The Formation of Hierarchical Decisions in the Visual Cortex

Highlights

- Monkeys can integrate information in parallel for multiple decisions

- During a later integration phase, different decisions influence each other

- Hierarchical decisions influence neurons in V1 and V4 of the visual cortex

- The interaction between decisions maximizes reward income

Authors

Jeannette A.M. Lorteije, Ariel Zylberberg, Brian G. Ouellette, Chris I. De Zeeuw, Mariano Sigman, Pieter R. Roelfsema

Correspondence

p.roelfsema@nin.knaw.nl

In Brief

Lorteije et al. found behavioral and physiological evidence that the primate brain can process multiple decisions in parallel. Besides decreasing overall processing time, parallel decisions can influence each other and, thereby, maximize the reward income.

CrossMark

CellPress

The Formation of Hierarchical Decisions in the Visual Cortex

Jeannette A.M. Lorteije,[1,2,11] Ariel Zylberberg,[1,3,4,5,11] Brian G. Ouellette,[1] Chris I. De Zeeuw,[6,7] Mariano Sigman,[3,8] and Pieter R. Roelfsema[1,9,10,*]

[1]Department of Vision and Cognition, Netherlands Institute for Neuroscience, Royal Netherlands Academy of Arts and Sciences, Meibergdreef 47, 1105 BA Amsterdam, the Netherlands
[2]Center for Neuroscience, Swammerdam Institute for Life Sciences, Faculty of Science, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, the Netherlands
[3]Laboratory of Integrative Neuroscience, Physics Department, Buenos Aires University, Intendente Güiraldes 2160, 1428 Buenos Aires, Argentina
[4]Institute of Biomedical Engineering, Faculty of Engineering, Buenos Aires University, Avenue Paseo Colón 850, 1063 Buenos Aires, Argentina
[5]Laboratory of Applied Artificial Intelligence, Computer Science Department, Facultad de Ciencias Exactas y Naturales, Buenos Aires University, Intendente Güiraldes 2160, 1428 Buenos Aires, Argentina
[6]Cerebellar Coordination and Cognition Group, Netherlands Institute for Neuroscience, Royal Netherlands Academy of Arts and Sciences, Meibergdreef 47, 1105 BA Amsterdam, the Netherlands
[7]Department of Neuroscience, Erasmus Medical Center, Wytemaweg 80, 3015 CN Rotterdam, the Netherlands
[8]Universidad Torcuato Di Tella, Almirante Juan Saenz Valiente 1010, C1428BIJ Buenos Aires, Argentina
[9]Department of Integrative Neurophysiology, Vrije Universiteit Amsterdam, De Boelelaan 1085, 1081 HV Amsterdam, the Netherlands
[10]Psychiatry Department, Academic Medical Center, 1105 AC Amsterdam, the Netherlands
[11]Co-first author
*Correspondence: p.roelfsema@nin.knaw.nl
http://dx.doi.org/10.1016/j.neuron.2015.08.015

SUMMARY

Intelligence relies on our ability to find appropriate sequences of decisions in complex problem spaces. The efficiency of a problem solver depends on the speed of its individual decisions and the number of decisions it can explore in parallel. It remains unknown whether the primate brain can consider multiple decisions at the same time. We therefore trained monkeys to navigate through a decision tree with stochastic sensory evidence at multiple branching points and recorded neuronal activity in visual cortical areas V1 and V4. We found a first phase of decision making in which neuronal activity increased in parallel along multiple branches of the decision tree. This was followed by an integration phase where the optimal overall strategy crystallized as the result of interactions between local decisions. The results reveal how sensory evidence is integrated efficiently for hierarchical decisions and contribute to our understanding of the brain mechanisms that implement complex mental programs.

INTRODUCTION

To solve a complex problem, we generally need to execute a nested sequence of actions. These problems are usually represented concisely as a decision tree, with a branching point for every decision (Anderson and Lebiere, 1998; Newell, 1990).

Consider the eight-puzzle of Figure 1, where a sequence of two correct decisions leads to the goal state. To solve the problem, one has to find an appropriate path in the problem space that corresponds to a sequence of decisions. Classical approaches such as breadth-first (starting with an evaluation of decisions at the top of the tree) and depth-first search (first following branches into the depth of the tree) examine the consequences of one decision at a time (Russell et al., 1995), but the efficiency of the problem solver can increase substantially when multiple decisions are explored in parallel. Here we asked how the visual cortex of monkeys forms complex, hierarchical decisions. Recent studies in human observers addressed the learning of hierarchical tasks and distinguished between model-based and model-free learning (Beierholm et al., 2011; Daw et al., 2005, 2011; Lee et al., 2014; Smittenaar et al., 2013). This study is different because we addressed the formation of hierarchical decisions in monkeys at the timescale of trials of a familiar task, and we did not examine the learning process.

Previous studies have demonstrated that single perceptual decisions may unfold over tens to hundreds of milliseconds (Gold and Shadlen, 2007; Romo and Salinas, 2003; Schall, 2001). The neuronal mechanisms for a single decision are well characterized by diffusion models that accumulate sensory evidence up to a bound, which signals the commitment to a choice (Drugowitsch et al., 2012, 2014; Huang et al., 2012; Ratcliff and McKoon, 2008). Previous studies that uncovered the neuronal processes for single decisions did not examine the neuronal mechanisms for hierarchical decision-making problems where the output of one decision is the input to the next (Anderson and Lebiere, 1998; Duncan, 2010; Zylberberg et al., 2011). We were particularly interested in problems where the decision maker has to consider multiple sources of stochastic evidence.

