

Robust sampling of decision information during perceptual choice

Hildward Vandormael^{a,1}, Santiago Herce Castañón^a, Jan Balaguer^a, Vickie Li^a, and Christopher Summerfield^a

^aDepartment of Experimental Psychology, University of Oxford, Oxford OX1 3UD, United Kingdom

Edited by Jacqueline Gottlieb, Columbia University, New York, New York, and accepted by Editorial Board Member Thomas D. Albright January 18, 2017 (received for review September 12, 2016)

Humans move their eyes to gather information about the visual world. However, saccadic sampling has largely been explored in paradigms that involve searching for a lone target in a cluttered array or natural scene. Here, we investigated the policy that humans use to overtly sample information in a perceptual decision task that required information from across multiple spatial locations to be combined. Participants viewed a spatial array of numbers and judged whether the average was greater or smaller than a reference value. Participants preferentially sampled items that were less diagnostic of the correct answer (“inlying” elements; that is, elements closer to the reference value). This preference to sample inlying items was linked to decisions, enhancing the tendency to give more weight to inlying elements in the final choice (“robust averaging”). These findings contrast with a large body of evidence indicating that gaze is directed preferentially to deviant information during natural scene viewing and visual search, and suggest that humans may sample information “robustly” with their eyes during perceptual decision-making.

decision-making | categorization | eye movements | numerical averaging | information sampling

A well-validated computational framework argues that decisions are made by integrating information toward a threshold, at which point a response is made (1, 2). Within this integration-to-bound framework, decisions are optimized by sequential sampling, at the cost of delaying responding (3). Studies using this framework in conjunction with psychophysical tasks have suggested that humans and monkeys make approximately optimal perceptual category judgments, with choice accuracy limited mainly by sensory noise or uncertainty in the external world (2, 4, 5). These psychophysical studies have been argued to offer a window on the general mechanisms by which humans make decisions (6).

However, paradigms in which stimulation is continuously available and controlled exclusively by the experimenter differ from everyday choices in real-world environments. For example, in the random dot motion (RDM) paradigm, subjects passively sample from a single, continuous, centrally presented stream of information and decide whether average dot motion tends in one direction or another (7). By contrast, decisions about the average feature in a natural scene are made by actively saccading (or covertly attending) to successive image locations and sampling the subset of information that is available there. For example, a foraging monkey might move its eyes freely to several locations in the forest canopy and average the available fruit yield to decide where to move. Real-world category decisions, thus, are affected not only by the quality of the sensory information and the dynamics of integration, but also by the policy that dictates how information is sampled and weighted when forming a decision. However, the nature of this policy remains largely unknown.

The determinants of saccadic eye movements have been extensively investigated during visual search (8) and natural scene perception (9). In visual search tasks, eye movements are attracted to image locations that contain perceptually salient or surprising items (10, 11) or to those points at which task-relevant information is maximal (12–15). In value-based choice tasks, in which for example participants judge which of two or more food items is preferred, choices and eye movements are interlinked, in

that items that are fixated more frequently are chosen more often (16, 17). Other models have suggested that more uncertain sources are sampled more thoroughly with the eyes (18) or that decisions and information-gathering mutually drive one another in a reciprocal cascade (19). However, little is known about the policy by which participants sample information during choice tasks that involve integrating decision information from multiple locations across the visual field. Various studies have investigated the sampling policy in differing domains, such as category learning (20) and classification tasks (21). Recently, researchers have highlighted the importance of investigating evidence accumulation as an active sampling process (22). We investigated this sampling process by measuring eye movements while human participants performed a categorization task that involved averaging numbers in a spatial array. We used fully visible symbolic numbers as stimuli, allowing us to assess how gaze placement was determined by the statistics of decision information, rather than the perceptual arrangement of the stimulus array.

When averaging perceptual features across multiple locations under central fixation, participants tend to downweight or ignore extreme values, such as a deviant expression in a crowd of faces (23) or outlying color values in an array of shapes (24). This has led to the proposal that “inlying” items (those that are closer to the category boundary) are processed with higher gain than outliers (24, 25). However, a related possibility is that observers tend to sample by preference those items that are closer to the category boundary. Testing this theory, we found humans prefer to fixate the inlying items in a stimulus array during a perceptual decision-making task—a policy that contradicts the well-established preference for fixating deviant items in visual search tasks.

