1 **Monkeys exhibit human-like gaze biases in economic decisions**

2 Shira M. Lupkin^{1,2} & Vincent B. McGinty^{1,2}

3 1. Center for Molecular and Behavioral Neuroscience, Rutgers University - Newark

4 2. Behavioral and Neural Sciences Graduate Program, Rutgers University - Newark

5

6 Correspondence: Vincent B. McGinty, vince.mcginity@rutgers.edu, 197 University Ave,

7 Newark, NJ 07102

8

9 *Acknowledgements:* W.T. Newsome for funding, material support, and thoughtful discussions; A.

10 Rangel, F. Molter, G. Rosenbaum, G. Karpov, E. Murray, D. Sharma, & A. Thomas for

11 thoughtful discussions and comments on the manuscript; J. Brown, E. Carson, S. Fong, A.

12 McCormick, M. Ortiz, J. Powell, J. Sanders, & D. Siegel for technical assistance. R. Kiani for

13 creation of the ROME toolbox. F. Molter and A. Thomas for creation of the *glambox* toolbox.

14

15 *Funding:* This work was supported by the Howard Hughes Medical Institute (W.T. Newsome),

16 National Institutes of Health Grant K01-DA-036659-01 (V.B.M.), the Busch Biomedical

17 Foundation (V.B.M.), and the Whitehall Foundation (V.B.M.).

18

19 *Author contributions:* VBM designed the study; SML and VBM performed the research. SML

20 and VBM analyzed the data. SML and VBM wrote the paper.

21

22

Abstract

In economic decision-making individuals choose between items on the basis of their perceived value. For both humans and nonhuman primates, these decisions are often carried out while shifting gaze between the available options. Recent studies in humans suggest that these shifts in gaze actively influence choice, manifesting as a bias in favor of the items that are viewed first, viewed last, or viewed for the overall longest duration in a given trial. This suggests the existence of a mechanism that links gaze behavior to the neural computations underlying value-based choices. In order to identify this mechanism, it is first necessary to develop and validate a suitable animal model of this behavior. To this end, we have created a novel value-based choice task for macaque monkeys that captures the essential features of the human paradigms in which gaze biases have been observed. Using this task, we identified gaze biases in the monkeys that were both qualitatively and quantitatively similar to those in humans. In addition, the monkeys' gaze biases were well-explained using a sequential sampling model framework previously used to describe gaze biases in humans—the first time this framework has been used to explore value-based choice mechanisms in nonhuman primates. Together, these findings suggest a common mechanism that can explain gaze-related choice biases across species, and open the way for mechanistic studies to identify the neural origins of this behavior.

47 Introduction

48 Economic decisions are ubiquitous in the natural world, and in natural settings humans
 49 and other primates often evaluate the available options by moving their eyes, rapidly shifting the
 50 focus of their gaze between the options as they deliberate. Over the past decade, several studies
 51 in humans have documented an association between gaze and value-based choices. In particular,
 52 when offered a choice between multiple items, people are more likely to select the first item that
 53 they view, the item they look at just prior to indicating their decision, and the item that they
 54 spend the longest time viewing over the course of their deliberation. Further, choice biases can
 55 be induced through manipulations of gaze (Gwinn et al., 2019; Liu et al., 2020, 2021; Pärnamets
 56 et al., 2015; Sui et al., 2020; Tavares et al., 2017). Together, these studies suggest that gaze plays
 57 an active role in the decision process, such that simply viewing a given option makes us more
 58 likely to choose it (see Krajbich, 2019, for a review).

59 However, despite increased interest in the role of gaze in decision-making, knowledge of
 60 the underlying neural mechanisms is still limited (c.f., Krajbich et al., 2021; Lim et al., 2011).
 61 One critical barrier to identifying these mechanisms is the lack of a suitable animal model of this
 62 behavior, and consequently a lack of studies addressing neural mechanisms at the cellular level.
 63 To this end, the aim of the present study was two-fold: first, to develop a novel behavioral
 64 paradigm for macaque monkeys that captures the essential features of tasks previously used in
 65 humans; and second, to determine whether monkeys exhibit gaze-based choice biases that are
 66 similar to those of humans.

67 While the neural mechanisms linking gaze and economic choice are not fully known, the
 68 underlying computations have been explored using sequential sampling models, a highly
 69 successful class of models used to explain simple decision behaviors (Ratcliff & McKoon,

2008). In these models, a stream of noisy evidence for each option is accumulated until one item accrues enough evidence to reach a pre-defined threshold, and a decision is rendered in favor of that item. For value-based tasks, the evidence accumulation rate can be formulated as a function of the relative value of the possible outcomes—the larger the difference in relative value, the faster that evidence accumulates for the higher-valued item. Over the last decade, the basic sequential sampling model framework has been extended to account for the influence of gaze behavior (i.e., overt attention) on the decision process. These models, collectively referred to as attentional sequential sampling models (aSSMs), assume that evidence for a given item accumulates at a faster rate while that item is being viewed, and accumulates slower when it is not being viewed (Krajibich et al., 2010; Krajibich & Rangel, 2011; Smith & Krajibich, 2018; Tavares et al., 2017; Thomas et al., 2019). Overall, these models have been successful in accounting for human gaze biases across several decision-making domains (Krajibich, 2019; Smith & Krajibich, 2018; Thomas et al., 2019), supporting the idea that gaze modulates the decision process through a dynamic, gaze-dependent reweighting of choice options.

To identify candidate neural substrates, intracranial recordings in nonhuman primates (NHPs) are an ideal approach, as they combine an animal model with human-like oculomotor behavior with the high temporal resolution necessary to accommodate the rapid pace of natural eye movements (Kravitz et al., 2013; Öngür & Price, 2000; Wise, 2008).

Moreover, NHPs are commonly used as animal models in neuroeconomics. However, existing paradigms contain elements that make them unsuitable for testing the relationship between gaze and choice behavior. More specifically, many NHP studies restrict eye movements and/or use eye movements as a means of reporting choice, rather than for inspecting the choice options, (c.f., Cavanagh et al., 2019; Hunt et al., 2018; Rich & Wallis, 2016). These common

restrictions preclude natural gaze patterns and therefore the ability to identify the relationship between gaze behavior and choice.

To address this, we have developed a novel behavioral paradigm in which monkeys are able to freely deploy their gaze to the targets on the screen and to indicate their choice by a manual response. Using this paradigm, we replicated the three core behavioral signatures of gaze-related choice biases. We found that the monkeys were more likely to choose the first item that they viewed, the last item that they viewed before deciding, and the items that received more overall gaze time. In addition, we use a computational model previously used to assess human gaze biases (Thomas et al., 2019) to determine whether similar computational mechanisms can explain choices in NHPs.

Results

Basic Psychometrics

We trained two monkeys to perform a value-based decision task in which they used eye movements to view the choice options and used a manual response lever to indicate their decision (Figure 1A). To begin a trial, the monkeys were required to fixate on the center of the task display while holding down the center of three response levers. After the required eye fixation and lever hold period was satisfied, two targets were presented on the left and right of the task display, each associated with a juice reward ranging from 1 to 5 drops of juice. Following the initial fixation period, gaze was unrestricted, allowing the monkeys to look at each target as many times and for as long as they liked before indicating a choice. To encourage the monkeys to look directly at the targets, and to limit their ability to perceive the targets using peripheral vision, the targets were initially masked until the first saccade away from the central

fixation point was detected (see *Methods*). To indicate their choice, the monkeys first lifted their hand off the centrally-located response lever, and then pressed the right or left lever to indicate which of the two targets they chose. The reaction time (RT) in each trial was defined as the period between the onset of the targets and the lift of the center lever.

Overall choice performance is illustrated in Figure 1D. Choices were nearly optimal for the easiest trials (value difference of +/- 3 or 4 drops of juice, > 98% of choices in favor of higher-value target) and occasionally suboptimal for more difficult trials (value difference of +/- 1 or 2 drops of juice, 87% higher-value choices for Monkey C, 86% higher-value choices for Monkey K). When the targets were of equal value, the probability of the monkeys choosing the left or right target was near chance (48% left choices for Monkey C, 46% for Monkey K). This pattern of choices indicates that the monkeys were behaving according to the incentive structure of the task, with a very low rate of guesses or lapses (Fetsch, 2016). Further, choice behavior was well-described by a logistic function parameterized by the difference between the values of the left and right items (fit lines in Figure 1D are from a logit mixed effects regression: $\beta_{\text{value-difference}} = 1.36$, CI = [1.32, 1.39]; $F(1,30355) = 8466.20$, $p < 1e-10$).

The average RT was modulated by choice difficulty: the monkeys responded slower when the two targets were near in value, and faster when the value difference was large (Figure 1E, linear mixed effects regression: $\beta_{\text{difficulty}} = -0.035$, CI: [-0.040, -0.030]; $F(1, 30355) = 170.97$, $p < 1e-10$). These results conform with patterns seen in both humans and NHPs in two-alternative forced choice tasks (e.g., Gold & Shadlen, 2007; Thomas et al., 2019). Taken together, the choices and reaction times are consistent with a deliberative process dependent upon the target values.

