining accuracy. The task made sure the detectability of all the targets was the same.

Stimuli

In the main experiment, the value associated with a particular target was indicated by its color. The color properties of the targets were derived from the hue ($H \in [0°,360°)$), saturation ($S \in [0,1]$), and brightness ($V \in [0,1]$) color space. In this space, hue is the attribute of a visual sensation according to which an area appear to be one of the perceived colors; saturation is the colorfulness of a stimulus relative to its own brightness; and brightness is the attribute of a visual sensation according to which an area appears to emit more or less light. Five different colors indicated five different values that could be earned ("cyan," hue 180°: zero points; "brown," 30°: 10 points; "green," 90°: 20 points; "blue," 240°: 40 points; "yellow," 60°: 80 points) (Figure 1B). The targets were approximately 1.5° × 1.5° in size and were always presented approximately 20° away from the central fixation point at angles 45°, 135°, 225°, or 315° relative to the horizontal. The background was approximately 16 × 22 cm in size and was uniformly gray and the brightness of each target exceeded that of the background by 4%. Since the targets had all the same brightness, the participants had to rely on color information alone to determine the relative value of each target. We manipulated the salience of this reward-related information by modulating color saturation independently of target value, as determined in the first pilot experiment. As discussed, the selection of high and low saturation levels for each hue was guided by the psychophysical data in the second pilot study.

Main experiment

At the start of the main experiment, participants received instructions on the nature of the task. They were informed that they would have the opportunity to earn "points" that accumulated over trials, and they were encouraged to maximize the number of points earned.

Participants then were presented with targets on a computer monitor (Figure 1A) in front of them that varied in value and salience. The task consisted of two types of trials: choice and no-choice trials. At the beginning of each trial, the participants were required to fixate the fixation point for 1500 ms. In choice trials, after the fixation point disappeared, two targets appeared in diagonally opposite locations on the screen. Participants chose a target by pressing a key on the keypad (Insert: left up; Delete: left down; Home:

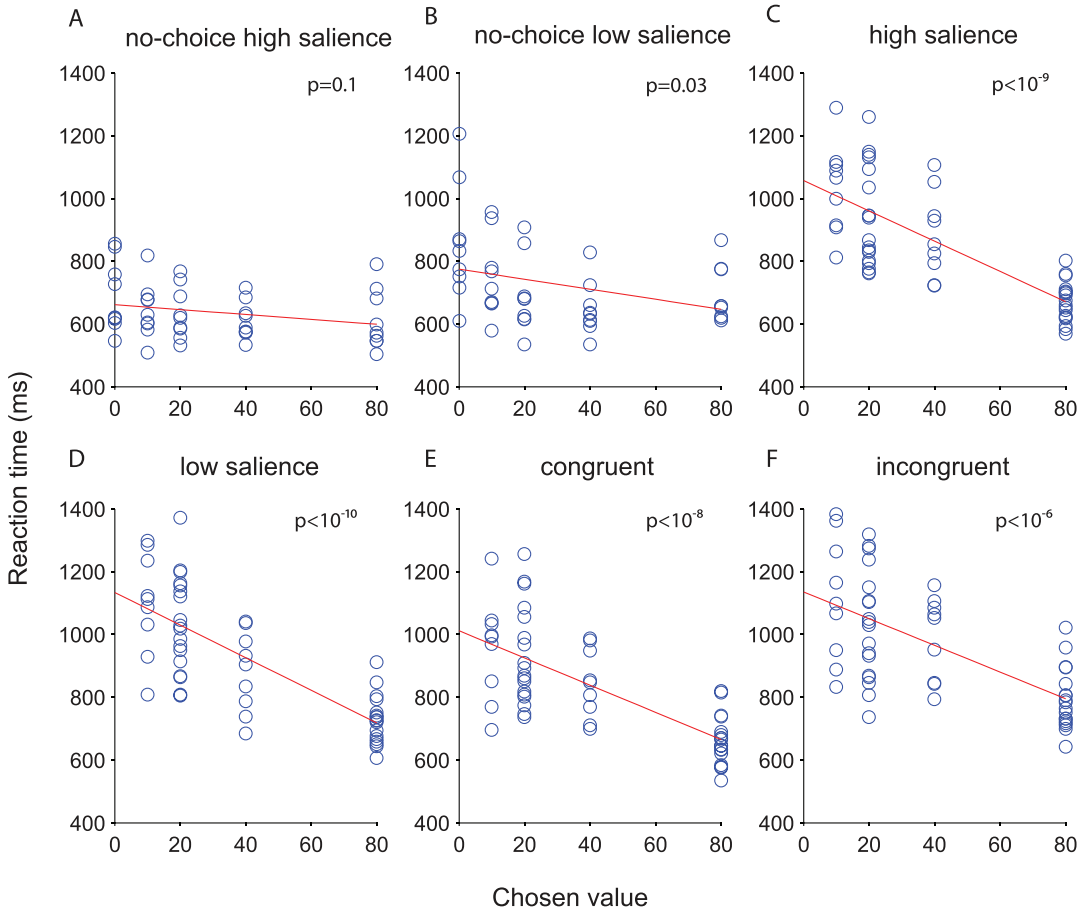

Figure 3. Effect of chosen value on reaction time. The mean reaction times for all participants are plotted against the value of the chosen targets in (A) no-choice high-salience trials, (B) no-choice low-salience trials, (C) high-salience trials, (D) low-salience trials, (E) congruent trials, and (F) incongruent trials. Each circle shows the mean reaction time of one participant in the respective condition. The red line shows the result of linear regression.

right up; End: right down; the relative locations of these keys on the keypad agree with the locations of the corresponding stimuli on the screen) that corresponded to the location of the desired option. Pressing any key other than these four was considered an invalid choice. No-choice trials were used as controls to test the influence of both salience and value without interference by choice. In no-choice trials, only one target appeared on the screen and participants had to press the key that corresponded to the target location. Pressing one of the other three keys was considered an error. The response deadline (2000 ms) was chosen generously to encourage participants to take as much time as necessary to choose the appropriate target. Following a valid key press, the amount of points associated with the chosen target and the non-chosen target were revealed on the monitor. Otherwise, no points were revealed, the next trial started after an inter-trial interval whose length was selected randomly (uniform distribution) from the range 1000–1500 ms.

Six comparisons were selected from the set of possible value differences between the two stimuli: zero versus 10 points, zero versus 20 points, 10 versus 20 points, 20 versus 40 points, 20 versus 80 points, and 40 versus 80 points. Each of these pairs was presented with equal frequency. These comparisons were selected so that across different trial types targets with medium values (10, 20, 40 points) had equal probability to be either larger or smaller in value than the alternative target. In addition, this set included comparisons with smaller and larger value differences. For each value comparison, there were four different salience-value combinations (Figure 1A): In low-salience trials, both targets were of low salience. In high-salience trials, both targets were of high salience. In congruent trials, the high-value target was of high salience, while the low-value target was of low salience. Finally, in incongruent trials, the high-value target was of low salience and the low value was of high salience. This resulted in 24 different combinations of choice trials (6 Value Comparisons $\times$ 4 Value-Salience Combinations). Together with 10 different types of no-choice trials (5 Values $\times$ 2 Salience Levels), there were 34 different trial types that formed one block of trials. The trials were

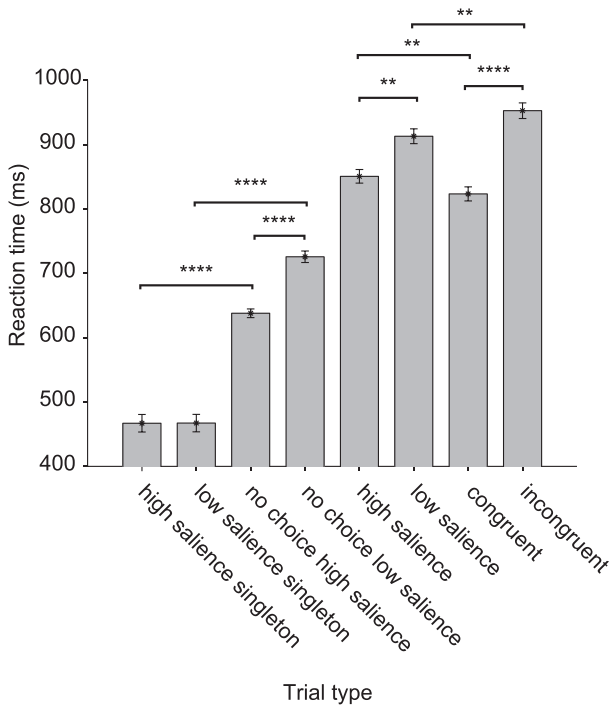

Figure 4. Mean reaction time across all participants on high-salience singleton trials (the singleton task), low-salience singleton trials (the singleton task), no-choice high-salience trials, no-choice low-salience trials, high-salience trials, low-salience trials, congruent trials, and incongruent trials. Asterisks indicate statistical significance of difference between conditions, * means $p \leq 0.05$, **means $p \leq 0.01$, *** means $p \leq 0.001$, ****means $p \leq 0.001$; in all cases from t tests. Sample size is nine. Error bars represent standard error of the mean.

presented in blocks, so that a trial of a particular type was not presented again until the next block. Within a block, the order of trials was randomized.

Participants were initially trained with high salience (32% saturation level) no-choice and choice trials. After they achieved an accuracy of 90%, the main experiment began. Each participant performed the same number of trials, 11 blocks with 34 trials each, 374 trials total.

Results

Influence of value on reaction time

Our behavioral results showed a significant effect of the chosen target value on reaction times in correct trials. The mean reaction time during both no-choice and choice trials reflected the value of the chosen target. For all types of choice trials, reaction time was significantly correlated with the chosen target value; high salience (t test, df = 34, $p < 10^{-9}$), low salience (t test, df = 34, $p < 10^{-10}$), congruent (t test, df = 34, $p < 10^{-8}$), and incongruent (t test, df = 34, $p < 10^{-6}$) (Figure 3). The correlation coefficients were significantly smaller than zero as tested by the t test in all cases; the larger the chosen value, the shorter the reaction time. This result reflected most likely the motivational drive of the chosen target value on the speed of decision making and response execution processes. Interestingly, the relative importance of this motivational drive was weaker on no-choice trials than on choice trials (slope of no-choice trials: high salience: −0.78 ms/point, low salience: −1.6 ms/point; slope of choice trials: high salience: −4.82 ms/point, low salience: −5.19 ms/point, congruent: −4.33 ms/point, incongruent: −4.25 ms/point), and in no-choice trials, the correlation between value and reaction time was significant only for low salience (t test, df = 43, $p = 0.03$), but not high-salience trials (t test, df = 43, $p = 0.1$) (Figure 3).