Significance

When making decisions, humans and other primates move their eyes freely to gather information about their environment. A large literature has explored the factors that determine where the eyes fall during natural scene perception and visual search, concluding that deviant or surprising perceptual information attracts attention and gaze. Here, we describe a sampling policy that paradoxically shows that in a decision-making task, the eyes are attracted to expected rather than unexpected information: When classifying the average of an array of numbers, human observers looked by preference at the numbers that were closest to the mean. This policy drove a behavioral tendency to discount the influence of outliers when making choices, leading to “robust” choices about the stimulus array.

Author contributions: H.V. and C.S. designed research; H.V. performed research; H.V., S.H.C., J.B., and V.L. contributed new reagents/analytic tools; H.V., V.L., and C.S. analyzed data; and H.V. and C.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. J.G. is a Guest Editor invited by the Editorial Board.

¹To whom correspondence should be addressed. Email: hildward.vandormael@sjc.ox.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1613950114/-DCSupplemental.

This “robust” sampling policy drove decisions to ignore or discount outlying information during perceptual decisions.

Results

Participants ($n = 54$) viewed eight two-digit numbers distributed quasirandomly around the screen (Fig. 1A). Numbers were drawn from a distribution with mean μ and SD σ , and participants were asked to compare their average to a reference value R that varied across blocks ($\mu = R \pm 2$ or $R \pm 4$ in low and high μ trials, respectively, and $\sigma = 6$ or 12 in low and high σ trials, respectively). Half of the participants indicated whether the average was higher (or not) than the reference; the others indicated whether the average was lower (or not). Numbers were interspersed among task-irrelevant items (letters), which introduced a gentle crowding effect and encouraged shifts of gaze during decision formation. By tracking gaze during task performance, we were able to ask how the information sampled with the eyes related to decisions, with a focus on whether participants sampled by preference those numbers that were inlying (close to reference) or outlying (far from reference).

We report data collapsed over two experiments that differed only in terms of the response deadline (3,000 ms and 5,000 ms in experiment 1 and experiment 2, respectively). Overall mean accuracy was 70.5% ($n = 54$). Performance depended on both the mean and variance of the distribution from which the numbers were drawn (Fig. 1B for accuracy and Table S1 for response time). Accuracy increased on high μ trials compared with low μ trials [$F_{(1, 53)} = 76, P < 0.001$] and decreased on high σ (more variable evidence) trials compared with low σ trials [$F_{(1, 53)} = 47, P < 0.001$], as previously reported (24).

Gaze Trajectories. We developed an approach to analyzing eye-tracking data (“landscape analysis”) that measured the decision information sampled by the eyes as a continuous, time-varying

quantity (the “evidence trajectory,” or ET) across each trial (see Fig. S1 and SI Methods for a detailed description of the analysis method). We first converted the stimulus array on each trial to a smooth topography (or landscape) of decision values, with peaks or troughs centered on the spatial position of each element and a height and sign that was determined by the difference between each element V_k and the reference R (positive peaks favoring the decision “higher than reference,” negative troughs favoring “lower”; Fig. 1C). The momentary evidence at each point in time was defined as the landscape value under the gaze location, and the ET constituted the time-series of values obtained as the eyes moved across the stimulus array on each trial (Fig. 1D and SI Methods). This is equivalent to conceiving of the gaze as a “soft” (or smooth) information-harvesting spotlight moving across the screen, allowing evidence to be accumulated up to a decision threshold. The width of the spotlight was estimated independently of further analyses under the assumption that participants’ gaze maximized the acquisition of total sensory evidence. We fit this value by searching for values that maximized “quality trajectory” (QT)—that is, the extent to which items were viewed irrespective of their task-relevant values (SI Methods). We also validated all analyses using a “hard” spotlight—that is, an aperture of fixed width with no spatial smoothing, obtaining qualitatively identical results (Fig. S2).