Figure 1. Decision-making task and performance in two monkeys. **A.** Abbreviated task sequence; see Figure S1 for full task sequence. The interval between target onset and center lever lift defines the decision reaction time (RT). The yellow and blue glyphs are choice targets, and the gray “+” shapes indicate the location of visual crowders designed to obscure the targets until they are viewed (fixated) directly. For clarity, in this panel crowders are shown only in gray and at a reduced scale; on the actual task display, crowders were multicolored and the same size as the targets, as in panel C. **B.** An example set of 12 choice targets, organized into 5 groups corresponding to the 5 levels of juice reward. **C.** Close-up view of a single target array, consisting of a central yellow choice target surrounded by six non-task-relevant visual crowders. Two such arrays appear on the display (panel A), each located 7.5° from the display center. **D-F.** Task performance. **D.** Fraction of left choices as a function of the left minus right target value (in units of juice drops). **E.** Reaction Time (RT) decreases as a function of difficulty, defined as the absolute difference between target values; **F.** Number of fixations per trial decrease as a function of difficulty. Filled circles show the mean values and smooth lines show the logistic (D) or linear (E-F) regression fits derived from mixed effects models (Table 1, rows 1, 3, and 4, respectively). Error bars are too small to be plotted. The data shown in blue are from Monkey C, and the red from Monkey K. A subset of the data in panels D-F were reported previously in McGinty & Lupkin (2021), and the illustrations in panels A-C were adapted from McGinty & Lupkin (2021).

Gaze Behavior

We characterized the fixations that the monkeys performed while viewing the targets before indicating a choice in each trial. A fixation was defined as a period of stable gaze upon one of the two targets. Consistent with fixation analyses in human studies, consecutive fixations upon the same target were merged into a single fixation epoch, and fixations onto non-target locations were not included in the analysis.

Basic Fixation Properties

On average, the monkeys made more than 2 fixations per trial (mean $\pm$ SEM across trials; Monkey C = 2.6 ± 0.01 fixations, Monkey K = 2.4 ± 0.01 fixations). Practically, this means that on the vast majority of trials, the monkeys viewed both targets before choosing (97.2% of trials for Monkey C, 92.6% for Monkey K). As in human studies, the mean number of fixations per trial increased as a function of trial difficulty (Figure 1F, linear mixed effects regression: $\beta_{\text{difficulty}} = -0.11$, CI = $[-0.13, -0.093]$; $F(1, 30355) = 165.11$, $p < 1e-10$).

Fixation Durations Depend on Target Value

The average fixation duration was dependent on the value of the target being fixated, and upon the order in the trial in which the fixation occurred. For the first fixation in each trial, duration increased as a function of fixated target value ($\beta_{\text{fixValue}} = 0.015$, CI = $[0.0029, 0.027]$; $F(1, 30355) = 5.89$, $p = 0.015$; Table 1, row 5). For middle fixations, defined as any fixation that was neither first nor final, the average duration decreased as a function of the value of the fixated target ($\beta_{\text{fixValue}} = -0.0079$, CI = $[-0.0088, -0.0069]$; $F(1, 18083) = 259.94$, $p < 1e-10$; Table 1, row 6), and also as a function of overall trial difficulty ($\beta_{\text{difficulty}} = -0.0020$, CI = $[-0.025, -0.015]$; $F(1, 18083) = 63.88$, $p < 1e-10$; Table 1, row 7). At the same time, middle fixation durations

increased as a function of relative value, defined here as the value of the fixated target minus the value of the non-fixated target ($\beta_{\text{relativeFixVal}} = 0.017$, CI = [0.013, 0.022]; $F(1, 18083) = 56.47$, $p < 1e-10$; Table 1, row 8). Together, these results support the notion that fixation behavior and valuation are not independent processes, and that the effects of value on viewing durations must be accounted for when computing and interpreting gaze biases (see *Cumulative Gaze-Time Bias*, below).

Final Fixations Are Prolonged

Additionally, we quantified the relative durations of the first, middle, and final fixations. First fixations were defined as the initial fixations following the targets onset; final fixations were defined as the fixation coincident with the initiation of choice (center lever lift); and middle fixations were defined as neither first nor final. Following conventions from human studies, final fixation durations are defined as the interval between fixation onset and the center lever lift; in other words, final fixation durations are truncated at the reaction time in each trial (Krajbich et al., 2010). For non-final fixations, the duration is defined in the conventional way, as the interval between fixation onset and the initiation of the next saccade.

Similar to human fixation patterns, first fixations were the shortest, (Figure 2C, light-shaded bars). However, unlike in humans, final fixations were longer than middle fixations (Figure 2C, medium- vs. dark-shaded bars). Because final fixations by definition occur latest in a trial, we confirmed that the duration difference was not attributable to absolute fixation order (1st 2nd, 3rd, etc.). To do this, we compared final to middle fixations looking only among fixations that occurred second in the trial, which were equally split between middle and final fixations. As illustrated in Figure 2D, when second fixations were the final fixations in a trial they were, on average, 200ms longer than when the second fixations were in the middle of the trial.

189 Thus, the final fixations in a trial are significantly prolonged when compared to non-final
 190 fixations of the same absolute order. This is a departure from human studies, where final
 191 fixations are consistently shorter than middle fixations (e.g., Krajbich et al., 2010; Krajbich &
 192 Rangel, 2011; Tavares et al., 2017). In aSSM models, shorter final fixations are explained as
 193 fixations that are “interrupted” by the accumulator reaching a bound before the fixations would
 194 have been terminated naturally by an eye movement to the other target. Clearly, the monkeys’
 195 fixation patterns require an alternative explanatory mechanism, which we detail below.

196

Figure 2. Fixation durations. **A.** Illustration of a trial with three fixations. The second fixation, highlighted in red, is considered a “middle” fixation because it is terminated by the initiation of a saccade to the other target. **B.** Illustration of a trial in which the second fixation is a “final” fixation, terminated at the moment the monkey initiates choice by lifting the center lever. **C.** Relative durations of all first, middle, and final fixations. Bar heights are the mean fixation durations and the error bars are the session-wise standard errors of the mean. **D.** Distributions of naturally terminated (gray) and choice terminated (black) fixations, considering only the fixations occurring second in the trial. Naturally terminated fixations: Monkey C: $n = 1545$, Monkey K: $n = 1281$; choice-terminated fixations: Monkey C: $n = 764$, Monkey K: $n = 792$.

Estimation and Truncation of Terminal Non-Decision Time (NDT).

One key observation that may help explain prolonged final fixations is the fact that final fixations are almost always on the item that is chosen by the monkey (95.8% of trials for Monkey C, 99.1% for Monkey K). In other words, immediately before a choice is physically initiated at the end of a trial, the monkeys appear to devote extra fixation time towards the item that will be chosen. Two features of primate reaching behavior may explain this phenomenon. First, skilled reaches typically require non-trivial planning time, which introduces a latency between the time that the subject commits to the movement and the time that the movement is actually detected by the experimenter; across a range of experimental preparations, this latency is estimated to be approximately 200ms (Arora et al., 2019; Cisek & Kalaska, 2005; Resulaj et al., 2009; Shenoy et al., 2011). Second, during this preparatory period, humans and NHPs exhibit “gaze anchoring”, a behavior in which the gaze becomes fixed upon the intended target of an upcoming reach until the reach has been completed (Neggers & Bekkering, 2001, 2000).

Taken together, this suggests a simple explanation for the surplus fixation upon the to-be-chosen items: it is the result of a *post-decision* gaze anchoring process. In other words, gaze becomes drawn to the chosen item for the interval of time after an internal decision has been made, but before the decision is physically reported.

Within the framework of sequential sampling models, this post-decision/pre-report interval is referred to as terminal non-decision time (NDT) and is usually attributed to motor preparation – in our case, preparation of either a left or right reach of the arm. Because the goal of this study was to characterize only the *pre-decision* relationship between gaze and decision outcome, it was necessary to minimize the influence of post-decision fixation behavior that occurs during the NDT. To this end we obtained for each monkey a bootstrapped distribution for

the median NDT, using the differences in duration between the longer final fixations and the shorter middle fixations (see Figure 2D and *Methods*). We took the upper 95th percentile of this distribution as a point estimate of the NDT for each monkey, and then truncated the gaze data in every trial by this amount. For Monkey C, the estimate was 196.7ms, and for Monkey K was 200.1ms. For all the analyses below examining gaze-related choice biases, as well as for the aSSM model fitting, the data from every trial were truncated by these amounts, thereby minimizing the contribution of post-decision fixation behavior.

Gaze-Related Choice Biases

Humans show three core choice biases related to gaze behavior. First, they tend to choose the item that they viewed first (initial fixation bias). Second, they tend to choose the item they spend more time looking at over the course of the trial (cumulative gaze-time bias). Last, they tend to choose in favor of the target they are viewing at the moment they indicate their choice. All three of these biases are evident in our monkey subjects, as detailed below.