Figure 1. Example of a Sequential Decision-Making Task
This eight-puzzle is solved by a sequence of two decisions. It is important to consider the deeper levels of the decision tree before making the first decision. Note that the figure represents a small fraction of the overall problem space.

In this situation, serial decision-making strategies might be prohibitive because the delays associated with successive decisions would add up. Will the brain use its massively parallel architecture to speed up problem-solving by considering multiple decisions at the same time?

To answer this question, we recorded neuronal activity in cortical areas V1 and V4 of monkeys that had to solve an explicit decision tree with stochastic evidence at every branching point. We found that the decision-making process was accompanied by neuronal selection signals in these areas because branches in the decision tree that were selected by the animal elicited stronger neuronal responses than branches that were not. Both behavioral and neuronal data revealed that multiple local decisions initially formed in parallel but were later integrated into a single global strategy. Furthermore, the data revealed interactions between the local decisions, which enabled a global strategy that maximized the reward income under uncertain sensory conditions.

RESULTS

Behavioral Performance
We trained three monkeys to solve an explicit decision tree with a series of bifurcation points (levels L1 and L2, Figure 2A; Movie S1). To solve the task, the monkeys had to select the brightest segment at the first bifurcation point (L1) and then the brightest segment at the second L2 bifurcation. The animals could ignore the third bifurcation (we will use L2′ as notation) connected to the darkest path at the L1 bifurcation. The monkeys were rewarded for making a single eye movement to the circle at the end of the correct path after a fixation delay of 500 ms. The sensory evidence for every decision was stochastic because the luminosity of the L1/L2 contour elements fluctuated around a base luminance at 20 Hz (50 ms/sample, Figure 2B). We adjusted the difficulty of the task by varying the difference in base luminance between the target (T) and distractor branches (D) (Figures 2A and 2B), which was highest in easy trials, smaller in intermediate trials, and smallest in difficult trials. In most trials, both decisions

were correct (branch TT, Figure 2C). In error trials, we inferred whether mistakes were made at the L1 (eye-movement to DT) or L2 bifurcation (TD) or both (DD). The monkeys made slightly more erroneous eye movements to the TD than to the DT target (Figure 2C). We determined the psychometrical functions at the L1 and L2 bifurcations by computing the strength of evidence in each trial (abscissa, Figure 2D). We weighted the luminance time series with psychophysical kernels (derived below) to emphasize time points with a stronger influence on the decision. As predicted by the lower error rate, the psychometric curve for the L1 decision was steeper than that for the L2 decision (Figure 2D).

Time Course of the Formation of the Decisions
We used a logistic regression analysis to assess the influence of sensory evidence at different time points on the decisions (Ahumada, 1996; Neri and Heeger, 2002; Figure 2E). Positive fluctuations in luminance should increase accuracy when they occur on the target branch and decrease accuracy when they occur on the distractor branch, but the influence of luminance ceases when the decision has been made. As found previously in studies of single decisions (Kiani et al., 2008; Ludwig et al., 2005), we observed that the first few samples of the target and distractor branches were most influential and that the impact of later samples on the animal's choice faded out gradually, presumably because they occurred when the decision had already been made. Importantly, the psychophysical kernels were similar for the L1 and L2 decisions, which suggests that they were taken in parallel (Ludwig et al., 2014). In trials with reaction times (RTs) longer than the median, the L1 and L2 kernels were protracted in time, but their shape remained similar (Figure S1D). We devised a bootstrap analysis to test the difference between the time courses of the L1 and L2 decisions. This test did not reveal a significant difference (p > 0.2, Figure S1A). Although this result is indicative of parallel decisions in the visual cortex, it is not conclusive. Specifically, it is possible that the monkeys take both decisions sequentially in each trial but that the order varies across trials (we will call this a "mixed strategy"). To address this possibility, we ran a control experiment with a single bifurcation (Figure 2F). According to a mixed strategy, the animals would start with L2′ in a portion of trials and then proceed to L1 and L2, implying a total of three sequential decisions in these trials. The mixing of trials with decision orders should cause a protracted accumulation of evidence compared with the control task. However, we found that the kernels of the control and hierarchical tasks were similar (L1, p > 0.4; L2, p > 0.5; Figures S1B and S1C). This result is incompatible with a mixed strategy (we will present neuronal evidence against mixed strategies below).