Not only did the absolute value of the chosen target have a significant effect on reaction time, but also the difference between chosen and non-chosen target. In our experimental design, we used only six out of the full set of all possible value combinations. Within this subset of choices, the value of the chosen target (i.e., the target with the higher value) was positively correlated with the value difference between chosen and non-chosen target. Therefore, we could not use a simple regression analysis to test whether the reaction time were correlated with value differences independent of chosen value. Nevertheless, a partial correlation analysis showed that, when we controlled for the contribution of the chosen value, the reaction time was still significantly correlated with the value difference between the chosen and the non-chosen target in all choice trials (Spearman partial correlations; high-salience trials: df = 42, $p < 10^{-4}$, low-salience trials: df = 42, $p = 0.015$, congruent trials: df = 42, $p = 0.016$, incongruent trials: df = 42, $p = 0.002$). This relationship was also negative (slope high-salience trials: −0.57 ms/point, low-salience trials: −0.33 ms/point, congruent: −0.33 ms/point, incongruent: −0.42 ms/point): the larger the value difference, the shorter the reaction time. This finding supports the view that participants compared the values of both targets before making a choice. Larger value differences made it easier to discriminate the more valuable target and resulted in faster responses, while smaller differences decreased the discriminability and required more time to select the correct response.

Influence of salience on reaction time

The salience of the reward information, i.e., the color saturation level of the targets, had a significant influence on reaction times when compared to the singleton task, the second pilot experiment (Figure 4).

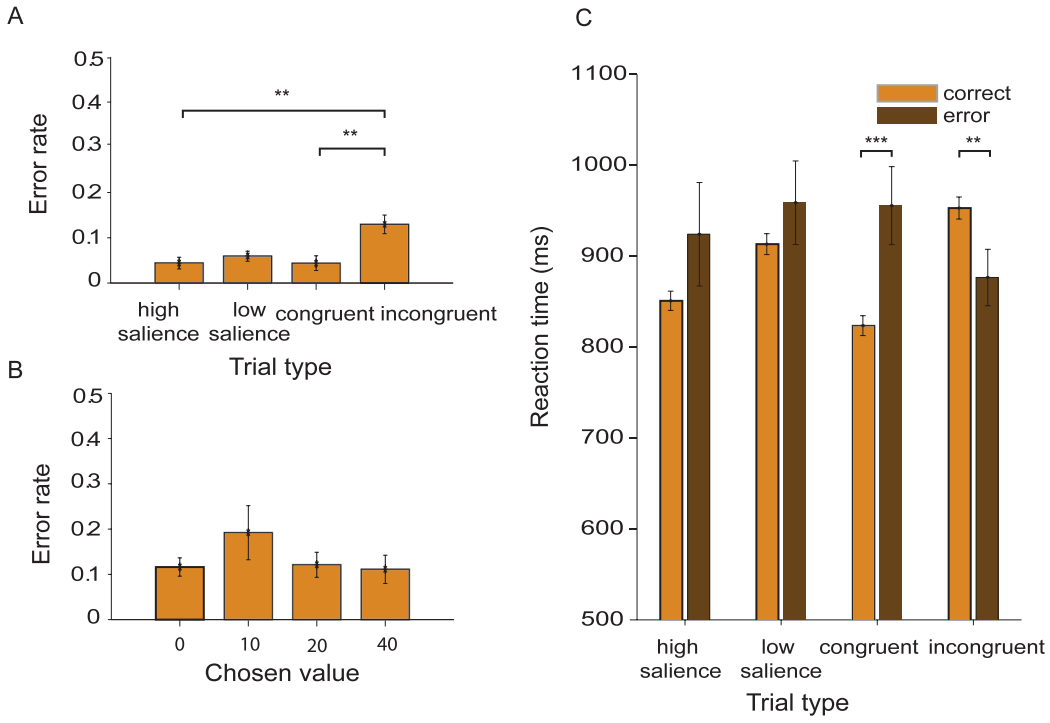

Figure 5. Error rates and mean reaction times on different trial types. (A) Mean error rate for all participants for different trial types. (B) Mean error rate in incongruent trial for all participants as a function of chosen value. (C) Mean reaction time for both correct (light orange) and wrong (dark orange) choice on different trial types. See Figure 4 for the meaning of the asterisks. Sample size is nine. Error bars represent standard error of the mean reaction time.

In the singleton task, the reaction time for high-salience targets (high salience singleton, mean: 467 ms) was identical to that for low-salience targets (low salience singleton, mean: 467 ms) and much faster than the reaction time in no-choice trials. On the other hand, in no-choice trials when there was no interference between salience and value, the reaction time for high-salience trials (mean: 638 ms) was significantly faster (Kolmogorov–Smirnov [K-S] test, $p < 10^{-10}$) than for low-salience trials (mean: 726 ms).

At first sight, the large latency difference between singleton and no-choice trials is surprising, since the only difference between the two trial types is that color is behaviorally meaningful in one (no choice), but not the other (singleton). However, this difference likely reflects a simple speed-accuracy trade-off caused by contextual differences in task demands. In the singleton task, the subjects could be sure that on any given trial there was only one target on the screen. The task was in essence to detect the changing location of the target as fast as possible. For this purpose, luminance contrast provided sufficient information, while target color could be safely ignored. In this situation, the subject's threshold for selecting a target could be lower than in the choice condition without affecting accuracy, due to the absence of distractors. In contrast, during our main experiment the no-choice trials were embedded in a larger number of choice trials. In this situation the subject's threshold for selecting a target had to be higher than in the choice condition, because in the majority of trials there were two competing targets whose value needed to be compared. In principle, there was an absence of distractors in the no-choice trials that was similar to the singleton trials. However, since the two trial types were randomized, the subjects could not be sure when a no-choice would occur and therefore could not adjust their response criteria selectively.

In choice trials, the reaction time in high-salience trials (mean: 851 ms) was significantly faster (K-S test, $p = 0.032$) than in low-salience trials (mean: 913 ms). Moreover, the congruency between differences in value and reward salience had a strong effect on the target selection process. In congruent as well as in incongruent trials, the two targets varied both in value and in salience. In congruent trials, the more valuable target had also more salient reward information. Thus, both the difference in value and in salience supported the same target. In contrast, in incongruent trials the more valuable target had less salient reward information. Here, the difference in value and in salience supported different targets. Accordingly, across all value levels, the reaction time on congruent trials (mean: 822 ms) was significantly faster (K-S test, $p < 10^{-13}$) than on incongruent trials (mean: 953 ms; Figure 4). This difference could not be explained merely by the fact

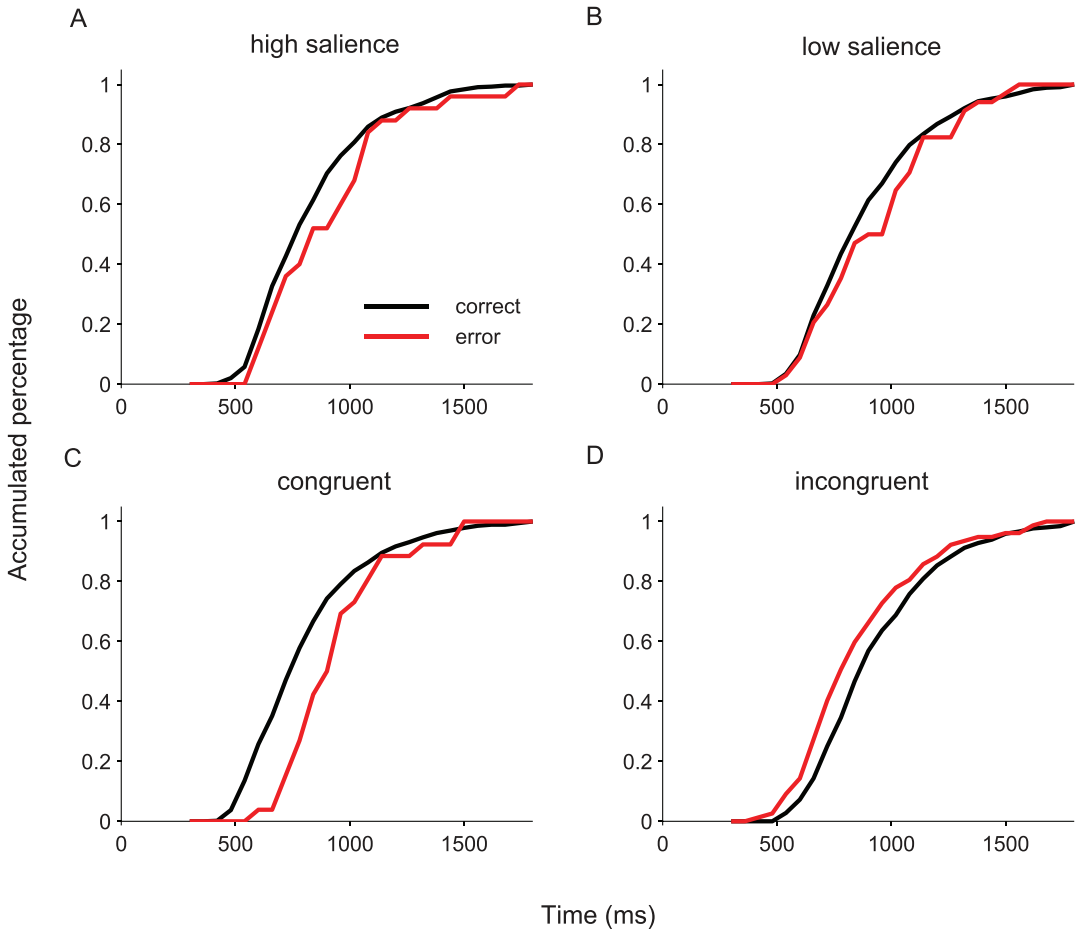

Figure 6. Cumulative reaction time distributions. Cumulative distributions of reaction time are plotted for correct (black) and error (red) trials on high-salience trials (K-S test, $p = 0.28$), low-salience trials (K-S test, $p = 0.25$), congruent trials (K-S test, $p = 0.001$), and incongruent trials (K-S test, $p = 0.01$).

that the chosen targets differed in salience. The reaction time on congruent trials was still significantly faster (K-S test, $p = 0.01$) than that on high-salience trials (Figure 4), even though the salience of the chosen targets were the same. Likewise, the reaction time on incongruent trials was also significantly slower (K-S test, $p < 0.01$) than that on low-salience trials (Figure 4).

Error trials

In the experiment, we did not set a very stringent time limit on the decision process. Therefore, the error rates, defined as the rate of choosing the lower valued target, were low in general. Nevertheless, we also saw an effect of congruency on error rate (Figure 5A). The error rates for high-salience, low-salience, and congruent trials were low (high: 4.3%; low: 5.8%; congruent: 4.4%). Specifically, the error rate for low-salience trials was not significantly higher (K-S test, $p = 0.25$) than the one for high-salience-trials. This indicated that the participants were still able to identify the color of the low-salience targets, although those targets were harder to identify. In contrast, the error rate for incongruent trials was much higher (13.1%). This is higher than the error rates for high salience (K-S test, $p = 0.01$) and congruent trials (K-S test, $p = 0.01$), as well as for low-salience trials although the latter difference was not significant (K-S test, $p = 0.07$).