Initial Fixation Bias

As is the case in humans, the monkeys were more likely to choose the target that they looked at first in a given trial. To quantify this bias, we used a logit mixed effects regression, which gives an estimate of the effect of an initial leftward fixation on the log odds of a leftward choice: $\beta_{\text{first-is-left}} = 0.64$, CI = [0.29; 1.0]; $F(1,30348) = 12.81$, $p = 3.44\text{e-}04$, see Table 1, row 10). The fact that $\beta_{\text{first-is-left}}$ is significantly greater than zero indicates that an initial leftward fixation is associated with a greater likelihood of a leftward choice. The effect of target value differences computed in this same model was $\beta_{\text{value-difference}} = 1.39$ (CI = [1.35, 1.42]; $F(1,30348) = 5643$). The ratio between these two regression estimates ($\beta_{\text{first-is-left}} / \beta_{\text{value-difference}}$), gives the

relative influence the initial fixation on choice as compared to the value differences of the targets: $0.64/1.39 = 0.47$. In other words, the monkeys chose as if the first-fixated item was, on average, worth 0.47 drops of juice more than its nominal value. Critically, because the targets were obscured until the monkeys initiated their first saccade (see *Methods*), their initial fixations were unbiased with regard to target value. Thus, the choice bias in favor of the initially fixated item cannot be attributed to a tendency for the monkeys to look at higher value targets first.

To test the causal role of the initial fixation on choice, we conducted additional experimental sessions in which we manipulated initial fixation by staggering the onset of the two targets (Monkey C: 19 sessions, $n = 6931$ trials; Monkey K: 9 sessions, $n = 5918$ trials). Staggered target onsets – either left first or right first – occurred randomly in 30% of the trials in these sessions (Figure 3A), with concurrent onsets in the other 70% of trials (Figure 3B). In the staggered-onset trials, the monkeys almost always directed their initial fixations to the target that appeared first (100% of staggered trials for Monkey C, 92% for Monkey K). As in the sessions without gaze manipulation described above, we found a main effect of initial fixation on choice (logit mixed effects regression: $\beta_{\text{first-is-left}} = 0.66$, $\text{CI} = [0.42; 0.90]$; $F(1,12828) = 28.509$, $p = 9.48\text{e-}08$; $\beta_{\text{first-is-left}} / \beta_{\text{value-difference}} = 0.55$ drops of juice; see Figure 3 and Table 1, row 12). Importantly, we found no difference in the magnitude of the initial fixation effect between the staggered-onset and standard trials $\beta_{\text{first-is-left:condition}} = -0.08$, $\text{CI} = [-0.29; 0.14]$; $F(1,12828) = 0.4823$, $p = 0.49$). Thus, by manipulating the direction of the initial fixation, it was possible to manipulate choice behavior. This suggests that the initial fixation direction, whether exogenously or endogenously driven, has a causal influence in choice.

Figure 3. Initial gaze bias for gaze-manipulation sessions. **A.** Depiction of gaze manipulation procedure. In 30% of trials, only a single target was initially presented, randomly assigned to the left or right of the display (left panel). Once a saccade was detected (or 250ms elapsed, whichever happened first), the second target appeared (right panel). **B.** Depiction of standard trials with masked targets appearing simultaneously on both the left *and* right of the display (left panel). Masks disappeared when the initial saccade was detected (right panel). **C-F.** Probability of choosing left as a function of the difference between the value of the item on the left and the value of the item on the right. Data were split according to the direction of the initial fixation. Blue and red indicate left and right initial fixation, respectively. Circles show the mean probabilities of left choice, error bars show the standard error of the mean over 19 sessions for Monkey C and 9 sessions for Monkey K. Lines show logistic fits from the mixed effects model (Table 1, row 12). **C & E.** show data from trials where the onset of the first target was staggered. **D & F.** show trials where the targets appeared simultaneously. **C-D.** shows data from Monkey C; **E-F.** shows data from Monkey K.

Cumulative Gaze-Time Bias

Here, we asked whether monkeys were more likely to choose items that were fixated longer in a trial. We first quantified the time spent looking at the two choice targets in every trial; critically, this was done using data that were truncated at the end to minimize the inclusion of NDT. Then, using a logit mixed effects regression (Table 1, row 2), we asked whether choice outcomes were dependent on the relative time spent looking at the left item as compared to the right item (“Time Advantage Left”). Consistent with human choice patterns, monkeys were more likely to choose the left item as the relative time advantage for the left item increased, and likewise more likely to choose right as the time advantage for the right item increased (Figure 4B,C; $\beta_{\text{time-advantage-left}} = 5.75$, CI = [5.24; 6.32]; $F(1,28691) = 442.22$, $p < 1 \text{ e-} 10$, Table 1 row 2). Because fixation durations are themselves dependent upon the target values (see *Basic fixation properties*, above), we repeated this analysis using choice probabilities that were corrected to account for the target values (see *Methods*). Using these corrected choice probabilities, the positive relationship between relative viewing times and choice was maintained. (Figure 4D,C, bottom row; linear mixed effects regression: $\beta_{\text{time-advantage}} = 0.49$, CI = [0.40; 0.58]; $F(1,28691) = 107.68$, $p < 1 \text{ e-} 10$, Table 1, row 9). Thus, the monkeys were more likely to choose items that were fixated longer, independent of the values of the targets in a given trial.

Figure 4. Cumulative gaze-time bias. **A.** Task schematic. The portion of the trial used to show the cumulative gaze-time bias is highlighted in red. **B-C.** Probability of choosing left as a function of the binned final gaze-bias advantage for the left item. **D-E.** The same as B-C, but using choice probabilities corrected to account for the target values in each trial (see *Methods*). For visualization, these graphs exclude those trials with a final time advantage greater than ± 3 standard deviations from the mean ($< 1\%$ of trials for each monkey). The trials were not excluded from the logit mixed effects regression reported in the text. In B-E, the bin boundaries were ± 0.05 , ± 0.15 , and $\pm \infty$. Error bars indicate SEM across sessions.

Final Fixation Bias

One key prediction of the aSSM framework is that the item being fixated at the end of the trial should be chosen more often than the non-fixated item. This final fixation bias is explained by the increased drift rate for fixated items: all else being equal, evidence in favor of the fixated item will accrue faster and will be more likely to reach the decision boundary. Hence, chosen items are more likely to be the items being fixated at the moment a decision is initiated (Krajbich et al., 2010; Krajbich & Rangel, 2011; Tavares et al., 2017).

We assessed this final fixation bias in our data in a manner similar to prior studies but with one important difference: as detailed above, the fixation data in each trial were truncated by ~200ms to minimize the influence of gaze anchoring behavior. Thus, in our data the location of the final fixation in each trial is defined as the target being viewed by the monkey at the *virtual* end of the trial, which in an aSSM framework can be interpreted as an estimate of when the decision boundary was reached (e.g., Resulaj et al., 2009). Using the truncated data, we found that the monkeys were slightly more likely to choose the target being viewed at the virtual trial end, consistent with the final fixation bias effects observed in humans (Figure 5, logit mixed effects regression: $\beta_{\text{last-is-left}} = 0.71$, CI = [0.19; 1.22]; $F(1,28690) = 7.27$, $p = 7.02\text{e-}03$).

303

Figure 5. Final Fixation Bias. **A.** Task schematic highlighting the portion of the trial used to show the final fixation bias (red box). **B-C.** Probability of choosing left as a function of the difference in item values, split according to the location of gaze (left or right) at the end of the trial. Circles show mean probabilities of left choice, error bars show the standard error of the mean over 25 sessions for Monkey C and 29 sessions for Monkey K. Lines show logistic fits from the mixed effects model (Table 1, row 11). Panel B shows data from Monkey C; panel C shows data from Monkey K.

304

Gaze-Weighted Linear Accumulator Model (GLAM)

The analyses above show broad consistency between NHP and human data with respect to the association between gaze and choice behavior. The question remains, however, as to whether these patterns emerge from a similar mechanism to those observed in humans. To this end, we asked whether an aSSM framework could explain the gaze and choice patterns we observed in NHPs. We chose the Gaze-weighted Linear Accumulator Model (GLAM; Molter et al., 2019), which allows evidence accumulation for the choice items to be modified by gaze behavior. Previously, the GLAM model was used to explain gaze-related choice biases in humans across multiple independent decision studies (Thomas et al., 2019). Using this model, we asked three questions with regards to our NHP data: First, does a model with a gaze bias mechanism fit our data better than a similar model without a gaze bias? Second, how do gaze biases estimated in our monkey subjects compare to prior human work? Third, can the model accurately predict the critical features of choice behavior, such as accuracy, RT, and the influence of gaze-on choice?

To answer the first question, we fit to each monkeys' data two variants of the GLAM that differed in their treatment of the gaze bias parameter, γ (see Equation 1 for model specification and parameter definitions). In the first GLAM variant, the gaze bias parameter γ was left as a free parameter to be fit, allowing gaze to modify the rate of evidence accumulation for fixated items. The second model fixed γ at 1, assuming *a priori* that gaze has no effect on the decision process. The maximum a posteriori (MAP) and 97.5 highest posterior density intervals (HPD) for the parameters estimated in both model variants are given in Table S1. The relative fit of the two variants was assessed using the difference in WAIC score (dWAIC, Vehtari et al., 2017), defined by subtracting the WAIC value for the model with γ fixed at 1 from the WAIC for the model

where γ was a free parameter. The dWAIC for Monkey C was -3106, and for Monkey K was -6055, indicating that the model that allowed for the presence of gaze biases is the better fit of the two model variants. In the gaze-bias model, γ values were as follows: Monkey C: MAP=0.28, HPD=[0.26,0.3]; Monkey K: MAP = 0.17, HPD = [0.15,0.19].