Decisions Depend on Gaze Location. Using this approach, we first compared cumulative ET on trials where participants responded “higher” and those where they responded lower, correcting for differences in available information on “high” and “low” trials (SI Methods). As expected, choices depended strongly on the decision information sampled, with average cumulative ET building up steadily until the response: participants responded lower when the sampled evidence was lower than the reference and higher when the sampled evidence was higher than the reference (Fig. 1D, Bottom). In other words, consistent with previous findings, choices were strongly affected by where participants looked (16).

Gaze Is Directed Preferentially to Inliers Over Outliers. Next, we used a similar landscape approach to determine the extent to which participants’ gaze favored items that were inlying or outlying (numerically closer or further from the reference). We generated a new set of landscapes in which the topography on each trial was determined by the rank of the absolute value of each element (number) within the stimulus array, so that more inlying elements had higher values and more outlying elements had lower values (Fig. 2A; similar results were obtained when this analysis was carried out using the absolute values instead of the ranks). The integral of the eye movement trajectory under these landscapes, which we call the “prototypicality trajectory” (PT), thus encoded participants’ time-varying preference to fixate those elements that were numerically similar (inliers; denoted by positive values) or dissimilar (outliers; negative values) to the reference. We found that inlying items were sampled preferentially, from 400 to 800 ms poststimulus presentation, indicated by a significantly positive PT value (Fig. 2B, black line and shading). This result survived correction for multiple comparisons across time points using a well-validated cluster-based approach (26).

Next, we evaluated this preference for inliers using a different approach. On each fixation, we defined the item being viewed as that which was closest to the current gaze location, allowing us to calculate the probability of gaze moving from an item of one rank (i.e., ordinal position in the array) to another. To avoid any biases that may arise from nonindependence of nearby items, we constructed a null distribution specific to our experiment by randomly permuting item-location assignments on each trial and repeating the analysis 1,000 times. Next, we computed the difference between gaze-transition probabilities in our dataset and under the null distribution. Finally, we calculated the probability

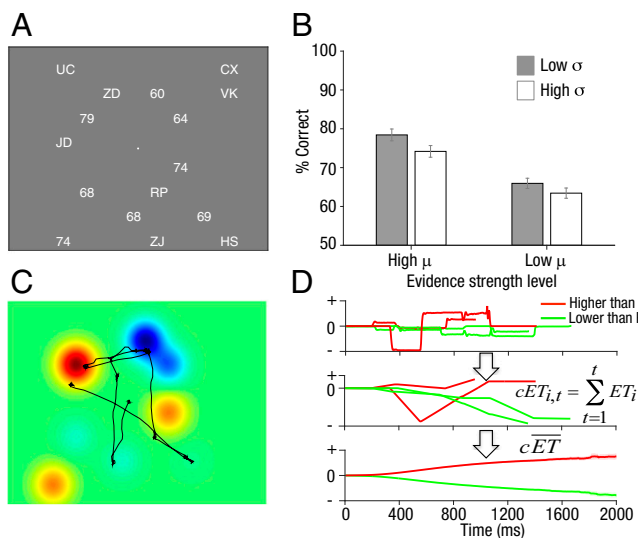


Fig. 1. Average accuracy and demonstration of the landscape approach. (A) Example trial display. Participants could sample the numbers freely by moving their eyes and had to indicate whether the average of all of the numbers on the screen was higher or lower than the reference value. (B) Performance increased for higher μ and lower σ . Bars are shown with 95% confidence intervals. (C) Landscape was created by convolving each stimulus location with a bivariate Gaussian basis function. Evidence landscape was created by setting the height of the Gaussian associated with each number location equal to the number value – reference value ($V_k - R$). The integral of gaze (black line) under these landscapes is used to calculate the ET. (D, Top) ET over time for four example trials (lines). (Middle) The cumulative ET over time (t) for these four example trials (i). (Bottom) The cumulative ET time series averaged across trials. Cumulative ET is on average positive when participants responded higher than reference and negative when participants responded lower than reference (green).