A Hierarchical or a Flat Decision Strategy?
We also considered the possibility that monkeys did not evaluate the hierarchical structure of the branching task but used a simpler heuristic instead. They might ignore the hierarchical structure altogether and treat the problem as a "flat" competition between four eye movement targets, as proposed recently for tasks with multiple alternatives without a hierarchical structure (Bollimunta and Ditterich, 2012). To test the plausibility of such a flat heuristic, we compared two models. The first (flat) model
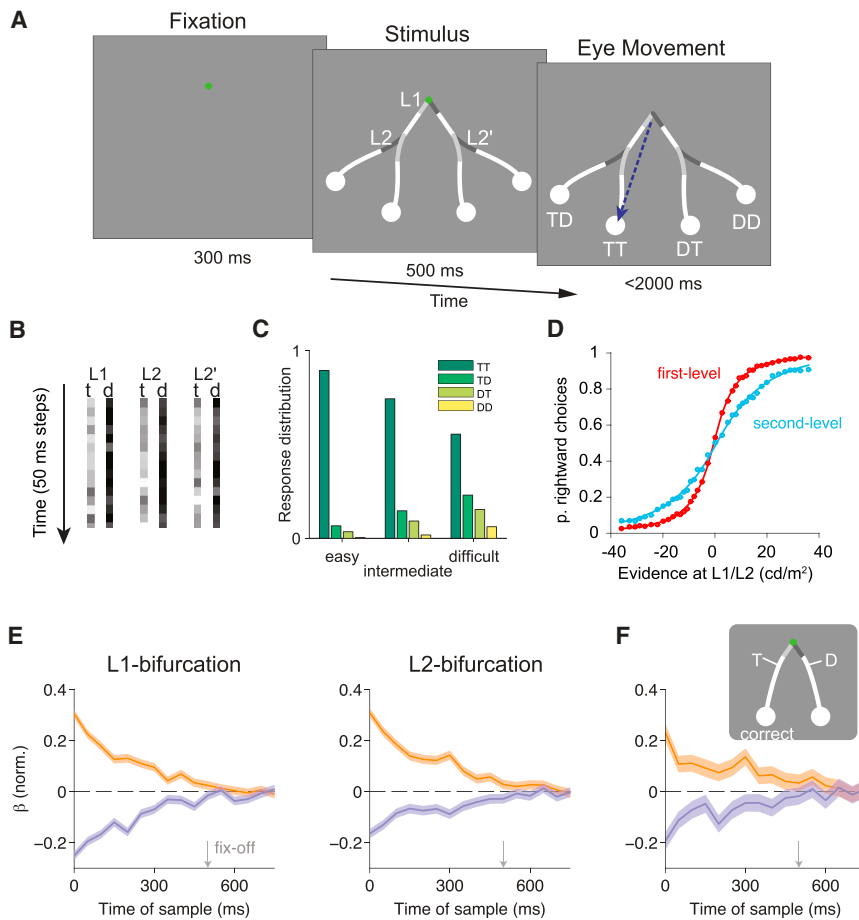
Figure 2. Hierarchical Decision-Making Task and Performance

(A) Behavioral task used to test hierarchical decision making in monkeys. After a fixation epoch, a stimulus appeared, with three bifurcations (L1, L2, and L2′) leading toward four circular markers. The two line segments beyond every bifurcation varied in luminance, and one was on average brighter (target) than the other (distractor). The monkey had to trace the curve from the fixation point (while maintaining fixation) and select the brightest path at each bifurcation. After 500 ms, the fixation point disappeared, cueing the monkey to make an eye movement to a marker (blue arrow). White contour segments and markers had a constant luminance, and the flickering segments are shown in gray. The letters at the right indicate the possible choices at the bifurcations: TT, correct choice at both bifurcations; TD, correct at L1 but error at L2; DT, error at L1 but correct at L2′; DD, both choices wrong.

(B) Example luminance fluctuations (50 ms/sample) at the flickering segments.

(C) Distribution of choices at the three difficulty levels.

(D) Psychometric curves that estimate sensory evidence by convolving the luminance signals (mean + fluctuations) with the psychophysical kernels (see Experimental Procedures). We constructed 39 bins with approximately equal numbers of trials based on the strength of evidence (~4,000 trials/data point). Curves are fits to the behavioral data.

(E) Influence of luminance fluctuations of the target (orange) and distractor (blue) segments at different time points on the L1 (left) and L2 decisions (right) as measured with a logistic regression. Shades indicate 95% confidence intervals for the regression coefficients. norm, normalized.

(F) The same analysis in a control task with a single decision.

See Movie S1 for an example stimulus and Figures S1 and S8.

is a diffusion-to-bound model with four decision variables, one for each choice (Figures 3A and 3B). Each decision variable (dv1–dv4, Figure 3B) summed luminance information across the two relevant branching points until it reached the bound. In the second (hierarchical) model, the accumulation of evidence for each bifurcation proceeded independently, followed by a second stage in which the output of the local accumulators was combined into a strategy (Figure 3D). To fit the models, we exploited the stochastic nature of the sensory stimulus using the sequence of luminances shown to the monkeys (as in Zylberberg et al., 2012). As expected, both models provide an accurate fit to the psychometric functions at the two bifurcations (Figures 3C and 3E; $R^2 > 0.99$ for every fit). However, the models make different predictions about how accuracy at one bifurcation influences accuracy at the other one. The flat model combines evidence for the L1 and L2 decisions and it therefore predicts that strong evidence at one bifurcation leads to a lower accuracy at the other one (Figure 3F; see Figure S2 for data of the individual monkeys; $p < 5.10^{-4}$, permutation test). Instead, the hierarchical model predicts that evidence for L1 and L2 accumulates independently ($p > 0.5$, permutation test) before they are combined

to determine the eye movement target. Therefore, L1 difficulty should not affect accuracy at L2. The monkeys' strategy was compatible with the hierarchical model and incompatible with the flat model because the L2 accuracy did not depend on L1 difficulty (permutation test, $p > 0.5$; see Experimental Procedures; Figure 3F).