Furthermore, we compared the reaction time distribution of the error and correct trials during choice trials. To quantify the differences, we plotted the cumulative distribution for both error trials and correct trials in all four trial conditions. As shown in Figures 5C and 6, in low- and high-salience trials, the reaction time (RT) on error trials (mean RT: low-salience trial: 959 ms, high-salience trial: 924 ms) tended to be longer than on correct trials (mean RT: low-salience trial: 913 ms, high-salience trial: 851 ms). Though the differences did not reach significance (K-S test, $p = 0.28$ and $p = 0.25$, respectively), the difference is significant (K-S test, $p = 0.048$) for the combined population. In congruent trials, this difference became larger. The reaction time for error trials (mean RT: 955 ms) was significantly

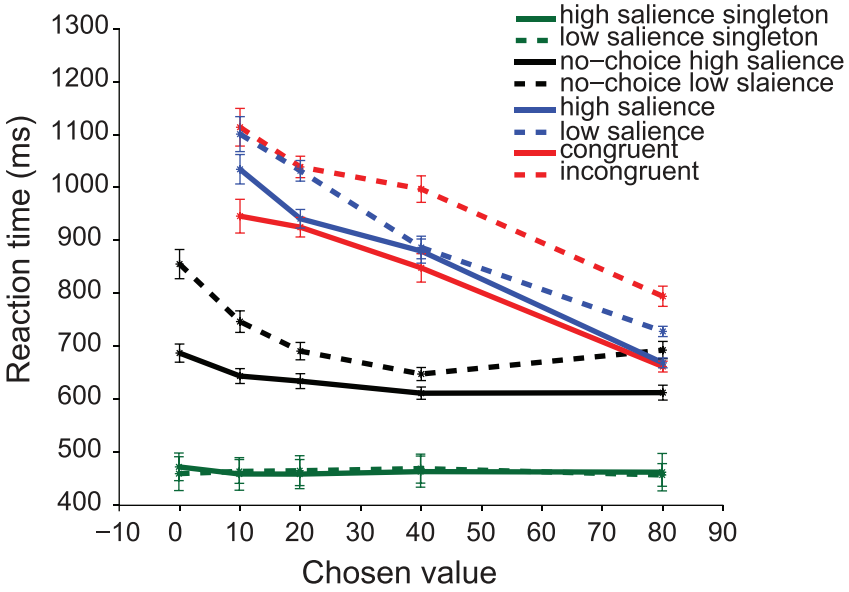

Figure 7. Mean reaction times for both second pilot and main experiment. Mean reaction time in the second pilot experiment are plotted against different targets with corresponding colors without value information on high-salience singleton trials (green solid line), low-salience singleton trials (green dotted line) in the singleton task (second pilot experiment). The reaction time in the main experiment are plotted against value of the chosen targets on no-choice high-salience trials (black solid line), no-choice low-salience trials (black dotted line), high-salience trials (blue solid line), low-salience trials (blue dotted line), congruent trials (red solid line), and incongruent trials (red dotted line). Sample size is nine. Error bars represent standard error of the mean reaction time.

longer (K-S test, $p = 0.001$) than for correct trials (mean RT: 823 ms). In contrast to all other trial types, in incongruent trials the reaction time for error trials (RT: 876 ms) was significantly shorter (K-S test, $p = 0.01$) than for correct trials (mean RT: 953 ms). Note that this shorter reaction time on error trials was not confounded by the chosen value on those trials. First, on error trials (by definition) a smaller value was chosen than on correct trials. Second, the error rate did not increase as the chosen value increased (Figure 5B). The chosen value on error trials was therefore on average not larger than the chosen value on correct trials. This specific difference in the reaction time distributions between congruent and incongruent trials turned out to be important, because, as we shall see below, it allowed us to distinguish between different types of accumulator models of the decision process.

Descriptive model of behavior

To summarize, in the main experiment, the reaction time across the six different trial types was correlated with the value of the chosen target, the salience of the reward information, and the contingency between value and salience (Figure 7). Across all chosen values in the main experiment, the reaction time was shortest for no-choice high-salience trials, increased successively for no-choice low-salience, congruent, high-salience, and low-salience trials, and was longest for incongruent trials. This was very different from the results in the second pilot experiment (singleton task), in which the results showed no difference between reactions to high- and low-salience targets (Figures 4 and 7).

We further used a descriptive model to quantify the trends we had observed in the behavior data. There were a number of factors that could contribute to the decision-making process, including (a) the value of the chosen target, (b) the salience of the chosen target, (c) the value of the non-chosen target, (d) the salience of the non-chosen target, (e) the value difference between chosen and non-chosen target, (f) the salience difference between chosen and non-chosen target, as well as the multiplicative interaction between salience and value for both (g) chosen and (h) non-chosen targets. In order to quantify the effect of each of these possible factors, we fitted a family of nested regression models to the reaction times in correct choice trials that included all possible iterations of the seven factors plus a baseline term. In order to combine the reaction time data across all participants, we normalized reaction times within each participant between zero and one (thus, the normalized RTs computed in Equation A11 cannot be directly compared with the actual RTs in Figure 4). By comparing the Bayesian information criterion value (BIC) and Akaike value (Burnham & Anderson, 2002; Busemeyer & Diederich, 2010) of each model (Table 1), we identified the best fitting model. Of all linear models tested, the lowest BIC value and lowest Akaike value occurred for the same model:

Rank	Variables resulting in model	BIC	AIC	Evidence ratio
1	v_{chosen}, s_{chosen}, dv, ds	−8711	−8734	1
2	v_{chosen}, $s_{non\text{-}chosen}$, ds, $v_{non\text{-}chosen}$	−8711	−8734	1
10	v_{chosen}, s_{chosen}, dv, $v_{non\text{-}chosen}$	−8710	−8733	1.23
13	v_{chosen}, s_{chosen}, dv	−8708	−8726	3.90
17	v_{chosen}, s_{chosen}, dv, ds, $v_{non\text{-}chosen} \times s_{non\text{-}chosen}$	−8705	−8732	23.90

Table 1. BIC table for descriptive behavior regression model. *Notes*: All possible combinations of eight independent variables: v_{chosen}, s_{chosen}, $v_{nonchosen}$, $s_{nonchosen}$, $v_{chosen} - v_{chosen}(dv)$, $s_{chosen} - s_{nonchosen}(ds)$, interaction between chosen value and salience ($v_{chosen} \times s_{chosen}$), interaction between non-chosen value and non-chosen salience ($v_{nonchosen} \times s_{nonchosen}$) were tested against the behavioral data. The first column shows the BIC rank of the model among all other 255 models. The second column shows the selected variables with the corresponding rank. The third column shows the BIC value, the forth column shows the AIC value, and the fifth column shows the BIC evidence ratio for each model compared with the best fitting model.

$$\begin{aligned} RT_{normalized} = {} & 0.5209 - 0.0015{*}v_{chosen} \\ & - 0.0373{*}s_{chosen} \\ & - 0.0018{*}(v_{chosen} - v_{nonchosen}) \\ & - 0.0188{*}(s_{chosen} - s_{nonchosen}), \end{aligned} \quad (1)$$

where v_{chosen} and $v_{nonchosen}$ are the point values of the chosen and non-chosen targets ($v_i \in (0, 0.1, 0.2, 0.4, 0.8)$), and S_{chosen} and $S_{nonchosen}$ are the salience values of the chosen and non-chosen targets ($S_i \in (0, 1)$), respectively. All four parameters (but none of the other four possibilities listed above) contributed significantly to the regression, including: (a) value of the chosen target (t test: $p < 10^{-7}$), (b) salience of the chosen target (t test: $p < 10^{-5}$), (c) value difference between chosen and non-chosen target (t test: $p < 10^{-4}$), and (d) salience difference between chosen and non-chosen target (t test: $p = 0.001$). Table 1 shows the BIC values, Akaike value, and evidence ratio (relative to the best-fitting model) for different models ranked by their fit to the behavioral data. From the evidence ratios it is clear that there were actually approximately 12 different regression models all containing four variables that all fitted the data almost as well as the best-fitting model. This phenomenon is likely related to the fact that the variables we chose were most likely not completely independent of each other such as the equation containing S_{chosen} and $S_{nonchosen}$ can be equally expressed as an equation containing S_{chosen} and dS. However, there was a clear drop in evidence for alternative three- or five-variable models.

Accumulator models

The behavioral results, confirmed by a linear regression model, showed that both value and salience of the targets as well as the congruency between them influence the decision-making process. To make progress towards understanding the underlying mechanisms, we modeled the decision process using accumulator models with four functionally related architectures (Figure 8). In addition to suggesting a functional mechanistic explanation of the underlying mechanisms, accumulator models have the additional advantage over descriptive regression models (like the one developed in the previous section) that they describe the entire distribution of behavioral data, rather than only their mean values. This modeling approach allowed us to address several questions beyond the identification of the behaviorally relevant factors. Most importantly, it allowed us to ask questions regarding the functional architecture of the decision-making mechanism that implements the value-based decisions. In the following simulation-based analysis, we focused specifically on two of these mechanistic questions. First, is there a role for inhibitory interactions between the processes representing the two targets? Second, how does the salience of the reward information influence the decision-making process? In addition, accumulator models incorporated in a natural fashion nonlinearity in the decision-making process, such as the threshold, which is not easy to be captured in a linear regression model.

Models 1–3 have the same complexity (six parameters). As described below, they are special cases of the general functional Model 4 that is slightly more complex with seven parameters. In order to determine the importance of particular factors, for Model 1–3, we systematically constrained one factor of the general model (Model 4), while allowing the other factors to change freely to achieve the best possible fit with the behavioral data. All other aspects of the functional architecture were held constant across the four different variants (see Appendix).

All models have two accumulator units, each of which integrates the input from the target whose choice they would trigger until it reaches the response threshold. The quanta size of the input for accumulation is determined by the target value. The rationale for this design choice follows immediately from the idea that what is accumulated during the decision process is support for a particular choice (here the value that is associated with selecting a particular target). In Model

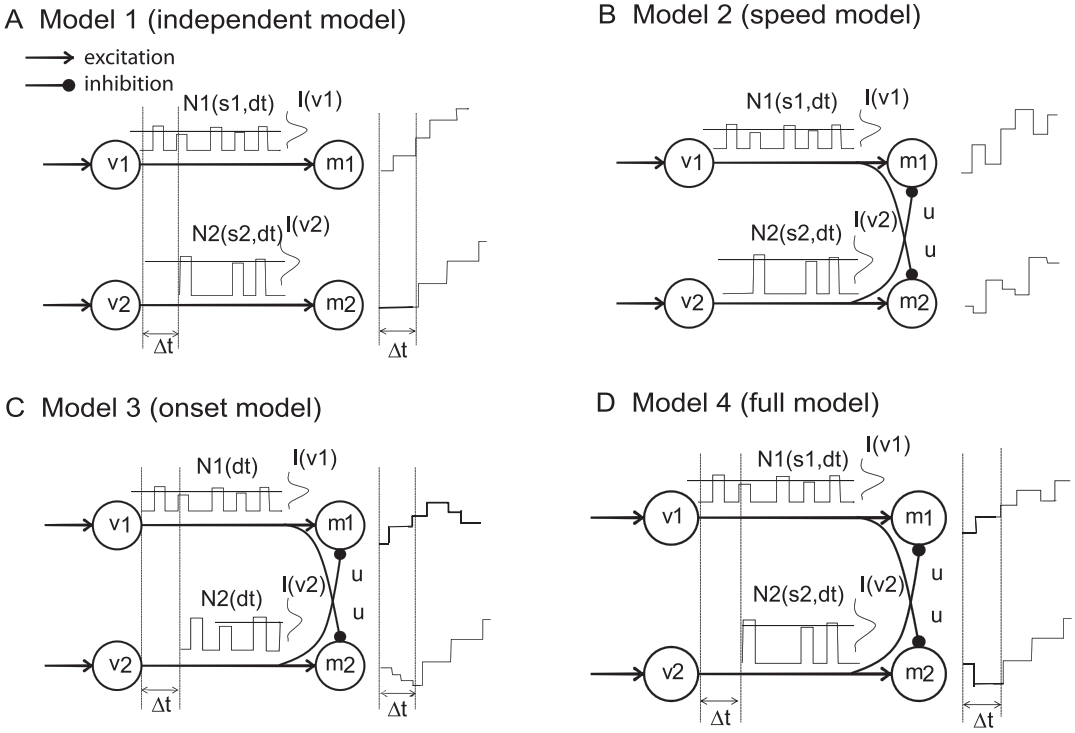

Figure 8. Architecture of the four accumulator models. (A) Model 1: independent model, without mutual inhibition between the targets. (B) Model 2: speed model. Salience influences the rate of accumulation but not its onset. (C) Model 3: onset model. Salience influences the onset of accumulation but not its rate. (D) Model 4: full model with feed-forward inhibition model salience influencing both the onset of accumulation and its rate. v_i are the units that transfer sensory input into value. m_i are the accumulators that integrate the input and trigger a motor response. Consistent with Equation A2, $\Delta t = t_1 - t_2$ is the onset difference generated by salience differences, N_i is the number of accumulations that occur in each accumulator mi, $I(v_i)$ is the rate of accumulation for each accumulator m_i, and u is the mutual inhibition parameter between two accumulators.