Since the GLAM had previously been used to explore gaze biases in four human decision-making datasets (Folke et al., 2016; Krajbich et al., 2010; Krajbich & Rangel, 2011; Tavares et al., 2017), we were able to directly compare the extent of the gaze bias, as indicated by γ , between humans and our monkey subjects. Treating each session as an independent sample, we found that the distribution of γ in each of our monkey subjects fully overlapped with the γ distributions from the human studies (Figure 6).

Figure 6. Distribution of gaze bias parameter (γ) across species. X-axis shows the range of possible values of the gaze bias parameter, γ . The gray distribution combines the γ values estimated for each subject in the four human datasets previously fit using the GLAM (Thomas et al., 2019). The means for each dataset are given as triangles above the distributions. The teal and magenta distributions show the session-wise distribution of γ parameters for Monkeys C, and K, respectively. Triangles in the corresponding colors show the means of the monkeys' distributions. Total number of subjects for each dataset are as follows: Folke, 2016 — N = 24 ; Tavares, 2017—N = 25; Krajbich, 2011—N = 30; Krajbich, 2010—N = 39. Total number of sessions for the monkeys were 29 and 25 for Monkeys C and K, respectively.

338

339

340

341

342

343

Finally, we assessed the model's predictive accuracy by performing out-of-sample simulations to predict core features of the behavioral data. To do this, the data were first split into even- and odd-numbered trials. Next, the models were estimated using only the even-numbered trials. Finally, these estimates were used to simulate each of the held-out trials 10 times, resulting in a set of simulated choice and reaction time measures that could be compared to the empirically observed behavior in the held out trials. As illustrated in Figure 7, the model predictions closely match the monkeys' actual performance (see also Table S2). Note that repeating this procedure with γ fixed at 1 (no gaze bias) produces accurate predictions of choices and reaction times but fails to capture the effect of cumulative gaze-bias on choice (not illustrated, see Thomas et al., 2019). Thus, an aSSM model that includes gaze bias provides a better fit than a model without a gaze bias, and also makes accurate out-of-sample predictions of choice and gaze behavior. Further, the magnitude of the gaze bias is comparable to gaze biases estimated in humans using the same model framework.

Figure 7. Out-of-sample predictions of the GLAM model for choice and gaze behavior. White bars show the empirically observed behavioral data from monkeys in odd-numbered trials. Blue dots and lines show GLAM model predictions for odd-numbered trials. Error bars show the standard error of the mean across sessions (SEM). **A,D.** Fraction of left choices as a function of the left minus right target value (in units of juice drops). **B,E.** Reaction time decreases as a function of difficulty; note that reaction times are shorter here than in Figure 1E, due to the use of NDT-truncated data. **C,F.** Corrected probability of choosing the left target as a function of the fraction of time spent looking at the left target minus fraction of time spent looking at the right target. The bin boundaries are 0, ± 0.33 , ± 0.67 , and ± 1 . Panels A-C correspond to Monkey C; panels D-F correspond to Monkey K.

Discussion

Over the last decade, there has been a growing appreciation for the active role that attention plays in a broad array of decision-making domains. However, the neural substrates of this relationship remain largely unknown due to the lack of a suitable animal model. The aim of the present study was to develop such a model using a novel value-based decision-making task for NHPs. We validated our animal model on two measures: behavioral control and the presence of human-comparable gaze biases. First, the monkeys' performance indicated good behavioral control: both choices and RTs were a graded function of task difficulty (defined by the difference in target values), with fast and nearly optimal choices in the easiest trials, and slower and more variable choices in harder trials. Such signatures of behavioral control are crucial to the development of any animal model of complex cognitive processes, as they allow the confident attribution of an animal's behavior to the construct of interest, and not to attentional lapses or guesses (Fetsch, 2016; Krakauer et al., 2017). For the present study, this suggests that the monkeys' behavior is ascribable to a decision process driven primarily by the values of the targets in each trial. As to the second criterion, our results show that the monkeys exhibit similar gaze biases to those observed humans, both in terms of the core behavioral measures and also in their ability to be mechanistically described using an aSSM. Taken together, this novel task is well suited for use in neural mechanistic studies, such as concurrent electrophysiological recordings or stimulation.

Comparison to Prior Studies

As noted, our data provide evidence that gaze has a similar influence on value-based choice in monkeys as it does in humans. However, there was one notable difference between our

findings and prior human work. Specifically, we found that the durations of final fixations were significantly longer than those of middle fixations, the opposite of what has been observed in humans. We propose that this discrepancy is not the result of a species-level difference, but instead can be attributed differences in the motor response modality in human tasks as compared to the current study. Human subjects typically use button boxes or keyboards, often with fingers resting on top of the response keys throughout the experiment; in contrast, the monkeys in our study were required to make a speeded full-arm reach between the initial (center) lever and the levers on either side to indicate their chosen target. When making targeted reaches, both humans and nonhuman primates tend to “anchor” their gaze onto the reach target in order to increase the speed and efficiency of the reach (Battaglia-Mayer et al., 2001; Dean et al., 2011; Neggers & Bekkering, 2000). In addition, this kind of reach typically requires a non-trivial movement planning time. Key presses, in contrast, do not require significant movement planning or hand-eye coordination, especially for subjects accustomed to pressing keys without visual guidance. Therefore, the most parsimonious explanation for the monkeys’ prolonged final fixations is that they reflect gaze anchoring during the post-decision movement planning interval (NDT) immediately prior to the initiation of the reach, resulting in increased fixation time on the to-be-chosen item.

To account for and minimize the influence of gaze anchoring, we derived an empirical estimate of the NDT for each monkey and truncated each trial using these values. In line with our hypothesis, our estimates were within the upper limits of the motor delays induced by gaze anchoring, i.e. approximately 200ms. Further, these values also correspond to the premotor latency associated with reaching movements, as recorded in the NHP premotor cortex (Churchland, 2006; Cisek & Kalaska, 2005). Finally, the truncation of our data did not impact

our ability to detect gaze biases nor to fit the data using an aSSM, which supports our hypothesis that the elongated final fixations are the result of post-decisional, NDT-like, process.

Beyond similarities to prior human studies, our results are also consistent with the handful of NHP studies in which some degree of free eye movement was permitted. Among these, both Hunt et al., 2018 and Rich & Wallis, 2016 imposed some restrictions on eye movements, making it difficult to assess the relationship between fully natural gaze patterns and choices. Cavanagh and colleagues (2019) used a fully free-viewing design, and also observed a bias in favor of initially fixated items. Our study confirms this result, but also significantly expands upon their conclusions, due to novel elements of our study design. First, by manipulating initial fixation direction, our study shows that initial fixation has a causal effect on the decision process. Second, our study requires the monkeys to sample the targets by viewing them foveally (rather than using peripheral vision), permitting us to accurately measure the time spent viewing each target and therefore demonstrate a clear relationship between relative viewing time and choice. Third, the results of the aSSM model-fitting provide a mechanistic explanation for our results, and permits a direct comparison to model fitting results in human studies.

Hypothesized Neural Substrates

The impetus for developing this animal model is to explore the neural underpinnings of the relationship between gaze and value-based choice. Here, we have shown that decision behavior is well fit by a sequential-sampling model framework, suggesting neural mechanisms that exhibit evidence accumulation dynamics. Although this is the first instance of a sequential sampling model being applied to monkeys making value-based decisions, there is ample

evidence for accumulator-like mechanisms in monkeys making perceptual decisions, with studies of dot motion discrimination perhaps being the most well-known (see Gold & Shadlen, 2007, for review). These studies identified the critical role of the extrastriate visual area MT for encoding instantaneous motion evidence, and the role of parietal oculomotor area LIP for encoding accumulated evidence. Generalizing from these studies, a natural hypothesis is that similar neural substrates may be identified for value-based decisions, i.e. that there may be distinct regions for encoding instantaneous and accumulated value evidence.

For the encoding of accumulated evidence, several findings point to regions involved in the preparation or execution of the motor response. For example, in NHP dot motion studies eye movements are typically used to report decisions, and accumulated evidence signals can be observed in oculomotor control regions such as area LIP, frontal eye fields, and the superior colliculus. This also is consistent with a recent report by Krajbich et al. (2021) showing that focal disruption of human frontal eye field activity reduces the effect of gaze on value-based choice and increases decision reaction times. Because the decision in our task is reported with a reach movement, likely sites for accumulator-like activity include motor cortical areas such as the arm representations within dorsal pre-motor cortex (Chandrasekaran et al., 2017).

Because decision behavior in this study was value-dependent, a natural hypothesis is that regions encoding economic value provide the input to the accumulator. Value-related signals have been observed in numerous cortical and subcortical regions, regions including, but not limited to, the amygdala (e.g., Paton et al., 2006), the ventral striatum (Kable & Glimcher, Lim), and the ventromedial frontal lobe (VMF), including both the orbitofrontal and ventromedial prefrontal cortices (Padoa-Schioppa & Assad, 2006; Rich & Wallis, 2016; Vaidya & Fellows, 2015).

Of these regions, the VMF has received the most scrutiny with regards to attentional biases. In humans, there is evidence of BOLD activity in this region related to attention-guided decision-making (Leong et al., 2017; Lim et al., 2011), and patients with lesions to this area show impairments in such tasks (Vaidya & Fellows, 2015). Evidence from NHP neurophysiology suggests a role for the orbitofrontal portion of the VMF, as it contains value signals that become amplified when gaze is directed towards a reward-associated visual target (Hunt et al., 2018; McGinty et al., 2016), a mechanism fully consistent with the effects of gaze posited by aSSM frameworks.