Neuronal Correlates of Decision Making in the Visual Cortex

What are the neuronal mechanisms underlying this parallel decision-making process? Previous studies have examined decision making in the parietal and frontal cortex (Ding and Gold, 2012; Gold and Shadlen, 2000; Kiani and Shadlen, 2009; Peck et al., 2009; Shadlen and Newsome, 2001), but in these areas, receptive fields (RFs) are large, which would make it difficult to monitor the L1 and L2 decisions separately. Therefore, in this study, we focused on areas V1 and V4 of the visual cortex, where RFs are smaller. We measured multi-unit spiking activity (MUA) with chronically implanted electrode arrays at a total of 109 recording sites in V1 in three monkeys and 40 sites in V4 in two of these monkeys. We constructed the stimuli so that the neurons' RFs
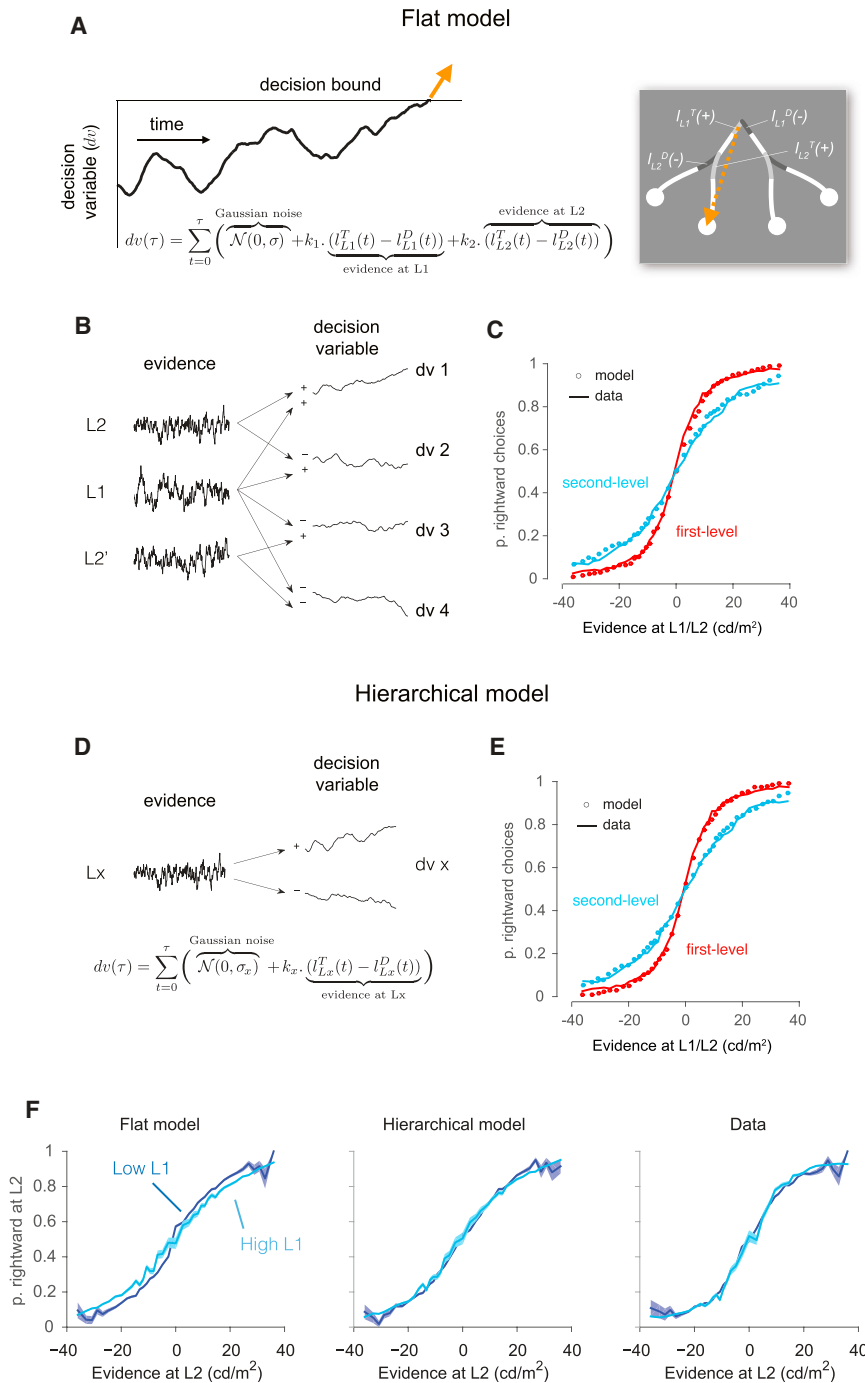
A

Flat model

decision bound

time

decision variable (dv)

$$dv(\tau) = \sum_{t=0}^{\tau}\left(\overbrace{\mathcal{N}(0,\sigma)}^{\text{Gaussian noise}} + k_1 \cdot \underbrace{(l_{L1}^{T}(t) - l_{L1}^{D}(t))}_{\text{evidence at L1}} + k_2 \cdot \overbrace{(l_{L2}^{T}(t) - l_{L2}^{D}(t))}^{\text{evidence at L2}}\right)$$

B

evidence

decision variable

L2

L1

L2'

+ + → dv 1

− + → dv 2

+ − → dv 3

− − → dv 4

C

p. rightward choices

○ model
— data

second-level

first-level

Evidence at L1/L2 (cd/m²)

Hierarchical model

D

evidence

decision variable

Lx

+ → dv x

−

$$dv(\tau) = \sum_{t=0}^{\tau}\left(\overbrace{\mathcal{N}(0,\sigma_x)}^{\text{Gaussian noise}} + k_x \cdot \underbrace{(l_{Lx}^{T}(t) - l_{Lx}^{D}(t))}_{\text{evidence at Lx}}\right)$$

E

p. rightward choices

○ model
— data

second-level

first-level

Evidence at L1/L2 (cd/m²)

F

Flat model

Hierarchical model

Data

p. rightward at L2

Low L1

High L1

Evidence at L2 (cd/m²)

Figure 3. Comparison between a Flat and a Hierarchical Decision-Making Model

(A) In the flat model, the decision process is modeled as a direct competition between the four alternative choices as a race between decision variables (one shown here). Each accumulator integrates luminance evidence across the L1 and the relevant L2 branching points with weights k_1 and k_2, respectively.