1 (independent model), salience influences both the onset and the rate of accumulation but the two accumulators are independent, with no inhibitory interaction between them (Figure 9A). In contrast, Models 2 (speed model) and 3 (onset model) are feed-forward inhibition models. Here, the two accumulator units also integrate the input from the target whose choice they would trigger, but in addition they receive inhibitory input from the alternative target, with the inhibitory strength determined by the behavioral fit. The strength can therefore approach zero, which includes the condition enforced in Model 1. The key difference between Models 2 and 3 is the mechanism by which salience influences the integration process. Model 2 assumes that salience influences the quality of the perceptual process output, and thus the probability of accumulation, which is independent of the input strength that is determined by value (Figure 9B). This results in a modulation of the mean drift rate, which is orthogonal to the effect of value, as supported by our behavioral analysis. In Model 3, on the other hand, salience is assumed to influence the onset time of the accumulation process (Figure 9C), but not the probability of accumulation (i.e., mean drift rate). Finally, salience is free to modify both onset and drift rate in Model 4 (Figure 9D).

We optimized the parameters in all four models using the observed reaction times in correct choice trials, which were used as the training set for parameter tuning. We then compared the simulated reaction time distribution with the training set reaction time distribution using Pearson chi-square statistics (Purcell et al., 2010; Van Zandt, Colonius, & Proctor, 2000). This method maximized the proportion of correct responses in addition to matching the distribution of observed RTs. In order to avoid over fitting of the training data and to test the models' capability of prediction, in addition, we compared the predicted behavioral performance with two test sets, neither of which was used during training. One was the observed behavior in no-choice trials; the other was the behavior in erroneous choice trials (the trials in which the participant chose the lower valued target). In addition, we also used the BIC to test both fitness and prediction of the models. The results of this analysis are consistent with the results using the chi-square criterion (Table 2).

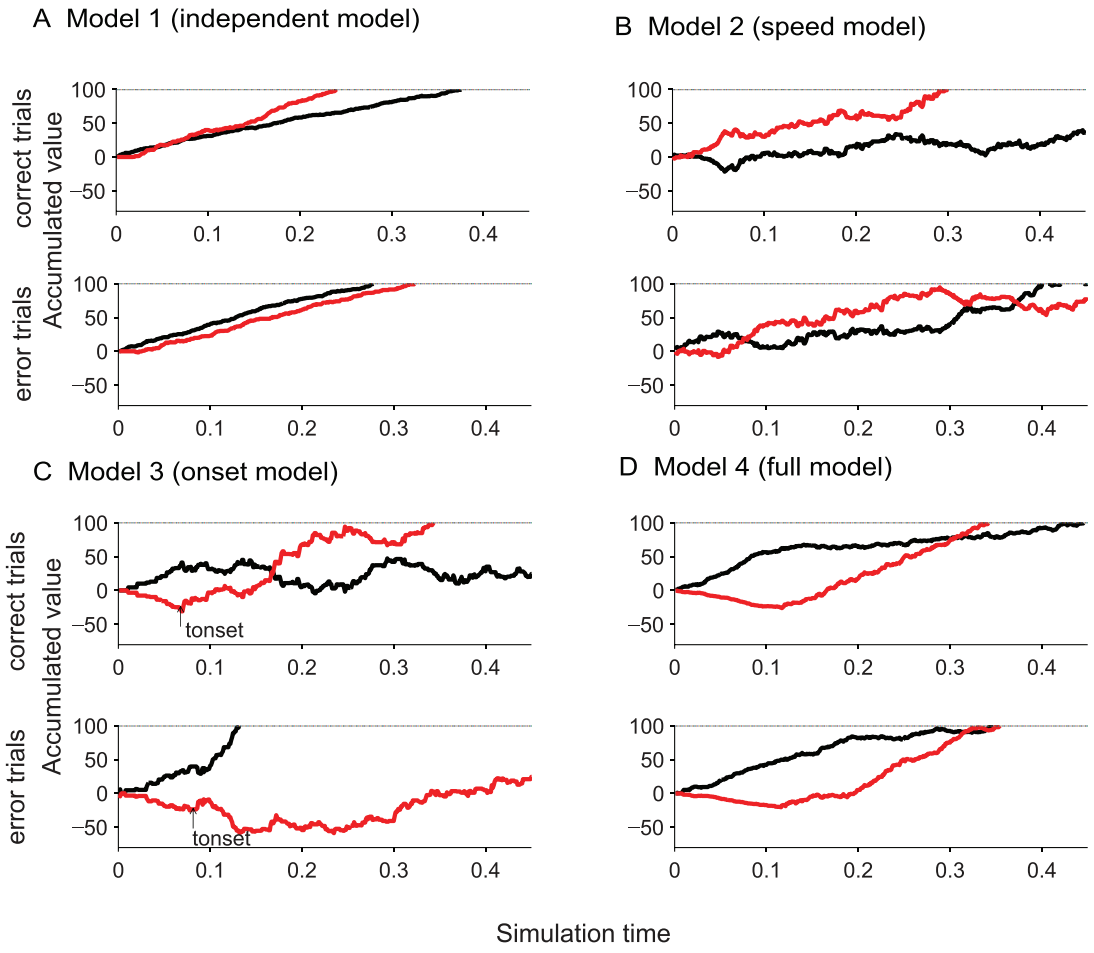

Figure 9. Examples of the time evolution of variables in independent models in incongruent trials. Within each plot, the upper panels are examples for correct trials, the lower panels for error trials. The paths for high-value low-salience targets are shown in red, and for low-value high-salience targets are in black. All competitions start at zero and threshold is always 100. (A) Model 1: independent model, no mutual inhibition between targets. (B) Model 2: speed model, salience influences the rate of accumulation. (C) Model 3: onset model, salience influences the onset of accumulation. (D) Model 4: the free model.

Mutual inhibition is necessary to explain behavior

The independent model (Model 1) did not explain the reaction times well. Its mean χ^2 fit (7.08) was significantly larger than that for the other two constrained models (Model 2: mean χ^2 fit, 2.44; t test, df = 29, $p < 10^{-8}$, Model 3: mean χ^2 fit, 2.24; t test, df = 29, $p < 10^{-10}$). More importantly, the predicted reaction times in no-choice trials did not fit the observed reaction times (Figures 10A and B). Specifically, a very general characteristic of the observed

	Choice trial		No-choice trial	
	BIC_{fit}	χ^2_{fit}	BIC_{test}	χ^2_{test}
Independent model	−111.86	7.08	−15.76	12.87
Speed model	−114.68	2.44	−39.30	8.38
Onset model	−115.51	2.44	−44.27	3.61
Full model	−114.37	3.50	−37.89	8.20

Table 2. Fitness of four different models in fitting reaction time on choice trials and predicting reaction time on no-choice trials. *Notes*: The first column shows the type of the model under testing. The second and third columns show the *BIC* value and χ^2 for the choice trials when optimizing model parameters during the training section. The fourth and fifth rows show the *BIC* value and χ^2 for the no-choice trials when testing the models' capability of prediction.

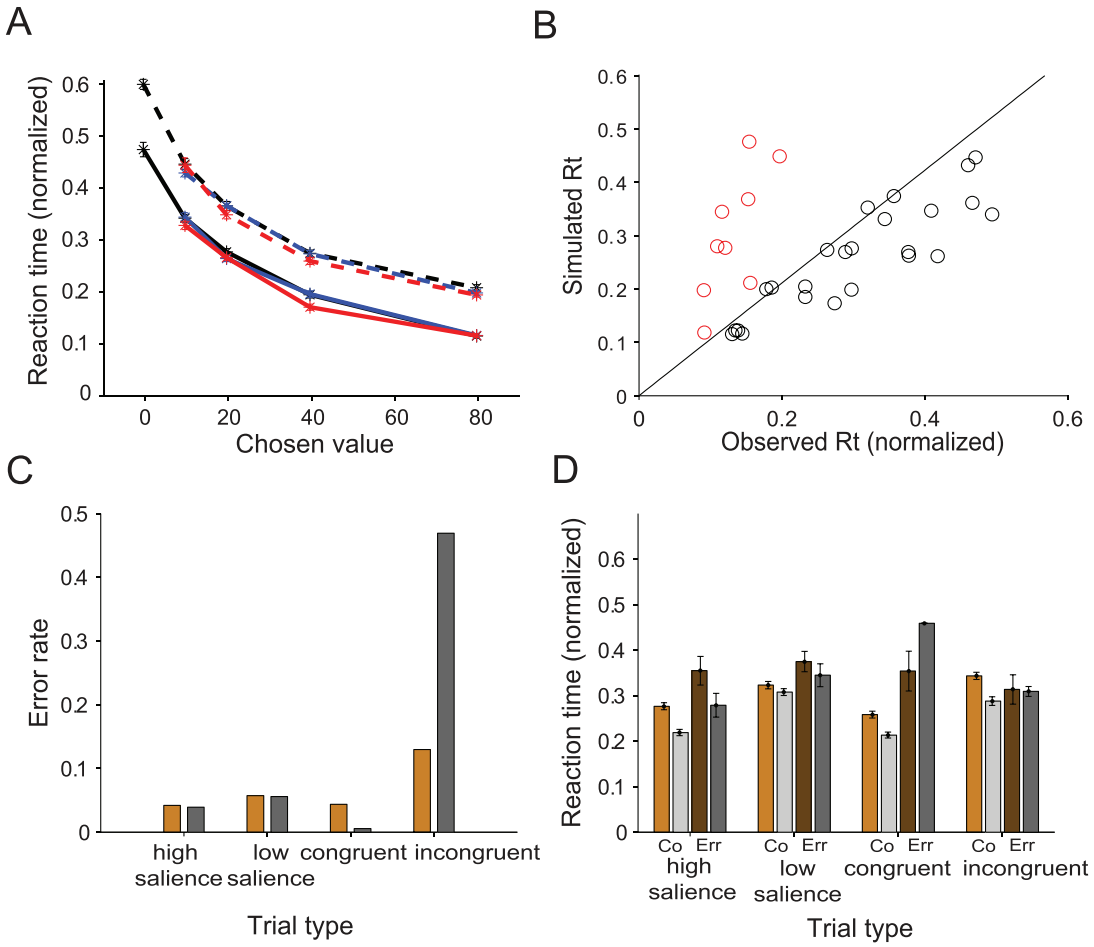

Figure 10. Predictions of Model 1 (independent model). (A) Predicted mean reaction times are plotted against the value of the chosen targets on different trial types. Each sample size is 100 simulations. Error bars represent standard error of the mean reaction time. Symbols are as in Figure 7. (B) Predicted mean reaction times are plotted against observed mean reaction time for 24 different choice trial types (black circle) and 10 different no-choice trial types (red circle). (C) Comparison between observed (light orange) and predicted error rate (gray) on different trial types. (D) Comparison between observed (correct trial: light orange; error trial: dark orange) and predicted (correct trial: light gray; error trial: dark gray) mean reaction times on different trials types. Symbols are as in Figure 5.