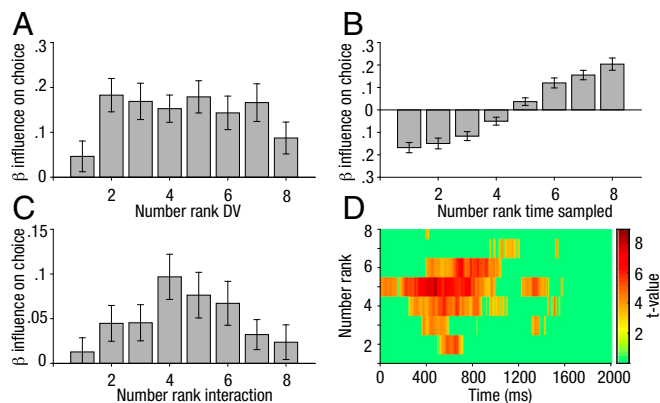


Fig. 4. Robust averaging of decision information. (A) Weights associated with each item rank, calculated by sorting the numbers associated with each item, and using probit regression to predict participants' choices. The inverted-u shape over ranks shows evidence for robust averaging (larger weights for inliers). (B) Weights associated with the time spent sampling each number rank. More time spent sampling low-ranked led to a tendency to respond lower than higher and vice versa. (C) Weights associated with the interaction between number rank and the time spent sampling these numbers. The inverted-u shape over ranks shows evidence that robust averaging is mediated by gaze duration. (D) Robust averaging due to the sampling policy as a function of time. Ranks are now shown on the y axis and t values by the green-red color scale. Robust averaging due to preferential sampling of inliers was maximal during the 400–800 ms time frame (larger weights for inliers). Bars on A–C are shown with 95% confidence intervals.

Given that the prototypicality analysis (above) showed that inlying items were preferentially viewed 400–800 ms poststimulus, we examined whether this downweighting, due to preferential sampling, is also most apparent in this time period. To do this, we conducted the previous regression for each sampling time point individually and tested which of these resulting coefficients (for the Gaze \times Stimulus interaction) were significantly different from zero. In Fig. 4D, these are plotted (t values > 3) on a heat map, showing that this downweighting of outliers effect is also maximally dependent on gaze during this 400–800 ms timeframe. These results suggest that robust averaging is due at least in part to the sampling policy with which participants actively sample inlying rather than “outlying” information.

Together, Fig. 4A and C show that (i) inlying stimuli (i.e., those close to the reference) have a stronger influence on choice than outlying stimuli (i.e., those that are much larger or much smaller than the reference) and (ii) that this effect is multiplicatively enhanced by gaze. Thus, any tendency that inlying stimuli might have had to drive choices more strongly is exaggerated when they are fixated, but this is not true for outlying stimuli. In other words, human observers are not only “robust averagers” but also “robust samplers”—selectively viewing inlying category information and thereby ensuring that this information is a more potent driver of choice. In further control analyses, we ruled out the possibility that this effect was due to a “counting” strategy, in which all items above or below the reference contributed equally to decisions (Fig. S3).

Computational Simulations. Next, to validate our approach, we used computational simulations to estimate how the ET and PT should vary under differing sampling policies. Using distributions of decision information (i.e., numbers) identical to those viewed by human participants (i.e., same μ and σ), we constructed simulated ET and PT vectors consisting of alternating periods of fixation and eye movements, prefixed by a delay period. The duration of all of these periods (initial delay, fixation, eye movement) was sampled from the distributions exhibited by human participants. During simulated fixations, simulated ET was set to the decision values $V_k - R$ (where

V_k is the numerical value of item k), and during delay and eye movements, ET values were set to zero. Simulated PT was constructed in the same manner by setting PT during a fixation of stimulus k to its value ranked within the array, as for the human analyses. Model choices occurred when the cumulative simulated ET reached a fixed threshold, as in standard integration-to-bound models (2). The sign of the cumulative ET at this fixed threshold determined model choice and thus accuracy. To ensure that human PT and model PT values were on the same scale, we multiplied the model PT with the average QT for the corresponding participant.