Conclusions

In real-world decision scenarios, both humans and NHPs depend heavily on active visual exploration to sample information in the environment and make optimal decisions. In this way, even simple decisions require dynamic sensory-motor coordination. The importance of these gaze dynamics is underscored by the growing number of human studies showing that gaze systematically biases the decision process — in particular, that decisions are biased in favor of items that are viewed first and/or longest in a given trial. At present, the cell-level neural mechanisms underlying this bias are unknown due to the lack of a suitable animal model and the resulting limitations on recording neural activity with sufficiently high spatiotemporal resolution. To fill this gap, we have developed a novel value-based decision-making task for macaque monkeys that allows them to freely view and select target stimuli in a self-paced manner. We validated this paradigm on two measures: behavioral control and the presence of human-comparable gaze biases. Thus, this novel paradigm strikes a balance between permitting the

expression of natural gaze behaviors while still maintaining a level of control suitable for rigorous neurophysiological investigations.

Methods

Subjects

The subjects were two adult male rhesus monkeys (*Macaca Mulatta*), referred to as Monkey C and Monkey K. Both subjects weighed between 13.5 and 15kg during data acquisition, and both were 11 years old at the start of the experiment. They were implanted with an orthopedic head restraint device under full surgical anesthesia using aseptic techniques and instruments, with analgesics and antibiotics given pre-, intra-, and post-operatively as appropriate. Data from Monkey K were collected at Stanford University (25 sessions, 14548 trials); data from Monkey C were collected at Rutgers University—Newark (29 sessions, 15809 trials). All procedures were in accordance with the Guide for the Care and Use of Laboratory Animals (2011) and were approved by the Institutional Animal Care and Use Committees of both Stanford University and Rutgers University—Newark.

Study Design and Apparatus

The subjects performed the task while head-restrained and seated in front of a fronto-parallel CRT monitor (120Hz refresh rate, 1024x768 resolution), placed approximately 57cm from the subjects' eyes. Eye position was monitored at 250 Hz using a non-invasive infrared eye-tracker (Eyelink CL, SR Research, Mississauga, Canada). Three response levers (ENV- 612M, Med Associates, Inc., St. Albans, VT) were placed in front of the subjects, within their reach. The “center” lever was located approximately 17cm below and 35cm in front of the display center, and the center points of the two other levers were approximately 9cm to the left and right of the center lever. The behavioral paradigm was run in Matlab (Mathworks, Inc., Natick, MA)

using the ROME toolbox (R. Kiani, 2012). Juice rewards were delivered through a sipper tube placed near the subjects' mouths.

Behavioral Task

Task Structure

Figure 1A shows an abbreviated task sequence; Figure S1A shows the full sequence as described here: Each trial began with a fixation point in the center of the screen. Upon appearance of the fixation point, the subjects were required to both depress the center lever and to maintain fixation within a radius of 3.5 degrees of visual angle from the center of the fixation point. The monkeys were required to use only one hand to perform the task (Monkey K used his right hand; Monkey C used his left). After the initial fixation and center lever press were both held for an interval between 1-1.5 seconds, the fixation point disappeared, and two target arrays were presented 7.5 degrees to the left and right of the fixation point. Each target array consisted of a choice target, an initial mask, and six non-relevant "visual crowders". The net effect of the mask and crowders was to obscure the target until the monkey looked directly at the target with a saccadic eye movement. See *Gaze Contingent Mask and Visual Crowders*, below, for details. The two targets shown in each trial were selected randomly without replacement from a set of 12 (see Figures 1B & S1C), and placed randomly on the left or right side of the display, giving 132 unique trial conditions.

Once the target arrays appeared, the monkeys were free to move their gaze as they desired and were free to initiate a choice at any time. To initiate a choice, the monkeys lifted their hand off the center lever, at which point the target arrays were extinguished. After releasing the center lever, the monkeys had 500ms to choose their desired target by pressing the left or right lever. After a delay (1s for Monkey K; 1 or 1.5s for Monkey C), the subjects received the

volume of juice reward associated with the chosen target (range 1-5 drops). If neither the left nor right lever was pressed within 500ms of the center lever lift, the trial ended in an error. Each trial was followed by an inter-trial interval randomly drawn from a uniform distribution between 2 and 4 seconds. Eye movements following the initial fixation were monitored but not enforced, and had no programmatic influence on the outcome of the trial. Reaction time (RT) was calculated as the time between target arrays onset and the release of the center lever.

Choice Targets

Choice targets were constructed as composites of two features, color and shape, and the reward value of a given target (range 1-5 drops) was determined by the sum of the reward values associated with its two features. (Figure S1C). Notably, this scheme allows for multiple visually distinct targets to share the same value. An example target set is shown in Figure 1B and Figure S1C. For a given set, the colors were selected by randomly sampling four equally-spaced hues from a standardized color wheel based on the CIELUV color space, and shapes were quasi-randomly selected from a library of 28 equal-area shapes.

New stimulus sets were generated frequently, to avoid over-training on a given stimulus set. For every new set, 1-2 training sessions (not used for analysis) were performed so that the monkeys could become familiar with the stimuli, followed by 1-3 sessions of data collection.

Gaze-Contingent Initial Mask and Visual Crowders

A key design goal of this task was to encourage the monkeys to fixate onto the choice targets directly using saccadic eye movements. To encourage this behavior, we used two methods to obscure the choice targets until they were fixated directly.

The first method was a gaze-contingent initial mask: when the target arrays first appeared in each trial (i.e. the first frame), the choice targets were not shown; rather, in their place were

two non-relevant, non-informative masking stimuli-drawn from a pool of visual crowder stimuli (described below). The masking stimuli remained until the monkeys' moved their eyes out of the initial fixation window by initiating a saccade, usually ~150ms after the onset of the arrays. Once the eyes left the initial window, the masking stimuli were replaced with the actual choice targets in the next display frame. Because the display frame rate (one frame every 8.25ms) was shorter than the typical time taken to complete a saccade (~30-40ms), the switch from the masks to the actual stimuli occurred while the monkeys' eyes were moving, so that the choice targets appeared on the display before the end of the saccade (See Figure S1A for a step-by-step illustration). Therefore, information about target values was only available after the monkeys initiated a saccade towards one of the two target arrays.

The second method was the use of 6 visual crowdors placed in close proximity in a hexagonal arrangement around each target (Figure 1C). These crowdors minimized the monkeys' ability to use peripheral vision to identify the choice targets (Crowder & Olson, 2015; Whitney & Levi, 2011), further encouraging them to make fixations directly onto the targets.

The initial mask stimuli and the visual crowdors were constructed using plus (+) or cross (x) shapes, colored using five hues randomly sampled from equally spaced points on a standardized color wheel. These stimuli convey no value information because the plus and cross shapes were never used to construct target stimulus sets, and because the use of multiple, random hues within each of the stimuli does not signify any one color that could indicate value.

The crowdors and initial target masking were highly effective at encouraging the monkeys to fixate both targets before initiating a choice: both targets were viewed in 97.2% of trials for Monkey C and 92.6% of trials for Monkey K. These methods also ensured that the first fixation following target onset was unbiased with respect to the target values: in trials where one

target had a larger value than the other, the monkeys directed their first saccade towards the higher value target 49.8% of the time, which was not significantly different from chance ($t(25782) = -0.7660, p = 0.4437$).

Gaze Manipulation

To test the causal nature of the relationship between initial fixation and choice, we conducted additional sessions in which we temporally staggered the onset of the two targets. (Monkey C: 19 sessions, $n = 6931$ trials; Monkey K: 9 sessions, $n = 5918$ trials). During these sessions, in 30% of the trials the first target array would appear randomly on either the left or right of the screen. The second target array appeared 250ms after the onset of the first, or when the monkeys began their initial saccades away from the central fixation point, whichever came first. The location of the first target array was randomized across trials. This manipulation was intended to encourage the monkeys to direct their attention towards the target that appeared first, as shown in prior human studies (Tavares et al., 2017). As expected, the monkeys directed their initial fixation to the cued target on 96% of the staggered-onset trials.

Behavioral Analysis

All analyses were performed in Matlab v. 2020a (Mathworks, Inc., Natick, MA). Unless otherwise noted, analyses were conducted using generalized mixed-effects regressions with random effects for monkey-specific slopes and intercepts. To fit these models, data were collapsed across sessions, unless otherwise noted. Full specifications for all models are given in Table 1. Regression estimates (β 's) and 95 % confidence intervals (CI) are reported for each effect. Significance for each effect was determined using a post-hoc ANOVA in which the regression estimates were compared to 0 (F - and p - values, reported in the text).

The onset and duration of fixations were recorded by the Eyelink CL software. For analysis purposes, a fixation was considered to be on a given target if it was located within 3-5 degrees of the target center. Two consecutive fixations made onto the same target were merged into one continuous fixation. Fixations that were not onto either target were discarded; however, these were extremely rare: in the vast majority of trials, the monkeys fixated exclusively upon the targets.