(B) Sketch of the flat model. The four decision variables (one for each possible shape of the target curve) are supported by different combinations of luminance evidence at the relevant bifurcations.

(C) Fits of the flat model to the psychometric functions for the first- (red) and second-level (cyan) decisions. Curves show data, and circles model fits.

(D) Sketch of a hierarchical model where decisions are made independently at each bifurcation. Two accumulators compete at each bifurcation, integrating luminance information (with opposite signs) and independently sampled Gaussian noise. The equation refers to the decision variable of the accumulator.

(E) Fits of the hierarchical model to the psychometric functions.

(F) Split of the second-level psychometric function by the strength of the evidence at L1 (median split). The flat model predicts higher accuracy at L2 when the evidence at L1 is low. The hierarchical model and the data do not show this bias. Shaded regions indicate SEM.

See Figure S2 for the data of individual monkeys.

fell on contour elements of constant luminance, downstream of the L1 bifurcation or the L2 bifurcation in the decision tree (Figures 4A and 4B). Figure 4A illustrates neuronal activity downstream of the L1 bifurcation. The target branch evoked stronger neuronal activity than the distractor branch in a time window of 200–500 ms in V1 and in V4 (V1: $p < 10^{-6}$, n = 92 recording sites; V4: $p < 5.10^{-5}$, n = 27, Wilcoxon signed-rank test). There was a weak 60-Hz oscillation visible in the average response in V1 (but not in V4), reflecting entrainment to the 60-Hz monitor refresh
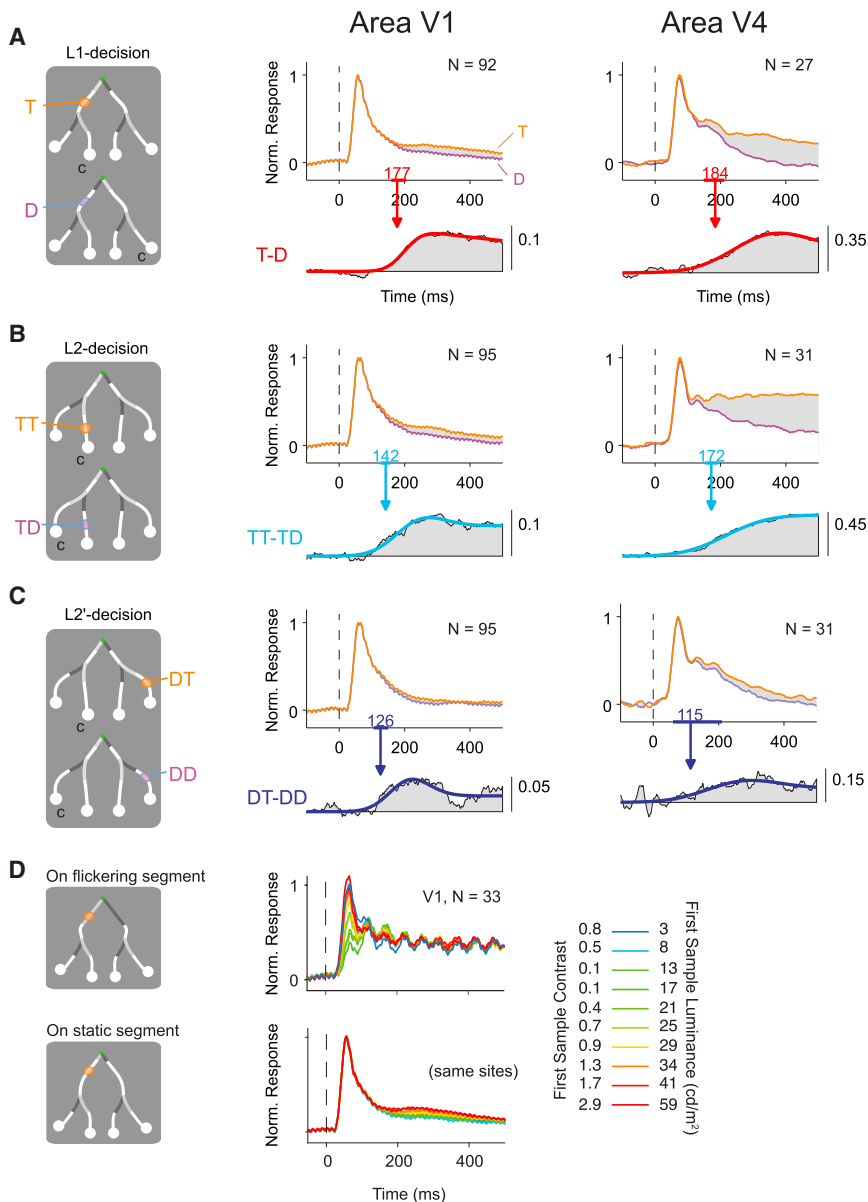
rate (Williams et al., 2004). Figure 4B illustrates neuronal activity elicited by the target and distractor branch downstream of the relevant L2 bifurcation. Again, the response evoked by the target branch was stronger (TT versus TD, Figure 4B; V1: $p < 10^{-6}$, n = 95; V4: $p < 5.10^{-6}$, n = 31). Therefore, V1 and V4 neurons carry signals related to the selection of target branches downstream of the L1 and L2 bifurcations.