This framework allowed us to manipulate the model sampling policy—that is, to simulate the order in which the array items were sampled. We implemented this via a single parameter π that dictated whether the observer preferred to view inliers before outliers ($\pi > 0$), outliers before inliers ($\pi < 0$), or sampled the elements in random order ($\pi \sim 0$). The parameter π encodes the (positive or negative) correlation between the observed sampling order and the most “prototypical” ordering—that is, from most inlying to most outlying. In Fig. 2C, we plot simulated PT for different values of π . The best fitting parameterization was determined per subject by calculating the mean squared error (MSE) between human and simulated PT for each value of π . Fig. 2B shows the simulated values of PT for the best fitting average π value over subjects (magenta line) plotted against human data. Human PT curves were best accounted for by positive π values, consistent with a policy that involved sampling inliers before outliers and thus with the analyses described above. Fig. S4 shows the distribution of best fitting values for π over subjects. Next, using the best fitting π values for each subject, we returned to the analysis of ET and compared the human and simulated ET values for the different conditions—that is, for different means μ and SDs σ of the decision values with respect to the reference. We observed a striking convergence between the human and model ET values (Fig. S5), providing an independent validation of our best fitting model.

Determinants of Human and Model Performance. Finally, we tested how model performance (accuracy) varied with π . Interestingly, performance for the model linearly decreased as the sampling policy π changed from a policy to sample outlying information first to a policy to sample inlying information (Fig. 2D). Thus, for an “ideal” model—in which decisions are only limited by the quality of information that is sampled from the screen—a policy that preferentially samples inlying information incurs a model performance cost relative to random sampling. To test whether this was the case for human observers, we conducted an additional analysis that evaluated how estimated human values of π predicted accuracy at the single-trial level. We computed single-trial estimates of π by correlating the order in which items were actually sampled with their ranked equivalent and then used logistic regression to estimate how these values predicted choice accuracy (Fig. S6). Both the sampling policy π [$t_{(53)} = 2.37$, $P < 0.05$] and the average prototypicality [$t_{(53)} = 11.3$, $P < 0.001$] were positive predictors of accuracy, indicating that unlike the model, participants performed better as they exhibited more robust sampling.

Discussion

In real-world settings, information arises in crowded natural scenes, replete with multiple features and objects that compete to drive decisions. Studies of saccadic control have revealed that attention is automatically captured by perceptually salient information, such as a target defined by a unique color in a visual search experiment (27), and may also be captured by semantically deviant information, such as an object in an unusual context (e.g., a sofa on the beach), although the latter claim remains controversial (11). However, the task at hand is a key determinant of saccadic control, and the eyes are drawn to those scene elements that are task-relevant, or valuable, or to those locations where the most information can be gleaned (10, 12–15,

28). Nevertheless, most decision-making studies either use a single, centrally presented source of information (such as a dot motion stimulus) precluding the detailed investigation of the policy by which participants sample decision information or assume that gaze allocation is a stochastic process that drives a preference for fixated alternatives, rather than vice versa (16, 17). Recently, researchers have highlighted the importance of investigating evidence accumulation as an active sampling process (22).

Here, we reveal an aspect of the sampling policy that determines where gaze is allocated during decision-making: faced with multiple spatially distinct sources of decision information, observers tend to ignore those sources that carry the most extreme information and focus instead on the inlying decision information—that is, that falling closest to the boundary that segregates one choice from another (or close to the mode of the overall information presented on one trial; our approach was not designed to distinguish these possibilities). This is surprising given that participants exhibit a well-characterized policy to saccade toward perceptually deviant information during natural scene viewing, or visual search (10, 11). The focus on inlying elements influenced human choices, with gaze driving a multiplicative effect to weight inlying elements more heavily, which in turn drove “robust averaging” in perceptual choices made during the experiment.

Robust averaging of perceptual information has been described before, for faces (23) and for colors and shapes (24, 25). Here, we add to these findings, revealing that humans engage in robust averaging of symbolic number. Importantly, we assumed that participants used an approximate number sense, meaning that numbers are treated as noisy estimates of magnitude, which can be averaged like continuous sensory signals (29). This assumption follows from the short duration for which numerical information was available and the low average dwell times observed for each number, which would preclude overt arithmetic approaches. Previous studies have left unresolved the question of whether robust averaging occurs because inlying items are sampled (i.e., overtly attended) by preference or whether the influence of outliers is attenuated at a later decision stage. For example, one model proposes that observers simply compute a posterior likelihood ratio conditional on the past distribution of features and their associated responses. This is equivalent to using a nonlinear decoder that “squashes” outlying information, muting its influence in the average and prolonging response latencies when decision information is heterogeneous (24). However, here we show that the robust averaging is driven at least in part by the policy with which gaze is directed during information integration. In other words, humans are not just robust averagers but robust samplers of decision information.