reaction time data was the increased reaction time latency on choice trials as opposed to no-choice trials (Figure 7). In contrast, the independent model predicted that the reaction times for choice trials were as fast as those in no-choice trials, due to the lack of inhibition from the non-chosen target. In addition, Model 1 overestimated the error rate on incongruent trials (Figure 10C) and it failed to predict the observation that on incongruent trials the reaction times on error trials were shorter than those on correct trials (Figure 10D). For no-choice trials, the mean χ^2 value (12.84) of Model 1 was significantly larger than that of the other two models (Model 2: χ^2: 8.38, t test, df = 29, $p < 10^{-15}$, Model 3: χ^2: 3.61, t test, df = 29, $p < 10^{-34}$). In sum, the analysis of Model 1 showed very clearly that mutual inhibition between two choices is important for target selection.

Salience affects the onset of accumulation

The speed model (Model 2) and onset model (Model 3) were both in good agreement with the training set, the observed reaction times of correct choice trials (Figures 11 and 12). Although the fitness of Model 2 (mean $\chi^2_{fit} = 2.44$) was slightly poorer than that of Model 3 (mean $\chi^2_{fit} = 2.24$), the difference was not significant (t test, df = 29, $p = 0.86$). The simulated reaction times of both Models 2 and 3 were correlated with the difference between chosen and non-chosen values, chosen values, non-chosen values, and the congruence between value and salience, which was consistent with the behavioral data.

However, the two models showed differences in their capability to predict the reaction time in the test set, the no-choice and erroneous choice trials. First, the two

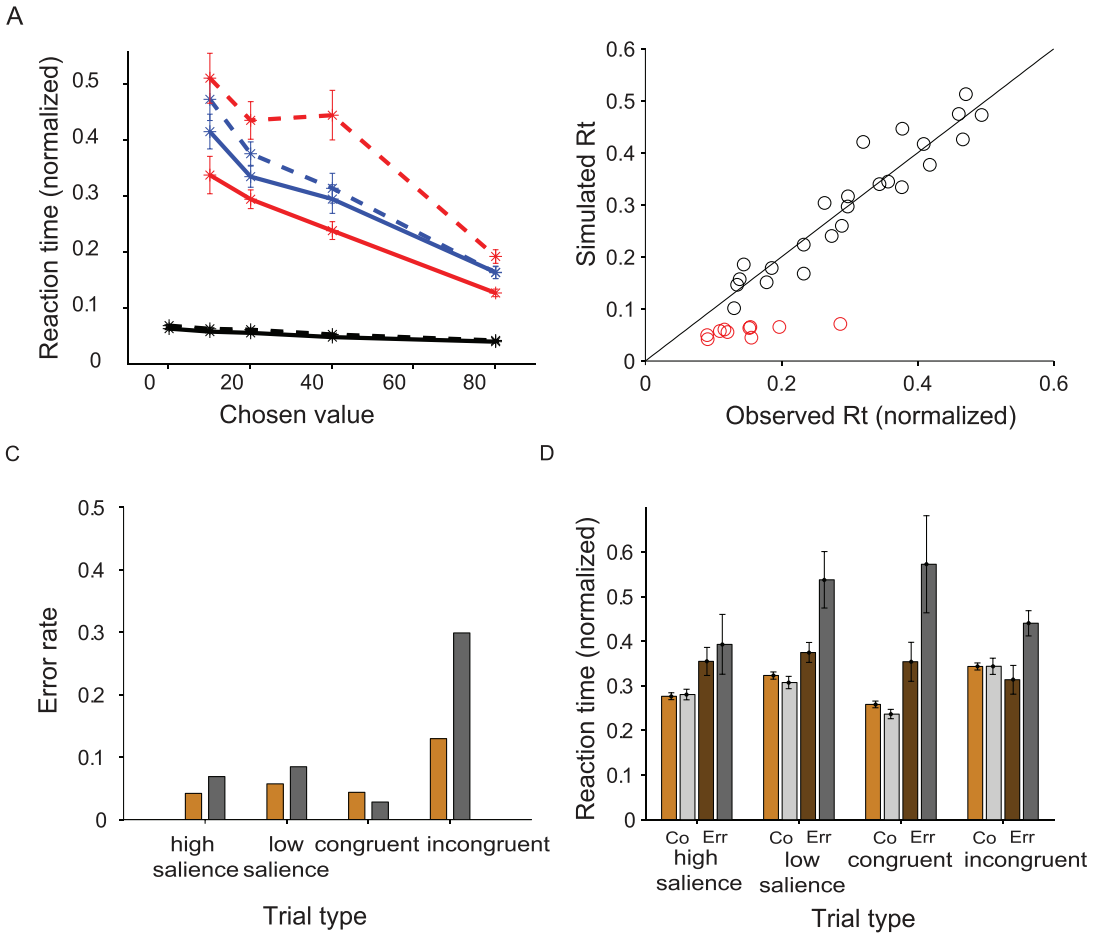

Figure 11. Predictions of Model 2 (speed model). Symbols are as in Figure 10.

models made different predictions for no-choice trials. In Model 2, in which salience modulates the accumulation rate, reaction time differences caused by salience differences were positively correlated with the time it took the accumulated activity to reach the threshold. Therefore, this model predicted that the reaction time difference caused by salience will be larger for choice trials than for no-choice trials (Figure 11A). On the other hand, in Model 3, in which salience modulates the accumulation onset, reaction time differences caused by salience were independent of how long it takes the activity to reach threshold. Therefore, this model predicted that the reaction time difference caused by salience will be similar for choice and no-choice trials (Figure 12A). Consequently, when the parameters of both models were adjusted so that they fit the reaction time differences in choice trials, the salience-induced reaction time differences in no-choice trials predicted by Model 2 should be smaller than the ones predicted by Model 3. This is exactly what the simulations showed (Figures 11A and 12A). Specifically, Model 2 predicted that reaction times in no-choice trials should not be significantly influenced by salience, as can be seen by comparing the solid and dotted black lines in Figure 11B. This prediction is not unreasonable. The luminance contrast of high-salience targets was similar to that of low-salience targets. This allowed the single targets to be equally localized. Furthermore, value information was not behaviorally relevant, since no choice was required. Therefore, one might expect that differences in salience do not necessarily influence reaction time similar as in the singleton task (Figure 7). In contrast, Model 3 predicted that reaction times to targets with high salience of the reward information should be faster than the ones to targets with low salience, as indicated in the corresponding plot in Figure 12A.

This difference in the predictions of the two models allowed us to compare how well they fit the observed behavioral data. A comparison of the actual (Figure 7) with the predicted reaction times (Figure 11A, 12A) showed that, overall, Model 3 agreed much better with observations than Model 2. The behavioral fit of Model 3 (mean $\chi^2_{test} = 3.61$) in no-choice trials was significantly (t test, df = 29, $p < 10^{-15}$) better than that of Model 2 (mean $\chi^2_{test} = 8.38$). In particular, the human

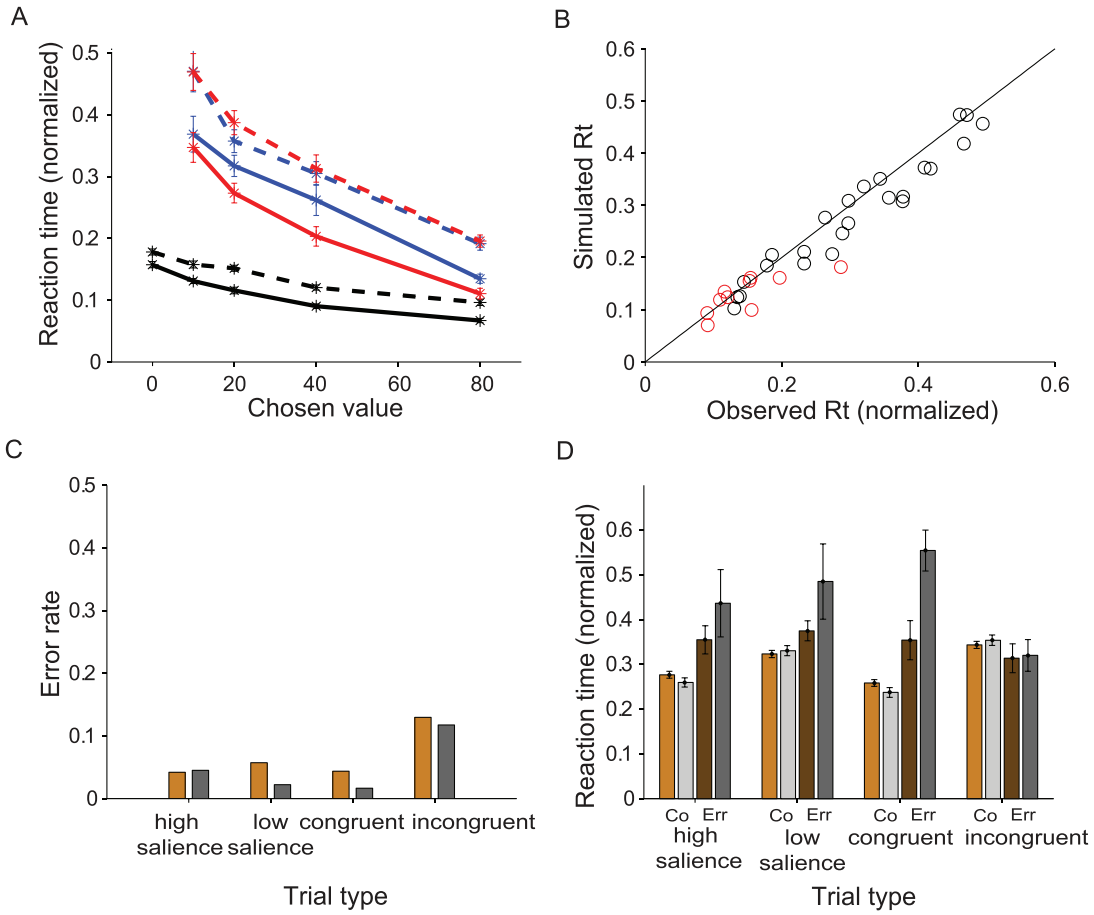

Figure 12. Predictions of Model 3 (onset model). Symbols are as in Figure 10.

participants showed consistently longer reaction times in no-choice trials with low-salience targets compared to trials with high-salience targets. Consequently, the comparison of predicted to observed reaction times in Figures 11B and 12B showed that for most no-choice trials (indicated by the red circles) the predicted RTs of Model 2 are too short, while the predictions of Model 3 for no-choice trials were as accurate as for choice trials.