For analysis purposes, fixations were designated as “first”, “final” or “middle”. First fixations were defined as the first target that was fixated after the onset of the target arrays; final fixations were defined as the target being fixated at the moment the monkey initiates a choice by lifting the center lever; and middle fixations were defined as all on-target fixations that were neither first nor final. In addition, for the analysis in Figure 2 and for estimation of NDT, fixations were further designated according to their absolute order in the trial, e.g. “first”, “second”, etc.

For analyses of the final fixation bias and the cumulative gaze-time bias, trials containing only a single fixation were excluded (Monkey C: N = 447 trials (2.8% of total), Monkey K: N = 1217 trials (8.4% of total)), as they could artificially inflate either or both of these biases. For the latter analysis, we computed in every trial the time advantage for the left item by subtracting the total duration of fixations onto the right target from the total duration of fixations onto the left target (Krajbich et al., 2010). Thus, for trials where the monkey allocated more gaze time to the left target, the time advantage takes a positive value, whereas it takes a negative value for trials where the monkey allocated more gaze time to the right.

Finally, for the cumulative gaze bias analyses shown in Figure 4C and 4E, choice proportions were corrected using the same procedure set forth in Krajbich et al., 2010. For each

trial, we coded the choice outcome as 0 for a right choice and 1 for a left choice. Then from each trial's choice outcome we subtracted the average probability of choosing left for all trials in that same condition, where the condition was defined as the left minus right target values. This correction removes from the choice data any variance attributable to the differences in the target values; as a result, any correlation between cumulative gaze and choice cannot be attributed to the target values. Put another way, this procedure allows us to measure the correlation between cumulative gaze-bias and choice separately within each value difference condition, so that the overall correlation reflects the aggregate of these within-condition correlations. The net result is that Figure 4C and 4E reflect the effect of cumulative gaze after accounting for the influence of the target values on gaze durations.

626

Behavioral Effect	Response Variable (Y)	Regressors (X1, X2, ...Xn)	Equation	Figure
Choice Performance (1)	Choice (0 = left, 1 = right)	X: Left value - right value	$Y \sim \text{logit}(1+X + (1+X \mid \text{monkey}))$	1D
Cumulative Gaze-Bias (2)	Choice	X: Time advantage left (total left gaze time - total right gaze time)		3A/C
Reaction Time x Difficulty (3)	Reaction time	X: Difficulty (left value - right value)		1E
Number of fixations x Difficulty (4)	Number of fixations	X: Difficulty		1F
Initial Fixation Duration x Fixated Target Value (5)	Initial fixation duration	X: Fixated target value	$Y \sim 1+X + (1+X \mid \text{monkey})$	-
Middle Fixation Duration x Fixated Target Value (6)	Middle fixation duration	X: Fixated target value		-
Middle Fixation Duration x Relative Target Value (7)	Middle fixation duration	X: Fixated target value - unfixated target value		-
Middle Fixation Duration x Difficulty (8)	Middle fixation duration	X: Difficulty		-
Value-Corrected Cumulative Gaze-Bias (9)	Value-corrected choice probability	X: Time advantage left		3B/D
Initial Fixation Bias (10)	Choice	X1: Left value - right value X2: Location of initial fixation (0 = right, 1 = left)	$Y \sim \text{logit}(1+X1 + X2 (1+X1 +X2 \mid \text{monkey}))$	-
Final Fixation Bias (11)	Choice	X1: Left value - right value X2: Location of final fixation (0 = right, 1 = left)		4
Initial Fixation Bias with Gaze-Manipulation (12)	Choice	X1: Left value - right value X2: Location of initial fixation X3: Trial type (0 = simultaneous onset, 1 = staggered onset)	$Y \sim \text{logit}(1+X1 + X2*X3 (1+X1 +X2*X3 \mid \text{monkey}))$	2

627

629 **Table 1. Mixed effects model specifications**

Terminal Non-Decision Time (NDT) Estimation and Truncation

Rationale

As detailed in Results, the durations of the monkeys' final fixations were longer than non-final fixations; humans, in contrast, show the opposite pattern (Krajovich et al., 2010, Tavares et al., 2017). We argue that this discrepancy can be attributed to differences in the motor response used in our task as compared to studies in humans. Specifically, human subjects typically use a keyboard or finger-actuated button box, while our monkey subjects were required to make a speeded full-arm reach between the initial center lever and left or right lever to indicate choice. This distinction is germane because of the role of gaze in goal-directed reaching: primates naturally direct their gaze towards the target of an upcoming reach, and when they do so the reach is more efficient than if the gaze is directed elsewhere (Arora et al., 2019; Hawkins et al., 2013; Neggers & Bekkering, 2001, 2000).

Therefore, we propose that the observed gaze behavior in our task is the result of two distinct processes that take hold at different times in the trial. First, gaze is used early in the trial to explore and evaluate the two targets, because direct fixation is the only way to learn the offer values (see *Gaze Contingent Initial Mask and Visual Crowders*, above). In sequential sampling terms, these early exploratory fixations are the manner in which evidence is made available to the accumulator. Once a decision threshold is reached, we propose that the initial exploratory process yields to a second process, which is the tendency to look at the to-be-chosen target in order to maximize the efficiency of the reach. If the monkey is already looking at the target for which evidence has exceeded the requisite threshold, he will maintain his gaze on that target. If instead he is looking at the other target at the time the threshold is reached, he will shift his gaze onto the to-be-chosen target. Once the gaze is fixed upon chosen target, the eyes remain locked

there until the reach is ultimately completed—a phenomenon known as gaze anchoring (Neggers & Bekkering, 2001, 2000). The net effect of this end-of-trial gaze anchoring is that the to-be-chosen target accrues a nontrivial amount of additional fixation time at the end of the trial, resulting in longer final fixations.

Gaze anchoring, by definition, is a post-decision process—in sequential sampling terms, a form of terminal non-decision time (NDT). As such, it was prudent that we account for this interval prior to evaluating the effects of gaze on the decision process itself. In the following section, we describe a procedure for estimating the duration of the terminal NDT, and removing this interval from the data.

Estimation of Terminal NDT

For each monkey, we estimated the duration of the NDT by quantifying the extra fixation time devoted to chosen targets during final fixations (those terminated by choice) as compared to middle fixations (naturally-terminated by a saccade). Final fixations durations are defined as the time between fixation onset and the center lever lift, whereas middle fixation durations are defined simply as the time between fixation onset and the next saccade. Since fixation durations depend on relative target values (see *Results*), to estimate NDT we used only those trials with equal-valued targets in order to eliminate effects of relative value on fixation durations. Further, to ensure that fixations were matched for absolute order we used only the fixations that occurred second in each trial: because the vast majority of trials had either two or three total fixations and, second fixations offered the largest available group of fixations for which it was possible to compare middle to final durations.

First, we categorized fixations by absolute order, identified the second fixation that occurred in each trial, and then divided these second fixations into two pools. One pool consisted

of middle fixations terminated by the initiation of a saccade onto another target (Figure 2A,D (gray distribution); Monkey C: $n = 1545$, Monkey K: $n = 1281$). The other pool consisted of final fixations terminated by a choice initiated in favor of the currently fixated target (Figure 2B,D (black distribution); Monkey C: $n = 764$, Monkey K: $n = 792$). For each of the two pools, we generated 1000 bootstrapped samples, with each sample consisting of 1000 fixations taken with replacement from each pool, and then found the median of each bootstrapped sample. We then subtracted the vector of medians for the naturally terminated fixations from the vector of choice-terminated fixations, resulting in a vector of 1000 differences, which we took as the bootstrapped distribution for the median NDT.

A separate distribution of median NDTs was obtained for each monkey. From these distributions, we selected as a point estimate the 95th percentile for truncation of the trial data. For Monkey C, the estimate was 197ms, and for Monkey K, it was 200ms. We elected to truncate with the 95th percentile NDT estimate in order to remove as much NDT time as possible from each trial, and to therefore minimize influence of gaze anchoring upon the various measures of gaze bias in decision-making.

Truncation of Non-Decision Time

The monkey-specific NDT estimates were used to truncate the gaze data from the end of every trial, marked by the center lever lift. In other words, in every trial the gaze data from Monkey C were treated as if the trial had actually ended 197ms before the center lever lift; all calculations related to the gaze sequence and timing were calculated with respect this *virtual* trial stopping point. Likewise, all gaze data for Monkey K were calculated as if every trial had ended 200ms before the center lever lift. In this way, the effects of gaze anchoring during the NDT were minimized. The truncated data were used to compute all of the gaze biases (the initial

fixation bias, the final fixation bias, and the cumulative gaze-time bias) and were also used as input to the computational model.

Computational Model

To gain insight into potential mechanisms underlying choice and gaze behavior in this task, we employed the Gaze Weighted Linear Accumulator Model (GLAM, Thomas et al., 2019). The GLAM is an aSSM characterized by a series of parallel accumulators racing to a single bound. These accumulators are co-dependent, such that their combined probability of reaching the bound was equal to 1. For two choice options, this formulation is mathematically equivalent to a single accumulator drifting between two bounds (Thomas et al., 2019, Krajbich et al., 2011). The effect of gaze on choice behaviors is implemented as follows. The GLAM posits that over the course of a trial the absolute evidence signal for each item, i , can have two states: a biased state during periods when an item is viewed, and an unbiased state when it is not. The average absolute evidence for each item, A_i , is a linear combination of these two states throughout the trial and is defined according to Equation 1.

$$A_i = g_i r_i + (1 - g_i) \gamma r_i \quad (1)$$

Here, r_i is the veridical value of target i (number of drops of juice), g_i is the fraction of gaze time that the monkey devoted to that target during a given trial, and γ is the crucial gaze bias parameter, which determines the extent to which the non-fixated item is discounted. This parameter has an upper bound of 1, which corresponds to no gaze bias. A value of γ less than 1 indicates biased accumulation in favor of the fixated item.