Why do humans sample and average robustly? We describe a computational simulation of the sampling policy adopted during our spatial averaging task that was able to faithfully recreate the relationship between the evidence sampled and the choices made by humans. Exploring model performance under a range of policies (spanning a tendency to sample inliers or outliers by preference) revealed that for an ideal model limited only by the information available on the screen, sampling outliers rather than inliers by preference boosted performance. This is consistent with the observation that when assigning samples to one of two Gaussian-distributed classes, eccentric or “off-channel” features are most diagnostic of the correct response (30, 31). It is very different, however, from the optimal policy that emerges as participants gradually learn a category boundary through repeated sampling (32) (see Fig. S7 for discussion).

However, in the human data, a tendency to sample inliers was a positive predictor of accuracy at the single-trial level (Fig. S6). This implies that a robust sampling policy has evolved to counter additional sources of loss that arise during the decision process, such as “late” noise that corrupts the information of integration, perhaps due to capacity limits on attention and working memory processes. For example, robust policies may lead to more stable

estimates of decision information, just as nonparametric statistical inference is robust to the inclusion of aberrant data points. Other recent work has suggested that apparent “suboptimalities” in human judgment may in fact maximize performance when this late noise is taken into account (33, 34). The effect of the interplay between the decision environment (i.e., task characteristics) and cognitive constraints on the relative (sub)optimality of various sampling policies is an interesting avenue for future research.

In summary, the current results highlight that perceptual judgments are not only constrained by the quality of the sensory information and the dynamics of integration but also by the sampling policy, which dictates how the information is selected for entry into the decision process. However, other questions remain open. In particular, the information processing steps by which fixations are guided to subsequent inliers remain unclear in our work, although this presumably relies on parafoveal processing of nonfixated locations. One possibility is that covert attention is first shifted to candidate items and the eyes are then “pulled” to this location (35, 36). Nor does our study directly address how saccade placement varies according to the spatial arrangement of items on the screen. For example, participants may prefer to saccade toward the center of mass of clusters of objects (37). Thus, gaze can indeed be driven by the spatial arrangement of items on the screen, but gaze is additionally driven by the decision values of the available evidence, exhibited in this work via a preference to sample inlying samples.

Methods

Participants. Sixty healthy human participants with normal or corrected-to-normal vision were recruited in two experiments (experiment 1, $n = 30$; experiment 2, $n = 30$). Participants provided written consent beforehand, and the study was approved by the Oxford University Medical Sciences Division Ethics Committee (approval no. MSD-IDREC-C1-2009-1). There were only minor differences between the two experiments (highlighted below), and so we collapse across their results. Six participants whose performance failed to differ from chance were excluded from all further analyses, leaving $n = 54$.

Stimuli. On each trial, participants viewed an array composed of eight two-digit numbers. Stimuli were generated using PsychToolbox (psychtoolbox.org) for MATLAB (Mathworks) and presented on a 17-inch screen (resolution, 1024 × 768) viewed from a distance of 60 cm. Numbers (~1.15° horizontal and ~0.95° vertical visual angle) were presented at randomly selected locations within an invisible 8 × 8 grid. Eight additional randomly selected grid locations contained letter pairs, which were distracters. No letter or number pairs were presented within the central four grid locations. An example stimulus array can be seen in Fig. 1A.

At the start of each block of 96 trials, participants were informed of the reference value, against which they compared the average number for all trials across that block. Subsequently, each trial began with the onset of a central fixation cross (500 ms), followed by the stimulus array, which remained on the screen for either up to 3,000 ms (experiment 1) or up to 5,000 ms (experiment 2). During array presentation, participants freely moved their eyes across the screen and chose which was higher (lower): the reference number or the average of the number array. The decision framing (decide whether numbers higher vs. lower than reference) was counterbalanced across participants within each experiment. Participants could respond at any point during the array presentation by pressing one of two keys (“F” or “J”), and they received auditory feedback in the form of a high or low tone lasting 100 ms (800/400 Hz, correct/incorrect). In experiment 2 only, participants were additionally paid in proportion to their accuracy. In the high frame, participants received a bonus corresponding to their choice (reference vs. numbers, in pence) on 20 randomly selected trials. In the low frame, the reward of participants started at £20, from which their choice was deducted on 20 randomly selected trials. Participants earned £10 on average for their participation. Participants performed eight blocks of 96 trials (768 trials in total) in an experiment lasting ~50 min.