To confirm that these response modulations were caused by the selection of target branches and not by a direct effect of flicker in the vicinity of the RF, we compared V1 responses elicited by the flickering segments to those elicited by the static contours with a constant luminance (n = 33 sites). We sorted trials based on the luminance of the first flicker sample. When the RFs of V1 neurons fell on a flickering segment, the influence of luminance appeared early, with a latency of ∼40 ms (Figure 4D), and neuronal activity was entrained by the flicker (Williams et al., 2004). Furthermore, the V1 response increased with contrast irrespective of whether it was caused by luminance values lower or higher than the background (15 cd/m²). The V1 responses evoked by the static segments were qualitatively

Figure 4. Time Course of Selection Signals in V1 and V4

(A–C) Receptive fields of neurons in V1 and V4 fell on the target (orange) or distractor curve (purple) beyond the L1 bifurcation (A), the relevant L2 bifurcation (B), or the irrelevant L2′ bifurcation (C). In correct trials, the average MUA response in V1 (left) and V4 (right) evoked by the target branches was stronger than that evoked by the distractor branches. The bottom panels show the selection signal computed as the difference in activity evoked by the target and distractor branches. We fitted a curve to this difference signal to estimate the latency. Colored horizontal lines on the x axis indicate 95% confidence intervals of the latency estimate.

D) Responses recorded from 33 V1 sites with RFs on the flickering segment (top) or on the static segment (bottom) were split according to the luminance of the first flicker sample in each trial. Colors indicate the luminance level of the first flicker sample (numbers on the right) and the Weber contrast (numbers on the left; $|(L_{segment} - L_{background})|/L_{background}$; where $L_{background} = 15$ cd/m^2).

See also Figures S3–S6.

different. Here the V1 activity increased monotonically with luminance in accordance with the selection rule that the monkeys had learned, and the influence of luminance appeared after a delay (Figure 4D; Figure S3). Therefore, the activity elicited by the static segments (Figures 4A–4C) depended on the selection of branches and not on a direct influence of flicker near the RF.

To investigate whether the late modulation of visual cortical activity is informative regarding different stages of decision making, we capitalized on the fact that we could deduce the locus of errors in the decision tree (Figure 5). When the monkey made an error at the L1 bifurcation, V4 activity evoked by the distractor segment downstream of this bifurcation was stronger than that evoked by the target (p < 0.05, Wilcoxon signed-rank test), and there was a trend for stronger distractor activity in V1 (p = 0.054) (Figure 5A). Similarly, when the monkey made an error

at the L2 bifurcation, V1 and V4 activity elicited by the corresponding distractor segment was stronger than that evoked by the target segment (p < 0.05 for both) (Figure 5B). Therefore, visual cortical neurons indeed convey information about the branches that are selected during the task.

To investigate whether the neurons also tracked the time course of decision making, we compared neuronal activity between trials where the first luminance sample pointed in the correct or wrong direction (luminance target higher/lower than distractor). The response modulation was strongest (L1 decision: V1, $p < 10^{-6}$ and V4, p < 0.05; L2 decision: V1, p < 0.005 and V4, p < 0.006; Wilcoxon signed-rank test) and appeared earliest (L1 decision: V1, $p < 10^{-6}$; L2 decision: V1, p < 0.005 and V4, p < 0.0005; not significant for the L1 decision in V4, p = 0.13; permutation test) when the first luminance sample pointed in the right direction (Figure S4). This result implies that the modulation of activity in V1 and V4 does not only reflect the outcome of the decisions but that it tracks the unfolding decision process at a high temporal and spatial resolution (because of the small RFs). Therefore, early and mid-level visual areas may contribute to the selection of contour elements, but our results are not incompatible with scenarios where feedback from higher brain regions modulates activity in V1 and V4 (Khayat et al., 2009; Pooresmaeili et al., 2014).

Previous studies usually attributed the modulation of activity in the early and mid-level visual cortex to shifts of attention (Ghose
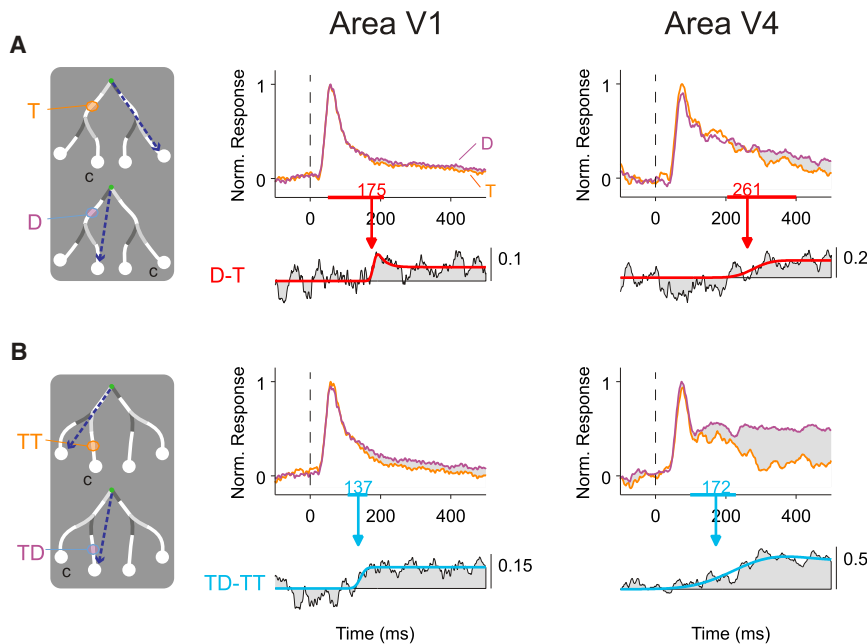
Figure 5. Population Responses for Error Trials under the Difficult Condition

(A) Errors at L1. During error trials, V4 activity evoked by the L1 distractor branch was stronger than that evoked by the target (p < 0.05, Wilcoxon signed-rank test), and in V1 there was a trend in the same direction (p = 0.054). The arrows in the left panel indicate the erroneous saccade.