Second, the two models also made different predictions with regards to the reaction times in error trials, providing another opportunity to determine which model provided a better description of behavior. The behavioral data showed that the reaction time in all trial types was significantly longer for erroneous than for correct responses, except for incongruent trials for which the pattern was the opposite (Figure 5C). Thus, there was an inversion of reaction time for congruent and incongruent trials with respect to correct and erroneous responses. Both models predicted correctly that in congruent trials the reaction time was significantly longer on error trials than on correct trials (Model 2: K-S test, $p = 0.004$; Model 3: K-S test, $p = 0.003$). On incongruent trials, however, Model 2 predicted that the reaction time on erroneous trials was also significantly longer (K-S test, $p = 0.002$) than on correct trials (Figure 11D). This is because an error can only occur when the accumulator associated with the low-value target happens to reach the threshold earlier than the one associated with the high-value target. Since in Model 2 the low-value accumulator tended to rise slowly, this could only happen if the competing high-value accumulator also rose slowly. Thus, in this model the reaction time on error trials had to be longer than on correct trials (Figure 9B).

In contrast, Model 3 accurately predicted shorter reaction times on erroneous than on correct incongruent trials (Figure 12D). Sensitivity analysis showed that the difference in onset time of the accumulation between high- and low-salience targets was significantly linearly correlated (t test, $\mathrm{df} = 47$, $p < 10^{-8}$, Sobol Index: 0.43) with the reaction time differences between error and correct trials on incongruent trials (Figure 13). In Model 3, errors were due to the earlier onset of accumulation for the low-value targets. This onset time difference, especially when it was large, created a window of opportunity for the low-value accumulation process during which it experienced no competition from the high-value accumulation process. Therefore,

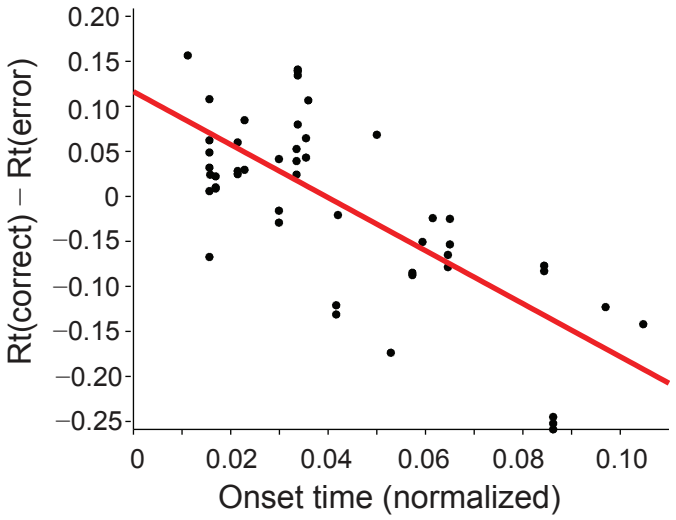

Figure 13. Effect of difference in the onset time of accumulation in Model 3 (onset model). The predicted mean reaction time difference between error and correct trials on incongruent trials is correlated with the simulated time of accumulation onset, shown in 50 different simulations (black circles), together with the regression line (red).

in this model the reaction times on error trials tended to be shorter than on correct trials (Figure 9C). The fact that this prediction, which was specific for Model 3, was confirmed by the behavioral data gives further support for the hypothesis that differences in the onset latency of accumulation (as in Model 3) rather than in the rate of accumulation (as in Model 2) explains the salience effect in our behavioral choice task.

Finally, we also optimized the full model, where salience could influence both onset and rate of the accumulation (Figure 14). This model has one parameter more than Models 2 and 3. We would therefore expect at least as good a fit to the training data as the best of the less complex models. However, Model 4 does not necessarily have better predictive power for the test data since it might still contain terms based on incorrect assumptions. Indeed, the full model did not significantly improve the accuracy of the reaction time fits in the training set (correct choice trial, mean $\chi^2_{fit} = 3.50$), and it resulted in decreased accuracy when predicting the reaction time in the testing set (no-choice trials, mean $\chi^2_{test} = 8.20$) over Model 3. This provided additional support for Model 3.

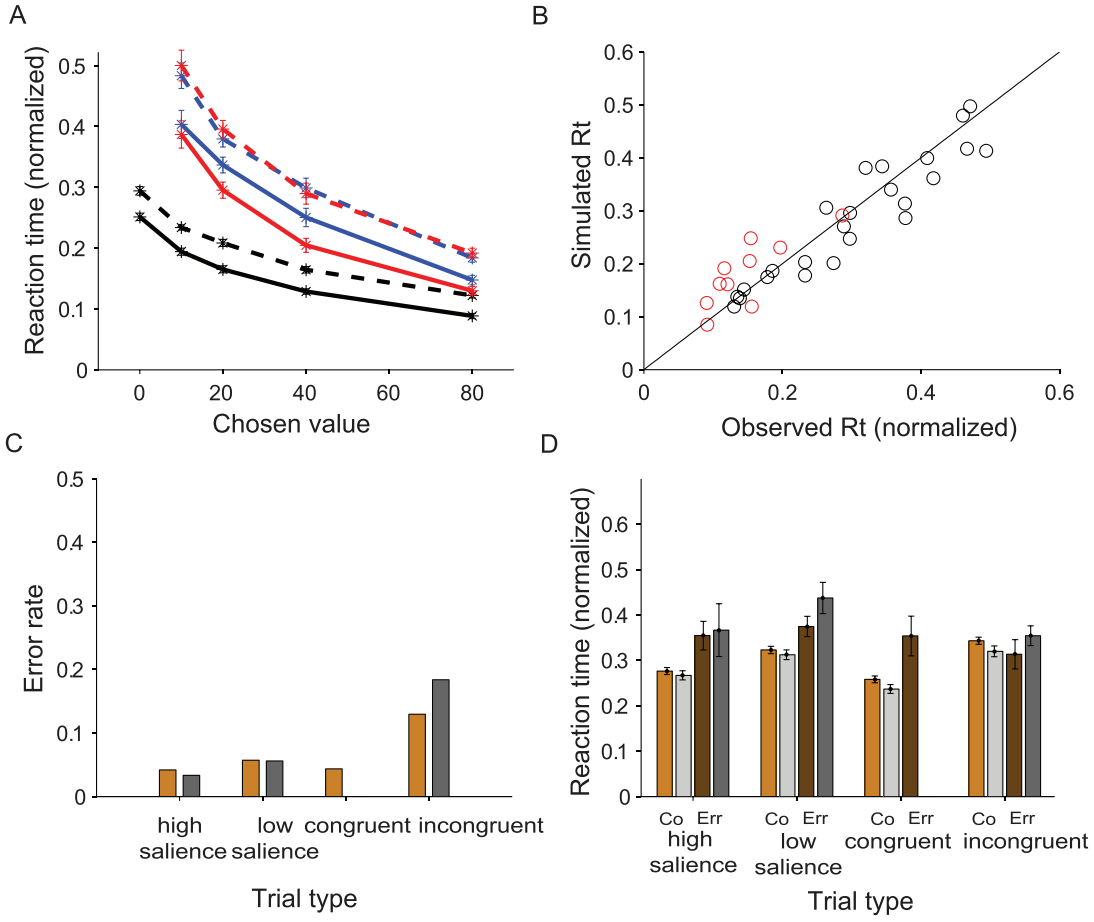

Figure 14. Predictions of Model 4 (full model). Symbols are as in Figure 10.

Discussion

We studied the influence of visual salience on value-based decision processes of human observers performing a two-alternatives forced-choice task. Most previous studies examined the effect of visual saliency (Berg et al., 2009) and value (Bendiksby & Platt, 2006; Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010; Milstein & Dorris, 2007; Platt & Glimcher, 1999) on saccades and decisions in isolation. In contrast, we manipulated saliency and value of targets simultaneously in order to investigate how these two factors interact and how they influence the decision process.

Our behavioral results showed that not only value, but also visual salience as well as congruency between value and visual salience influence the decision process. Specifically, we found that reaction time across all trial types is correlated with all three of these variables. Furthermore, congruency had an effect on error rates as well as on reaction times in error trials. These findings are in broad agreement with recent behavioral studies of eye movements in macaque monkeys and humans that also manipulated salience and value information simultaneously (Markowitz et al., 2011; Navalpakkam et al., 2010; Schutz et al., 2012). This convergence of findings across species and effector systems indicates that these behavioral trends are robust and reflect basic selection mechanisms in the primate brain. However, despite the similarity of the behavioral findings, the mechanism underlying these phenomena can differ depending on whether salience and value information share the same feature dimension as will be discussed later.

Onset time differences

The speed and the onset model are based on different hypotheses about how salience modulates the decision process. In the speed model, salience influences accumulation speed by modulating the probability of an increase in activity, which influences how likely it is that the accumulator responds to the input independent of its strength. A possible neurophysiological interpretation could be that salience modulates the likelihood that an individual neuron responds to synaptic input, or how many neurons out of the entire pool of decision neurons respond to the input in a given time interval and for a given stimulation. This interpretation does not seem to be unreasonable, given our current understanding of primate decision-making mechanisms. It is therefore quite noteworthy that our analysis ruled out this model so convincingly. Instead another model, namely the onset model, explained the observed reaction time distribution of error trials and no-choice trials much better.

In the onset model, salience influences the onset time of accumulators to accumulate value information instead of influencing accumulation speed. Specifically, this stochastic model suggests an earlier accumulation onset time when processing high- compared to low-salience targets. What might be the sources of this earlier onset? One possibility is that the difference is due to an attentional shift. Specifically, high-salience targets might attract attention first, as predicted by bottom-up attentional guidance models (Itti, Koch, & Niebur, 1998). Focus on one of the targets might in turn result in an advantage for the accumulation process associated with this target, similar to the logic underlying a recent model of how value-based decisions are guided by visual attention (Krajbich & Rangel, 2011; Lim, O'Doherty, & Rangel, 2011). However, in our paradigm, attentional guidance can only explain the behavioral results on choice trials but not on no-choice trials. In the latter, attention should always be focused on the only target present, but we found that reaction times still differed between high- and low-salience trials (Figure 7). Even though saliency-based models (Itti et al., 1998) predict slightly faster deployment of attention to more salient than to less salient targets, the effect is likely too small to explain the size of the observed reaction time difference in no-choice trials. Allocation of attention cannot, then, be the main reason for the onset differences in the accumulation process.

An alternative hypothesis is that the onset difference is caused by differences in the visual processing time required for computing the value of high- and low-salience targets (Ratcliff, Hasegawa, Hasegawa, Smith, & Segraves, 2007; Ratcliff & Smith, 2011). Low-salience targets might need more time to be identified than high-salience targets. This hypothesis is consistent with our behavioral data, since the salience effect on reaction time is similar for choice and no-choice trials. It is also supported by neurophysiological findings of clear effects of contrast, but not of attention, on visual response latencies in primates (Lee, Williford, & Maunsell, 2007). Our results suggest therefore that the onset time of accumulation is time locked to the end of the visual processing. This implies a sequential form of information processing, similar to a recent model of target selection in frontal eye fields (Purcell et al., 2010). What might keep the accumulators from integrating evidence earlier? One possibility is that a specific gating mechanism blocks accumulation until a certain degree of difference has developed in the input sources. Alternatively, in our model, information about the location of potential targets drives the accumulators equally well and the two accumulators inhibit each other sufficiently to suppress any activity increase. Only

when value information is added is the symmetry broken and the differentiated accumulation process can start.