Design. The reference number R on each block was drawn randomly from a uniform distribution between 25 and 75. The sample numbers were drawn from one of eight Gaussian distributions whose mean μ differed from R by either ± 2 or ± 4 and whose SD σ was either 6 or 12. Sample values were resampled to ensure

that the sample mean and variance fell within the desired value for that trial with a tolerance of 0.7 and 0.1, respectively. Each of the eight trial types occurred equally often, and their presentation order within a block was random.

Eye-Tracking. All eye movement data were recorded monocularly using an SR Research EyeLink 1000 eye-tracking system at a sampling rate of 1,000 Hz. Participants placed their heads on a chin rest that was mounted on the table at a fixed distance of 68 cm from the screen. Before the start of each new block, a standard calibration/validation procedure was executed, in which participants tracked a dot that appeared at 13 evenly spaced locations across the screen. Data were converted to vectors of x, y position in the frame of reference of the screen exported to MATLAB using the EyeLink toolbox and analyzed using in-house scripts. Missing data (due to blinks or eye movements away from the screen) were ignored (i.e., set to NaN) in all analyses. Eye movement traces were cut off at reaction time for all analyses described above. All of our analyses assume that the eye-tracker faithfully estimates the pixel position of the gaze on the screen.

Analysis. We analyzed the eye-tracking data using both a combination of conventional methods (e.g., identifying the location and dwell time of individual fixations) and a technique that treated the stimulus array as a topography of decision information that was sampled continuously with the gaze (landscape analysis). We describe this method in detail in *SI Methods* (Fig. S1). The methods associated with other analyses are detailed in accompanying supplementary sections that are referenced in the main text (Figs. S3, S6, and S7).

Weighting of Decision Information on Choice. Probit regression was used to predict human choices (higher or lower than reference) as a function of three sets of predictors: (i) the decision values of each number ranked from lowest to highest, (ii) QT for each number QT_k (once again, with k ranked from lowest to highest), and (iii) the interaction between these two quantities. The QT for each number QT_k was computed as the sum of the integral under an element-specific QT landscape—that is, a landscape constructed using each element in isolation (with the height of the Gaussian set to 1)—up to the response time (Fig. S1). QT_k is then proportional to the time spent sampling element k .

The regression model we used was as follows, where in each case:

$$p(\text{high}) = \phi \left[b + \sum_{k=1}^8 w1_k \bullet (V_k - R) + \sum_{k=1}^8 w2_{k,t} \bullet TQ_{k,t} + \sum_{k=1}^8 w3_{k,t} \bullet (TQ_{k,t} \bullet (V_k - R)) \right],$$

where ϕ denotes the cumulative normal distribution, k the sample number, and t the time point of the QT vector. For the data shown in Fig. 4 A–C, we dropped the indexing by t and used the sum of QT for all time points up to the response.

ACKNOWLEDGMENTS. This work was supported by European Research Council Starter Grant 281628 (to C.S.), Economic and Social Research Council Funding Award 1369566 (to H.V.), and Wellcome Trust Grant 0099741/Z/12/Z (to S.H.C.).