The gaze-adjusted absolute evidence signals for each item are then used as inputs to an evidence accumulation process, as follows: First, each absolute value signal (A_i) is converted to

a relative evidence signal, R_i , through a logistic transform (see Molter et al., 2019; Thomas et al., 2019 for details and rationale). Then, evidence for each item (E_i) at time t is accumulated according to Equation 2, where v and σ indicate the speed of accumulation and the standard deviation of the noise, respectively, and R_i is the relative evidence signal. N indicates a univariate normal distribution. A choice is made when the evidence for one of the items reaches the threshold.

$$E_i(t) = E_i(t - 1) + vR_i + N(0, \sigma^2); E_i(0) = 0 \quad (2)$$

We fit two variants of the GLAM to each monkey's data. In the first model, the gaze bias parameter γ was left as a free parameter allowing gaze to modulate accumulation. The second model assumed no gaze bias, which was accomplished by setting $\gamma = 1$. Model fit was assessed using the Widely Applicable Information Criterion (WAIC, Vehtari et al., 2017) which accounts for the differences in the complexity of the two model variants. Specifically, we used both the signed difference of WAIC values and the WAIC weights for each model, which can be interpreted as the probability that each model is true, given the data (Vehtari et al., 2017). For the signed difference in WAIC scores, a negative value indicates a better fit for the gaze bias model (for more details see Thomas et al., 2019). In addition to the relative model fit, we assessed the model's predictive accuracy by estimating the parameters for the better of the two model variants using only even numbered trials, and then using these parameter estimates to simulate the results of each of the held-out odd-numbered trials 10 times.

We quantified the goodness-of-fit of the simulations by comparing three outcome metrics: choice, RT, and corrected choice probability. Note that since value is removed from the corrected choice probability metric, the only input to the model that could modulate it was the relative amount of gaze time devoted to each target. Thus, we can use the corrected choice

probability to assess the model’s ability to emulate the net effect of gaze alone on choice. For each metric we computed a mixed effects regression (logistic for choice, linear for RT and the corrected choice probability), in which a vector including both the empirically observed data and the simulations was explained as a function of a binary variable indicating whether a data point came from the empirically observed data, or from the model simulations (i.e., $Metric \sim 1 + simulated_or_observed$). If the fixed effect of this variable was not significantly different from 0, then the simulated data did not deviate from the empirically observed data. The mean deviations for each metric are given by the beta-weights and are reported in Table S2 along with the confidence intervals. As with the other regression models used in this paper, “Monkey” was included as a random effect.

All GLAM parameter estimation procedures were identical to those set forth in Thomas et al., 2019 with one exception: Thomas and colleagues assumed a fixed rate of random choice of 5%. This was changed to 1% to reflect the extremely low lapse rate for the easiest trials shown by both monkeys. The lapse rate for each monkey was estimated by finding the percentage of trials in which the monkey chose a 1-drop target when a 5-drop target (the highest possible value) was also available (mean lapse rate: Monkey C = 0.64%, Monkey K = 0.33%).

Code Accessibility

The custom MATLAB scripts used for behavior and eye tracking analyses are available upon request. The *glambox* toolbox that was used for the computational modeling can be accessed at <https://github.com/glamlab/glambox>.

References

- Arora, H. K., Bharmauria, V., Yan, X., Sun, S., Wang, H., & Crawford, J. D. (2019). Eye-head-hand coordination during visually guided reaches in head-unrestrained macaques. *Journal of Neurophysiology*, 122(5), 1946–1961. <https://doi.org/10.1152/jn.00072.2019>
- Battaglia-Mayer, A., Ferraina, S., Genovesio, A., Marconi, B., Squatrito, S., Molinari, M., Lacquaniti, F., & Caminiti, R. (2001). Eye–Hand Coordination during Reaching. II. An Analysis of the Relationships between Visuomanual Signals in Parietal Cortex and Parieto-frontal Association Projections. *Cerebral Cortex*, 11(6), 528–544. <https://doi.org/10.1093/cercor/11.6.528>
- Cavanagh, S. E., Malalasekera, W. M. N., Miranda, B., Hunt, L. T., & Kennerley, S. W. (2019). Visual fixation patterns during economic choice reflect covert valuation processes that emerge with learning. *Proceedings of the National Academy of Sciences*, 116(45), 22795–22801. <https://doi.org/10.1073/pnas.1906662116>
- Chandrasekaran, C., Peixoto, D., Newsome, W. T., & Shenoy, K. V. (2017). Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nature Communications*, 8(1), 1–16. <https://doi.org/10.1038/s41467-017-00715-0>
- Churchland, M. M. (2006). Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. *Journal of Neuroscience*, 26(14), 3697–3712. <https://doi.org/10.1523/JNEUROSCI.3762-05.2006>
- Cisek, P., & Kalaska, J. F. (2005). Neural Correlates of Reaching Decisions in Dorsal Premotor Cortex: Specification of Multiple Direction Choices and Final Selection of Action. *Neuron*, 45(5), 801–814. <https://doi.org/10.1016/j.neuron.2005.01.027>
- Crowder, E. A., & Olson, C. R. (2015). Macaque monkeys experience visual crowding. *Journal of Vision*, 15(5), 14–14. <https://doi.org/10.1167/15.5.14>
- Dean, H. L., Martí, D., Tsui, E., Rinzel, J., & Pesaran, B. (2011). Reaction Time Correlations during Eye–Hand Coordination: Behavior and Modeling. *Journal of Neuroscience*, 31(7), 2399–2412. <https://doi.org/10.1523/JNEUROSCI.4591-10.2011>
- Fetsch, C. R. (2016). The importance of task design and behavioral control for understanding the neural basis of cognitive functions. *Current Opinion in Neurobiology*, 37, 16–22. <https://doi.org/10.1016/j.conb.2015.12.002>
- Folke, T., Jacobsen, C., Fleming, S. M., & De Martino, B. (2016). Explicit representation of confidence informs future value-based decisions. *Nature Human Behaviour*, 1(1), 1–8. <https://doi.org/10.1038/s41562-016-0002>
- Gold, J. I., & Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience*, 30(1), 535–574. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>
- Gwinn, R., Leber, A. B., & Krajbich, I. (2019). The spillover effects of attentional learning on value-based choice. *Cognition*, 182, 294–306. <https://doi.org/10.1016/j.cognition.2018.10.012>

- 806 Hawkins, K. M., Sayegh, P., Yan, X., Crawford, J. D., & Sergio, L. E. (2013). Neural Activity in
807 Superior Parietal Cortex during Rule-based Visual-motor Transformations. *Journal of*
808 *Cognitive Neuroscience*, 25(3), 436–454. https://doi.org/10.1162/jocn_a_00318
- 809 Hunt, L. T., Malalasekera, W. M. N., de Berker, A. O., Miranda, B., Farmer, S. F., Behrens, T.
810 E. J., & Kennerley, S. W. (2018). Triple dissociation of attention and decision
811 computations across prefrontal cortex. *Nature Neuroscience*, 21(10), 1471–1481.
812 <https://doi.org/10.1038/s41593-018-0239-5>
- 813 Krajbich, I. (2019). Accounting for attention in sequential sampling models of decision making.
814 *Current Opinion in Psychology*, 29, 6–11. <https://doi.org/10.1016/j.copsyc.2018.10.008>
- 815 Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and
816 comparison of value in simple choice. *Nature Neuroscience*, 13(10), 1292–1298.
817 <https://doi.org/10.1038/nn.2635>
- 818 Krajbich, I., Mitsumasu, A., Polania, R., Ruff, C. C., & Fehr, E. (2021). A causal role for the
819 right frontal eye fields in value comparison. *ELife*, 10, e67477.
820 <https://doi.org/10.7554/eLife.67477>
- 821 Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship
822 between visual fixations and choice in value-based decisions. *Proceedings of the*
823 *National Academy of Sciences*, 108(33), 13852–13857.
824 <https://doi.org/10.1073/pnas.1101328108>
- 825 Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017).
826 Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*, 93(3), 480–490.
827 <https://doi.org/10.1016/j.neuron.2016.12.041>
- 828 Kravitz, D. J., Saleem, K. S., Baker, C. I., Ungerleider, L. G., & Mishkin, M. (2013). The ventral
829 visual pathway: An expanded neural framework for the processing of object quality.
830 *Trends in Cognitive Sciences*, 17(1), 26–49. <https://doi.org/10.1016/j.tics.2012.10.011>
- 831 Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic Interaction
832 between Reinforcement Learning and Attention in Multidimensional Environments.
833 *Neuron*, 93(2), 451–463. <https://doi.org/10.1016/j.neuron.2016.12.040>
- 834 Lim, S.-L., O'Doherty, J. P., & Rangel, A. (2011). The Decision Value Computations in the
835 vmPFC and Striatum Use a Relative Value Code That is Guided by Visual Attention.
836 *Journal of Neuroscience*, 31(37), 13214–13223.
837 <https://doi.org/10.1523/JNEUROSCI.1246-11.2011>
- 838 Liu, H.-Z., Lyu, X.-K., Wei, Z.-H., Mo, W.-L., Luo, J.-R., & Su, X.-Y. (2021). Exploiting the
839 dynamics of eye gaze to bias intertemporal choice. *Journal of Behavioral Decision*
840 *Making*, 34(3), 419–431. <https://doi.org/10.1002/bdm.2219>
- 841 Liu, H.-Z., Zhou, Y.-B., Wei, Z.-H., & Jiang, C.-M. (2020). The power of last fixation: Biasing
842 simple choices by gaze-contingent manipulation. *Acta Psychologica*, 208, 103106.
843 <https://doi.org/10.1016/j.actpsy.2020.103106>