Functional architecture of value-based decision making in primates

Our analysis using stochastic accumulator models aimed at investigating the neuronal mechanisms underlying behavior. Obviously, simple modeling studies can only provide indirect evidence about mechanisms implemented in the brain. Nevertheless, the fact that our models show qualitative differences in their ability to explain behavioral data allows us to rule out entire families of functional architectures. The failure of Model 1, the independent model, shows that mutual inhibition between choices is required for the decision mechanism. The comparison between the two feed-forward inhibition models, speed and onset model, shows that increased visual salience leads to an earlier start of accumulation rather than an increased rate of accumulation. That the rate of accumulation plays a minor role, if any, is underscored by the fact that including it in addition to the change in onset time (Model 4) does not improve performance significantly. These findings allow us to formulate a strong hypothesis about the type of decision architecture that underlies the effects of visual salience and value information on choice behavior.

One hypothesis is that value and salience independently drive action selection and compete for access to the motor system. Indeed, we observed that early choices reflected more strongly visual salience, while later choices were more driven by value, consistent with previous findings by Markowitz et al. (2011). These authors suggested that the reaction time difference is due to the time-varying balance between stimulus- and reward-driven selections. This was reasonable, since in their study salience influenced both location and value-related information simultaneously. Information about target location influences the preparation of motor acts directed towards these locations (Schiller & Tehovnik, 2005). Information about target identity, on the other hand, is crucial for the generation of value information, which in turn affects competition between different value options (Padoa-Schioppa, 2011). However, in our study, we designed stimuli whose salience influenced only the discrimination of target identity, but did not affect the ability to locate a target. Therefore, automatic selection processes driven by stimulus location were equally strong for all targets as confirmed by our second pilot study (singleton task; see Figures 4 and 7). If value and salience were competing during the decision process, the choice behavior should show a joint dependence on sensory and goal-directed processes. In contrast, our behavioral results do not show an interaction between visual salience and value.

This suggests a second hypothesis that there is a combined value and salience map within the visuomotor system (Navalpakkam et al., 2010; Schutz et al., 2012). Under this hypothesis, the salience map formed is first influenced by bottom-up factors (in our case luminance contrast with the background) and is only later modified by value information. This implies that bottom-up salience alone can influence behavior independent of value information, in particular during early responses. However, in our study we were interested in dissociating the effect of salience on value-based decision making from the one on motor generation in general. Indeed, the salience manipulation that we used does not seem to have influenced motor behavior very much, since the reaction time for both high- and low-salience targets is the same in the singleton task (Figures 4 and 7). Moreover, if the salience map is modulated by the attention captured by the objects associated with reward (Anderson, Laurent, & Yantis, 2011), we would expect to see a joint dependence between reward and salience, which is not shown in the result. Therefore, the salience effect in our experiment is less likely to be the result of early activation in the salience map driven by bottom-up factors.

Instead, these findings support a third alternative hypothesis about the interaction of salience and value in decision making; namely that value-based action selection is a serial iterative decision process. First, lower sensory areas process the stimuli and derive value and location information from them. The salience of the relevant visual feature influences how long this process takes but has no influence on the output of this stage. The processed information is then sent to accumulators in a higher comparative area that selects the final (and motor) response. At this point the value input has been stripped of other content, such as saliency. This serial processing hypothesis contains the predictions of the combined saliency map hypothesis as a special case. When the sensory features carrying both target location and identity information are of varying salience, the targets with more salient features will influence the decision-making stage earlier and there will be a higher likelihood that the subject will choose the more salient target (in particular, if the accumulation process was fast and there was less time for the less salient target to reach the accumulation stage). Behavior will appear as if bottom-up salience of the targets alone can influence choice independent of value information. Thus, for this set of experimental conditions, both hypotheses explain the behavior equally well. In contrast, our present experimental findings can only be explained by the serial processing hypothesis but not the combined saliency map hypothesis. This

attribute of the serial processing hypothesis is attractive. Nevertheless, at this point we do not know if this conjecture is indeed correct, since we have not tested subjects in both conditions and determined if our model can really explain data across both situations.

The higher order areas involved in the final decision likely include structures representing subjective value, such as orbitofrontal (Padoa-Schioppa & Assad, 2006) and ventromedial prefrontal cortex (Kim, Hwang, & Lee, 2008; Lim et al., 2011), and/or visual-motor association structures, such as the lateral intraparietal area (Platt & Glimcher, 1999; Sugrue, Corrado, & Newsome, 2004), supplementary eye field (So & Stuphorn, 2010), or the supplementary motor area (Scangos & Stuphorn, 2010). Given our results, it is worth investigating whether and how these areas integrate both visual salience and value information during value-based decision making when salience influences selectively the value perception of an object, but not awareness of its presence.

Keywords: decision making, accumulator, reaction time

Acknowledgments

We are grateful to H. Egeth for helpful comments on the manuscript. This work was supported by the National Eye Institute through grants R01-EY019039 and R01EY016281, and by ONR though grant N000141010278.

Commercial relationships: none.
Corresponding author: Veit Stuphorn.
Email: veit@jhu.edu.
Address: Solomon H. Snyder Department of Neuroscience, Johns Hopkins University School of Medicine and Zanvyl Krieger Mind/Brain Institute, Baltimore, MD, USA.

References

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences, USA, 108*(25), 10367–10371.

Bendiksby, M. S., & Platt, M. L. (2006). Neural correlates of reward and attention in macaque area LIP. *Neuropsychologia*, *44*(12), 2411–2420.

Berg, D. J., Boehnke, S. E., Marino, R. A., Munoz, D. P., & Itti, L. (2009). Free viewing of dynamic stimuli by humans and monkeys. *Journal of Vision*, *9*(5):19, 1–15, http://www.journalofvision.org/content/9/5/19, doi:10.1167/9.5.19. [PubMed] [Article]

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review, 113*(4), 700–765.

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference: A practical information-theoretic approach*. New York, NY: Springer.

Busemeyer, J. R., & Diederich, A. (2010). *Cognitive modeling*. Thousand Oaks, CA: Sage.

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience, 29*(37), 11560–11571.

Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology, 48,* 269–297.

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience*, *31*(17), 6339–6352.

Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *20*(11), 1254–1259.

Kass, R. E., & Wasserman, L. (1995). A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *Journal of the American Statistical Association, 90*(431), 928–934.

Kim, S., Hwang, J., & Lee, D. (2008). Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron*, *59*(1), 161–172.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences, USA, 108*(33), 13852–13857.

Lee, J., Williford, T., & Maunsell, J. H. (2007). Spatial attention and the latency of neuronal responses in macaque area V4. *Journal of Neuroscience*, *27*(36), 9632–9637.

Lim, S. L., O'Doherty, J. P., & Rangel, A. (2011). The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. *Journal of Neuroscience*, *31*(37), 13214–13223.

Markowitz, D. A., Shewcraft, R. A., Wong, Y. T., &

Pesaran, B. (2011). Competition for visual selection in the oculomotor system. *Journal of Neuroscience*, *31*(25), 9298–9306.

Masciocchi, C. M., Mihalas, S., Parkhurst, D., & Niebur, E. (2009). Everyone knows what is interesting: Salient locations which should be fixated. *Journal of Vision, 9*(11):25, 1–22, http://www.journalofvision.org/content/9/11/25, doi:10.1167/9.11.25. [PubMed] [Article]

Mihalas, S., Dong, Y., von der Heydt, R., & Niebur, E. (2011). Mechanisms of perceptual organization provide auto-zoom and auto-localization for attention to objects. *Proceedings of the National Academy of Sciences, USA*, *108*(18), 7583–7588.

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. *Judgment and Decision Making*, *5*(6), 437–449.

Milstein, D. M., & Dorris, M. C. (2007). The influence of expected value on saccadic preparation. *Journal of Neuroscience*, *27*(18), 4810–4818.

Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences, USA, 107*(11), 5232–5237.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal, 7,* 308–313.

Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. *Annual Review of Neuroscience, 34,* 333–359.

Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, *441*(7090), 223–226.

Parkhurst, D. J., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, *42*(1), 107–123.

Parkhurst, D. J., & Niebur, E. (2003). Scene content selected by active vision. *Spatial Vision*, *16*(2), 125–154.

Pike, A. R. (1966). Stochastic models of choice behaviour: Response probabilities and latencies of finite Markov chain systems. *British Journal of Mathematical and Statistical Psychology, 19*(1), 15–32.

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233–238.

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review*, *117*(4), 1113–1143.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.

Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology* , *97*(2), 1756–1774.

Ratcliff, R., & Smith, P. L. (2011). Perceptual discrimination in static and dynamic noise: The temporal relation between perceptual encoding and decision making. *Journal of Experimental Psychology, 139*(1), 70–94.

Sawa, T. (1978). Information criteria for discriminating among alternative regression models. *Econometrica,* (46), 1273–1291.

Scangos, K. W., & Stuphorn, V. (2010). Medial frontal cortex motivates but does not control movement initiation in the countermanding task. *Journal of Neuroscience*, *30*(5), 1968–1982.

Schiller, P. H., & Tehovnik, E. J. (2005). Neural mechanisms underlying target selection with saccadic eye movements. *Progress in Brain Research, 149,* 157–171.

Schutz, A. C., Trommershauser, J., & Gegenfurtner, K. R. (2012). Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences, USA, 109*(19), 7547–7552.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, *86*(4), 1916–1936.

So, N. Y., & Stuphorn, V. (2010). Supplementary eye field encodes option and action value for saccades with variable reward. *Journal of Neurophysiology*, *104*(5), 2634–2653.

Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, *304*(5678), 1782–1787.

Van Zandt, T., Colonius, H., & Proctor, R. W. (2000). A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin and Review*, *7*(2), 208–256.

Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, *13*(1), 37–58.

White, B. J., & Munoz, D. P. (2011). Separate visual

signals for saccade initiation during target selection in the primate superior colliculus. *Journal of Neuroscience*, *31*(5), 1570–1578.

Wolfe, J. M. (1998). Visual search. In H. Pashler (Ed.), *Attention* (Vol. 1., pp. 17–73): London UK: University College London Press.

Ziehn, T., & Tomlin, A. (2009). GUI-HDMR: A software tool for global sensitivity analysis of complex models. *Environmental Modelling and Software, 24,* 775–785.