(B) Errors at L2. We included trials with a correct L1 decision and an incorrect L2 decision. In V1 and V4, the L2 distractor branch elicited stronger activity in error trials (p < 0.05 for both). The analyses only included trials from sessions with at least eight incorrect trials.

these methods yielded similar results (Supplemental Experimental Procedures; Figure S6).

A prediction of parallel models is that evidence should accumulate for all decision points, even for branches that are not part of the correct path. In contrast, a serial algorithm would first take the L1 decision and then trim the tree, focusing on the relevant L2 decision (TT versus TD, Figure 2A) and ignoring the irrelevant L2′ decision. We therefore examined whether evidence accumulated for the irrelevant L2′ decision (DT versus DD). We found that neuronal activity for the DT branch was indeed stronger than for the DD branch, both in V1 (p < 10^{-6}, Wilcoxon signed-rank test) and V4 (p < 5.10^{-3}) (Figure 4C), in accordance with a parallel decision-making process.

Integration of the L1 and L2 Decisions into a Strategy

How are the three decisions integrated into a strategy? That is, how does the monkey select the eye-movement based on these local selection signals? Above, we mentioned the possibility of a flat strategy that predicts parallel evidence accumulation for each of the four possible shapes of the target curve. However, it is also conceivable that the decision process leverages on specific interactions between selection signals because committing to the L1 decision might have an effect on the two L2 decisions. A direct comparison of selection signals between the relevant L2 and irrelevant L2′ decisions revealed that the initial time courses of these decisions were similar, but, at a latency of 229 ms in V1 (with a 95% confidence interval of 175–273 ms) and 225 ms in V4 (95% confidence interval, 175–261 ms), the selection signal for the L2 branches became stronger than that for the L2′ branches (Figure 6; V1, p < 5.10^{-6}; V4, p < 5.10^{-5}; Wilcoxon signed-rank test; Figure S7). The difference between the latencies of the influence of L1 on the L2 decisions (TT-TD-[DT-DD]) and the L2 decision itself was significant (p < 0.05 in V1 and V4, Montecarlo procedure; Supplemental Experimental Procedures).

These results imply that committing to one of the L1 choices amplified the selection signals for the corresponding L2 decision. Therefore, after the initial phase of parallel decision making, V1 and V4 neurons started to code the outcome of the entire decision process after ~230 ms. The delayed influence of the L1 decision, which amplified the relevant L2 selection signal relative
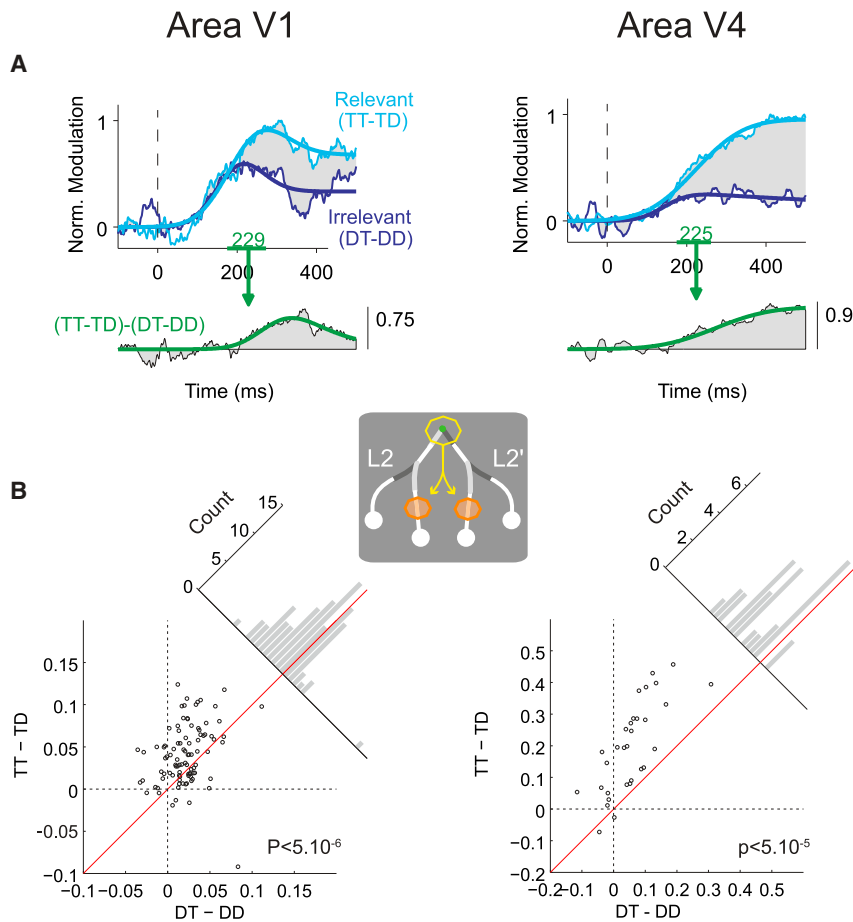
and Maunsell, 2002; Luck et al., 1997; Moran and Desimone, 1985; Motter, 1994; Roelfsema et al., 1998). The present results suggest that decision making causes a similar modulation of neuronal activity. The enhanced representation of contour elements in V1 and V4 appears to reflect the monkey's belief that they are part of the relevant path (Nienborg and Roelfsema, 2015). We can interpret this response modulation as a shift of attention and as a decision about the bifurcation point, and we do not know whether it is possible to dissociate these interpretations. In what follows, we will use the neutral term "neuronal selection signal" to describe the difference in activity evoked by selected and ignored contour elements.