- Gold JJ, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Ratcliff R, McKoon G (2008) The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Comput* 20(4):873–922.
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113(4):700–765.
- Gold JJ, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5(1):10–16.
- Beck JM, et al. (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60(6):1142–1152.
- Shadlen MN, Kiani R (2013) Decision making as a window on cognition. *Neuron* 80(3):791–806.
- Newsome WT, Paré EB (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* 8(6):2201–2211.
- Eckstein MP (2011) Visual search: A retrospective. *J Vis* 11(5):14.
- Torralba A, Oliva A, Castelholano MS, Henderson JM (2006) Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychol Rev* 113(4):766–786.
- Towal RB, Mormann M, Koch C (2013) Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proc Natl Acad Sci USA* 110(40):E3858–E3867.
- Vö ML-H, Henderson JM (2010) The time course of initial scene processing for eye movement guidance in natural scene search. *J Vis* 10(3):1–13.
- Itti L, Baldi P (2009) Bayesian surprise attracts human attention. *Vision Res* 49(10):1295–1306.
- Eckstein MP, Schoonveld W, Zhang S, Mack SC, Akbas E (2015) Optimal and human eye movements to clustered low value cues to increase decision rewards during search. *Vision Res* 113(Pt B):137–154.
- Navalpakkam V, Koch C, Rangel A, Perona P (2010) Optimal reward harvesting in complex perceptual environments. *Proc Natl Acad Sci USA* 107(11):5232–5237.
- Manohar SG, Husain M (2013) Attention as foraging for information and value. *Front Hum Neurosci* 7:711.
- Krajbich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13(10):1292–1298.
- Krajbich I, Rangel A (2011) Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc Natl Acad Sci USA* 108(33):13852–13857.
- Cassey TC, Evens DR, Bogacz R, Marshall JA, Ludwig CJ (2013) Adaptive sampling of information in perceptual decision-making. *PLoS One* 8(11):e78993.
- Shimojo S, Simion C, Shimojo E, Scheier C (2003) Gaze bias both reflects and influences preference. *Nat Neurosci* 6(12):1317–1322.
- Rehder B, Hoffman AB (2005) Eyetracking and selective attention in category learning. *Cognit Psychol* 51(1):1–41.
- Meier KM, Blair MR (2013) Waiting and weighting: Information sampling is a balance between efficiency and error-reduction. *Cognition* 126(2):319–325.
- Gottlieb J, Hayhoe M, Hikosaka O, Rangel A (2014) Attention, reward, and information seeking. *J Neurosci* 34(46):15497–15504.
- Haberman J, Whitney D (2010) The visual system discounts emotional deviants when extracting average expression. *Atten Percept Psychophys* 72(7):1825–1838.
- de Gardelle V, Summerfield C (2011) Robust averaging during perceptual judgment. *Proc Natl Acad Sci USA* 108(32):13341–13346.
- Michael E, de Gardelle V, Nevado-Holgado A, Summerfield C (2015) Unreliable evidence: 2 sources of uncertainty during perceptual choice. *Cereb Cortex* 25(4):937–947.
- Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* 164(1):177–190.
- Yarbus AL (1967) *Eye Movements During Perception of Complex Objects* (Springer, New York).
- Najemnik J, Geisler WS (2008) Eye movement statistics in humans are consistent with an optimal search strategy. *J Vis* 8(3):1–14.
- Brezis N, Bronfman ZZ, Usher M (2015) Adaptive spontaneous transitions between two mechanisms of numerical averaging. *Sci Rep* 5:10415.
- Navalpakkam V, Itti L (2007) Search goal tunes visual features optimally. *Neuron* 53(4):605–617.
- Scolari M, Serences JT (2009) Adaptive allocation of attentional gain. *J Neurosci* 29(38):11933–11942.
- Markant DB, Settles B, Gureckis TM (2016) Self-directed learning favors local, rather than global, uncertainty. *Cogn Sci* 40(1):100–120.
- Tsetos K, et al. (2016) Economic irrationality is optimal during noisy decision making. *Proc Natl Acad Sci USA* 113(11):3102–3107.
- Howes A, Warren PA, Farmer G, El-Deredey W, Lewis RL (2016) Why contextual preference reversals maximize expected value. *Psychol Rev* 123(4):368–391.
- Inhoff AW, Rayner K (1986) Parafoveal word processing during eye fixations in reading: Effects of word frequency. *Percept Psychophys* 40(6):431–439.
- Henderson JM, Ferreira F (1990) Effects of foveal processing difficulty on the perceptual span in reading: Implications for attention and eye movement control. *J Exp Psychol Learn Mem Cogn* 16(3):417–429.
- He PY, Kowler E (1989) The role of location probability in the programming of saccades: Implications for “center-of-gravity” tendencies. *Vision Res* 29(9):1165–1181.
- Morvan C, Maloney LT (2012) Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS Comput Biol* 8(2):e1002342.
- Dodge R (1900) Visual perception during eye movement. *Psychol Rev* 7(5):454.
- Cheadle S, et al. (2014) Adaptive gain control during human perceptual choice. *Neuron* 81(6):1429–1441.