Appendix

Regression analysis

We quantified the degree to which reaction times reflected different qualities of the presented targets by fitting a family of linear regression models to the reaction times and determining the best-fitting model. We first normalized data by dividing all reaction times of a given participant by the largest value in correct choice trials. The normalized reaction times were then used for regression analysis. The eight tested factors were: (a) value of the chosen target, (b) value of the non-chosen target, (c) salience of the chosen target, (d) salience of the non-chosen target, (e) value difference between chosen and non-chosen target, (f) salience difference between chosen and non-chosen target, (g) multiplicative interaction between value and salience of the chosen target, and (h) multiplicative interaction between value and salience of the non-chosen target. We identified the best fitting model by minimizing the Bayesian information criterion (BIC),

$$BIC = n\times\log\left(\frac{RSS}{n}\right) + K\times\log(n), \tag{A1}$$

where n is the number of trials (a constant in our case), k is the number of fitting parameters, and RSS the residual sum of squares after fitting (Burnham & Anderson, 2002; Busemeyer & Diederich, 2010). A lower numerical BIC value indicates better fit of a model, with a lower residual sum of squares indicating better predictive power, and a larger K penalizes less parsimonious models. This procedure is related to a likelihood-ratio test, and equivalent to choosing a model based on the F statistic (Sawa, 1978). It provides a Bayesian test for nested hypotheses (Kass & Wasserman, 1995).

We also used AIC test, which is similar to BIC test that deals with the trade-off between the complexity of the model and the goodness of fit of the model. Different from BIC, AIC penalizes the number of parameters less strongly. AIC value is defined as:

$$AIC = n\times\log\left(\frac{RSS}{n}\right) + 2K. \tag{A2}$$

BIC evidence ratio test represent the relative likelihood of the best model versus Model I (Burnham & Anderson, 2002). Evidence ratio is the ratio of BIC weight and is defined as:

$$\frac{w_1}{w_i} \equiv \frac{1}{e^{-1/2\Delta_i}} \equiv e^{1/2\Delta_i}. \tag{A3}$$

w_1 is the BIC weight for the best model that equals to one. Δ_i is the BIC difference for each model:

$$\Delta_i = BIC_i - BIC_{min}, \tag{A4}$$

in which BIC_{min} is the model with the smallest BIC value.

An alternative procedure would have been a series of sequential F tests, but this approach, while exact, requires the assumption of data with a normal distribution that is not true for the reaction time distribution. We decided therefore to use model selection criteria based on information theories, because they were computationally straightforward and potentially more robust. An additional advantage was that we could compare all 255 models simultaneously, using a consistent criterion (Burnham & Anderson, 2002).

Accumulator models

Behavioral data were fitted using accumulation models, also called race or drift-diffusion models (Pike, 1966; Ratcliff, 1978; Vickers, 1970),which have been shown to explain decision making in value-judgment tasks (Milosavljevic et al., 2010). We studied four related models that are schematically shown in Figure 8. Models 1–3 (independent model, speed model, and onset model, respectively) are variants of a more general model, Model 4 (full model), and are derived from it by constraining one or more of its parameters.

We started with the description of the most general model, Model 4 (Figure 8D). In choice trials, it consists of two integrators m_1 and m_2 that accumulate evidence in favor of the two possible choices as described by a stochastic differential equation:

$$\begin{aligned}\{dm_1 &= I_1(1+dw_1)H(t-t_1)N_1(s_1,dt)\\ &\quad -uI_2(1+dw_2)H(t-t_2)N_2(s_2,dt),\\ \{dm_2 &= I_2(1+dw_2)H(t-t_2)N_2(s_2,dt)\\ &\quad -uI_1(1+dw_1)H(t-t_1)N_1(s_1,dt),\end{aligned}$$

$$m_1(0) = m_2(0) = 0. \tag{A5}$$

In the simulations, we always chose $dt = 1ms$. The first

term in Equation A5 describes how each integrator accumulates evidence from its corresponding target. It has four factors. The first, I_1, and I_2, is the mean quanta of accumulation. We assumed that this quanta is determined by the normalized value of the target, $v_i \in$ (0, 0.1, 0.2, 0.4, 0.8). The second, $(1 + dw_i)$ introduces variability for each accumulated quantum. As dw_i is drawn from a normal distribution, at each time the instantaneous accumulation, $I(1 + dw_i)$ comes from a Gaussian distribution whose mean is equal to its standard deviation and both are determined by

$$\begin{cases} I_1 = a_0 + a_1 v_1 \\ I_2 = a_0 + a_1 v_2 \end{cases}, \tag{A6}$$

and a_0 and a_1 are fitted coefficients. We equated standard deviation and mean of the distribution to minimize the number of free parameters.

The third factor is the onset time of accumulation. This is determined by $H(t - t_i)$ in Equation A5. $H(t)$ is the Heaviside step function:

$$H(t) = \begin{cases} 1, & t \geq 0 \\ 0, & t < 0 \end{cases}. \tag{A7}$$

t_i is negatively correlated with the salience of the target. $t_i = t_0 + t_i s_i$, where s_i is the salience of target I_i, and t_0, and t_1 are fitted coefficients.

The fourth factor is the probability of accumulation, which is also determined by the salience of the target:

$$\begin{aligned} &\{P(N_1(t + dt) - N_1(t) = 1) \\ &\quad = P(N_1(dt) = 1) = (\lambda_0 + \lambda s_1)dt; \\ &\{P(N_2(t + dt) - N_2(t) = 1) \\ &\quad = P(N_2(dt) = 1) = (\lambda_0 + \lambda s_2)dt; \end{aligned} \tag{A8}$$

where λ_0 and λ are fitted coefficients. No accumulation occurs in accumulator m_i. if $N_i = 0$ and the probability of this happening is determined by s_i through Equation 8. The mean rate of accumulation is determined by $I_i(\lambda_0 + \lambda s_i)$. By having the probability of accumulation term, the salience of value information could influence accumulation rate in a way that is orthogonal to the effect of value itself. That is the influence of salience is independent of the strength of value information.

Finally, u in Equation A5 is the inhibition coefficient. It determines the strength of inhibition between two integrators.

In addition to the general model defined in Equations A5–A8, we now introduce the more restricted Models 1–3. Each is obtained from the general model by removing some degrees of freedom. Model 1 (Figure 8A) is an independent model in which accumulators for each of the two alternatives integrate evidence independently. The inhibition coefficient u (Equation A5) was set to zero in this model. Information for each integrator is therefore accumulated according to the simplified version:

$$\begin{cases} dm_1 = I_1(1 + dw_1)H(t - t_1)N_1(s_1, dt), \\ dm_2 = I_2(1 + dw_2)H(t - t_2)N_2(s_2, dt), \end{cases}$$

$$m_1(0) = m_2(0) = 0. \tag{A9}$$

Models 2, 3, and 4 are feed-forward inhibition models (Purcell et al., 2010; Shadlen & Newsome, 2001). We assume that the amount of momentary evidence for accumulation is positively correlated with the value of the corresponding targets and we take into account the influence of the salience of the target based on two different hypotheses. In Model 2 (speed model) (Figure 8B), we assume that salience influences the rate, but not the onset time of accumulation. Therefore, onset difference, t_i, was set to zero for Model 2. The stochastic differential equations for integrators accumulating evidence can then be simplified as:

$$\begin{cases} dm_1 = I_1(1 + dw_1)N_1(s_1, dt) - uI_2(1 + dw_2)N_2(s_2, dt), \\ dm_2 = I_2(1 + dw_2)N_2(s_2, dt) - uI_1(1 + dw_1)N_1(s_1, dt), \end{cases}$$

$$m_1(0) = m_2(0) = 0. \tag{A10}$$

In contrast, in Model 3 (onset model) (Figure 8C) we assume that salience influences the onset time rather the rate of accumulation. The accumulation rate was determined by the probability to receive momentary evidence, $P(N_i(dt))$. Accordingly, in this model the probability of accumulation for each integrator in a given time interval, and therefore the rate of accumulation of this accumulator, is independent of the salience level of the corresponding target.

$$P(N_1(dt) = 1) = P(N_2(dt) = 1) = \lambda_0. \tag{A11}$$

The stochastic differential equations for integrators accumulating evidence can be simplified as

$$\begin{aligned} \{dm_1 &= I_1(1 + dw_1)H(t - t_1)N_1(dt) \\ &\quad - uI_2(1 + dw_2)H(t - t_2)N_2(dt), \\ \{dm_2 &= I_2(1 + dw_2)H(t - t_2)N_2(dt) \\ &\quad - uI_1(1 + dw_1)H(t - t_1)N_1(dt), \end{aligned}$$

$$m_1(0) = m_2(0) = 0. \tag{A12}$$

Traditionally the accumulators in decision making have been modeled as a Wiener process and a continuous accumulation of information. Since we want to test the hypothesis that value influences accumulated quanta magnitude and salience accumulated quanta rate, we have to use the more general description presented above. However, it should be noted that in Model 3, the accumulation quanta frequency is independent of stimulus properties and thus could be mapped more easily to a Wiener process.

Having defined the models for choice trials, we now define the general model for no-choice trials. It is identical to the choice-trial model (Equation A5) but only one dependent variable (m_1) is defined. Thus, there is only one integrator m_1 and only one input unit, with all values corresponding to the second target set to zero.

Simulation methods

We adopted standard model fitting techniques to optimize the values of parameters that provided the best fit to the behavioral data. For each set of parameter values, we generated 100 simulated trials to produce predicted RT distributions for all 24 choice trial conditions (4 Trial Types × 6 Values). We used Pearson chi-square statistics to quantify the discrepancies between the observed and predicted cumulative correct RT distribution (Van Zandt et al., 2000)

$$\chi^2 = n\sum_i \frac{(o_i - p_i)^2}{p_i}. \quad \text{(A13)}$$

RTs were binned corresponding to the cumulative probabilities of 0.1, 0.3, 0.5, 0.7, and 0.9. o_i is the observed proportion of RTs in these bins, p_i is the corresponding proportions from the respective model, and n is the number of correct trials in a given experiment that is the total number of data points in the observed (and simulated) RT distribution. The advantage of this fitting method is that it maximizes the proportion of correctly fitted responses in addition to matching the distribution of observed RTs (Purcell et al., 2010; Van Zandt et al., 2000).

We used a hybrid scheme implemented in MATLAB (The MathWorks, Natick, MA) to find values of the free parameters that minimize the chi-square value. We first ran a genetic algorithm (GA) to find parameter values near a global optimum. Then the solution derived by the GA was used as an initial starting point for the simplex algorithm (Nelder & Mead, 1965) that is more efficient for a local search. For data sets with 24 conditions, all conditions were fitted simultaneously by summing the individual chi-square statistics for all of them. For each model and data set, we ran this hybrid routine 30 times with different, randomly selected initial conditions to mitigate the problem of finding local minima.

To evaluate the fitness of the four models, we also compared their predictions to behavior. To do this, we first used the best fitting parameters for a given model to simulate 100 RT distributions for all conditions, which included 24 conditions in choice trials and 10 conditions in no-choice trials. We repeated this procedure 30 times for all four models, generating 30 χ^2 values separately for correct choice trials [χ^2_{fit}] and no-choice trials [χ^2_{test}]. Thus, the χ^2_{fit} values tested the fitness using those observed data that had previously been used in the parameter optimization of the models. In contrast, the χ^2_{test} values tested the fitness using those observed data that had not been used in the parameter optimization and provided an independent assessment of the predictability of the different models. In addition, we also compared the percentage of error trials and the RTs in error trials in different types of choice trials to predictions from the optimized models. Next, we performed a global sensitivity analysis (Mihalas, Dong, von der Heydt, & Niebur, 2011; Ziehn & Tomlin, 2009) in MATLAB on seven parameters (u, a_0, a_1, t_0, t_1, a, λ_0, and λ in Equations 5–8) to specify which parameters contributed to the differences between models and to what extent.