- 844 McGinty, V. B., & Lupkin, S. M. (2021). Value signals in orbitofrontal cortex predict economic
845 decisions on a trial-to-trial basis. *BioRxiv*, 2021.03.11.434452.
846 <https://doi.org/10.1101/2021.03.11.434452>
- 847 McGinty, V. B., Rangel, A., & William T. Newsome. (2016). Orbitofrontal Cortex Value Signals
848 Depend on Fixation Location during Free Viewing. *Neuron*, 90(6), 1299–1311.
849 <https://doi.org/10.1016/j.neuron.2016.04.045>
- 850 Molter, F., Thomas, A. W., Heekeren, H. R., & Mohr, P. N. C. (2019). GLAMbox: A Python
851 toolbox for investigating the association between gaze allocation and decision behaviour.
852 *BioRxiv*, 741678. <https://doi.org/10.1101/741678>
- 853 National Research Council (US) Committee for the Update of the Guide for the Care and Use of
854 Laboratory Animals. (2011). *Guide for the Care and Use of Laboratory Animals* (8th
855 ed.). National Academies Press (US). <http://www.ncbi.nlm.nih.gov/books/NBK54050/>
- 856 Neggers, & Bekkering. (2001). Gaze Anchoring to a Pointing Target Is Present During the Entire
857 Pointing Movement and Is Driven by a Non-Visual Signal. *Journal of Neurophysiology*,
858 86(2), 961–970. <https://doi.org/10.1152/jn.2001.86.2.961>
- 859 Neggers, S. F., & Bekkering, H. (2000). Ocular gaze is anchored to the target of an ongoing
860 pointing movement. *Journal of Neurophysiology*, 83(2), 639–651.
861 <https://doi.org/10.1152/jn.2000.83.2.639>
- 862 Öngür, D., & Price, J. L. (2000). The Organization of Networks within the Orbital and Medial
863 Prefrontal Cortex of Rats, Monkeys and Humans. *Cerebral Cortex*, 10(3), 206–219.
864 <https://doi.org/10.1093/cercor/10.3.206>
- 865 Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic
866 value. *Nature*, 441(7090), 223–226. <https://doi.org/10.1038/nature04676>
- 867 Pärnamets, P., Johansson, P., Hall, L., Balkenius, C., Spivey, M. J., & Richardson, D. C. (2015).
868 Biasing moral decisions by exploiting the dynamics of eye gaze. *Proceedings of the*
869 *National Academy of Sciences*, 112(13), 4170–4175.
870 <https://doi.org/10.1073/pnas.1415250112>
- 871 Paton, J. J., Belova, M. A., Morrison, S. E., & Salzman, C. D. (2006). The primate amygdala
872 represents the positive and negative value of visual stimuli during learning. *Nature*,
873 439(7078), 865–870. <https://doi.org/10.1038/nature04490>
- 874 Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-
875 choice decision tasks. *Neural Computation*, 20(4), 873–922.
- 876 Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-
877 making. *Nature*, 461(7261), 263–266. <https://doi.org/10.1038/nature08275>
- 878 Rich, E. L., & Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex.
879 *Nature Neuroscience*, 19(7), 973–980. <https://doi.org/10.1038/nn.4320>
- 880 Shenoy, K. V., Kaufman, M. T., Sahani, M., & Churchland, M. M. (2011). A dynamical systems
881 view of motor preparation: Implications for neural prosthetic system design. *Progress in*
882 *Brain Research*, 192, 33–58. <https://doi.org/10.1016/B978-0-444-53355-5.00003-8>

- Smith, S. M., & Krajbich, I. (2018). Attention and choice across domains. *Journal of Experimental Psychology: General*, 147(12), 1810–1826.
<https://doi.org/10.1037/xge0000482>
- Sui, X.-Y., Liu, H.-Z., & Rao, L.-L. (2020). The timing of gaze-contingent decision prompts influences risky choice. *Cognition*, 195, 104077.
<https://doi.org/10.1016/j.cognition.2019.104077>
- Tavares, G., Perona, P., & Rangel, A. (2017). The Attentional Drift Diffusion Model of Simple Perceptual Decision-Making. *Frontiers in Neuroscience*, 11.
<https://doi.org/10.3389/fnins.2017.00468>
- Thomas, A. W., Molter, F., Krajbich, I., Heekeren, H. R., & Mohr, P. N. C. (2019). Gaze bias differences capture individual choice behaviour. *Nature Human Behaviour*.
<https://doi.org/10.1038/s41562-019-0584-8>
- Vaidya, A. R., & Fellows, L. K. (2015). Ventromedial Frontal Cortex Is Critical for Guiding Attention to Reward-Predictive Visual Features in Humans. *Journal of Neuroscience*, 35(37), 12813–12823. <https://doi.org/10.1523/JNEUROSCI.1607-15.2015>
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1413–1432.
<https://doi.org/10.1007/s11222-016-9696-4>
- Whitney, D., & Levi, D. M. (2011). Visual crowding: A fundamental limit on conscious perception and object recognition. *Trends in Cognitive Sciences*, 15(4), 160–168.
<https://doi.org/10.1016/j.tics.2011.02.005>
- Wise, S. P. (2008). Forward frontal fields: Phylogeny and fundamental function. *Trends in Neurosciences*, 31(12), 599–608. <https://doi.org/10.1016/j.tins.2008.08.008>

911

Supplemental Material

Figure S1. Full task sequence **A.** Illustration of the complete task sequence. Note that for illustration purposes, in this panel the background is white, and the surrounding crowders are shown at reduced saturation relative to the centrally located targets. However, on the actual task display, the background was dark, and target luminance was reduced relative to the crowders (see panel B). **B.** Close-up view of a single target array, consisting of a central yellow choice target surrounded by six non-task-relevant crowders. Two such arrays appear on the display (panel A), each centered 7.5° from the display center. **C.** To generate stimulus sets, four shapes are selected without replacement from 7 possible; then each is rotated randomly by 0, 90, 180, 270 degrees. The four rotated shapes are then combined with four colors selected randomly from equally-spaced points on a color wheel defined within the RGB gamut in the CIELUV color space. The grid on the left shows the color and shape primitives as rows and columns, respectively. Each primitive was assigned a value from 0 to 4, corresponding to the number of drops of juice each would signal on its own. The 12 targets that made up a given stimulus set were taken from the off-diagonal of this grid; their value determined according to their shape and color. The result is a set of targets ranging in value from 1 to 5 drops of juice.

912

913

914

	Monkey C		Monkey K	
	Bias Model	No Bias Model	Bias Model	No Bias Model
γ-MAP	0.28	1	0.17	1
HPD 2.5	0.26	1	0.15	1
HPD 97.5	0.3	1	0.19	1
σ-MAP	0.24	0.26	0.22	0.25
HPD 2.5	0.24	0.26	0.22	0.25
HPD 97.5	0.24	0.26	0.23	0.25
τ-MAP	0.32	0.23	0.35	0.24
HPD 2.5	0.31	0.22	0.34	0.24
HPD 97.5	0.33	0.23	0.36	0.25
v-MAP	2.6	2.62	2.52	2.55
HPD 2.5	2.59	2.61	2.51	2.54
HPD 97.5	2.61	2.63	2.53	2.56

Table S1. Parameter estimates from both GLAM variants. *Maximum a posteriori (MAP) and 97.5 highest posterior density (HPD) interval for each parameter obtained from fitting each variant to all data available for each monkey. Note that for “Bias Model”, γ was a free parameter in the model, while in the “No Bias Model”, γ was fixed at 1. γ : gaze bias parameter; σ : noise parameter; τ : logistic scaling parameter; v : velocity parameter. See Methods for further details on parameter definitions.*

Metric	Estimate	Confidence Interval	Figure
Choice (0 = left, 1 = right)	0.017	[-0.063, 0.097]	6 A/D
Reaction Time	0.0023	[-0.0046, 0.0092]	6 B/E
Value-corrected choice probability	3.40E-17	[-0.0060, 0.0060]	6 C/F

Table S2. Comparison of GLAM simulations to held out trials. *Estimates and confidence intervals are the result of mixed effects models where each metric was regressed against a binary variable indicating whether the data were from the predicted trials or the observed trials. Confidence intervals including zero indicate no significant difference between predicted vs. observed, i.e. that the predicted data are a good fit to out of sample observations. The choice metric was assessed using a logistic regression. The others were linear. All models included “Monkey” as a random effect. All p-values > 0.5.*