Time Course of the Neuronal Selection Signals

Our behavioral results suggested that evidence for multiple decisions accumulates in parallel. To estimate the time course of the L1 and L2 decisions at the neuronal level, we subtracted the response evoked by the distractor branch from that evoked by the target branch, fitted a curve to the response difference, and estimated latency as the time point where the fitted function reached 33% of its maximum (Roelfsema et al., 2003; Figures 4A and 4B; see Figure S5 for data on individual monkeys). The latency of the L1 decision was 177 ms (95% confidence interval, 159–205 ms) in V1 and 184 ms in V4 (149–212 ms), and the latency of the L2 decision was 142 ms in V1 (126–156 ms) and 172 ms in V4 (153–193 ms) (Figures 4A and 4B). We could directly compare the latencies of the two decisions for a subset of recording sites and found that the latency of the L2 decision was slightly earlier than that of the L1 decision (V1, p < 0.02; V4, p < 0.05; permutation test). Nevertheless, the time courses of the selection signals at the two bifurcations overlapped strongly. Because the time point where the modulation reaches 33% of its maximum is somewhat arbitrary, we also used other criteria, including measures for the onset of modulation, but

Area V1

A

Area V4

(A) Selection signals for the relevant L2 (cyan) and irrelevant L2′ (dark blue) decisions. The selection signal for the relevant L2 decision became stronger at a latency of 229 ms in V1 and 225 ms in V4, reflecting the influence of L1 on the L2 selection signals (inset).
(B) Comparison of the decision-related activity for the relevant L2 decision (TT-TD) and irrelevant L2′ decision (DT-DD) across individual recordings sites in V1 (left) and V4 (right). Neuronal activity was averaged in a window of 200–500 ms after stimulus onset.
See also Figures S6 and S7.

B

$$Ev(tr, Lx) =$$
$$\sum_t w_{norm}(t)(Lum_{T,Lx}(tr, t) - Lum_{D,Lx}(tr, t)),$$

(Equation 1)

where $Lum_{T,Lx}(tr, t)$ and $Lum_{D,Lx}(tr, t)$ are target and distractor luminance at branch L_x in trial tr, and $w_{norm}(t)$ is the psychophysical kernel at time t.

We measured the difference in difficulty between the L2-decisions as

$$Diff(tr) = |Ev(tr, L2)| - |Ev(tr, L2')|$$
(Equation 2)

and computed the L1 psychometric function for the highest and lowest quartiles of the $Diff(tr)$ distribution (Figure 7A). This

to the irrelevant L2′ decisions, supports a hierarchical decision-making strategy and cannot be explained by a flat competition between four alternatives (Figures 3A and 3B).

An important advantage of the first phase of parallel decision making, in addition to its speed, is that it permits "backward influences" of decisions deeper in the tree onto earlier ones. Specifically, if sensory evidence is weak at the L1-bifurcation, then it is optimal to compare the relative difficulties of the two L2 bifurcations and select the L1 branch leading to the easier L2 decision. As an extreme example, consider the case where the evidence for the L1 decision is ambiguous, evidence for the left L2 decision is ambiguous too, but the right L2 decision is easy. In this situation, choosing one of the two targets on the left yields a reward with a probability of 25%, whereas choosing the target on the right yields a reward with a probability of 50%. We therefore investigated whether the monkeys' choices at the L1 bifurcation were sensitive to the relative difficulty of the two L2 decisions (Fleming et al., 2013). We estimated the quality of evidence $Ev(tr, L_X)$ for each trial tr at a bifurcation L_X during the epoch of evidence integration, weighting the time series of luminance differences between target and distractor by the psychophysical kernels. Therefore, luminance samples at the start of the trial received a stronger weight in our estimate of the quality of the evidence than samples that were presented later (Figure 2E):

analysis revealed a bias at L1 toward the easier L2 bifurcation ($p < 10^{-6}$, likelihood ratio test, nested logistic regression), which was most pronounced when the L1 decision was difficult. We estimated the time course of the backward influence with a logistic regression and found that it was similar to the time course of the L1 and L2 decisions themselves (Figure 7C). Therefore, the backward influence of the L2 decisions onto the L1 decision is an intrinsic feature of a strategy that builds on a parallel decision-making process. To determine the optimality of the resulting strategy, we fitted a model to the psychometric curves for the L1 and L2 decisions (continuous curves, Figure 2D) and used it to estimate the reverse influence that would maximize reward (Figure 7B). The reward maximization strategy required that the backward influence of L2 evidence is strongest when the L1 decision is difficult, just as observed in the monkeys' behavior.

A prediction of this tradeoff is that the neuronal selection signal at the (relevant) L2 bifurcation should not only depend on the local evidence (Figure S4) and the L1 decision (Figure 6) but also on the evidence at the (irrelevant) L2′ bifurcation. We therefore examined whether the luminance difference (T-D) of the first sample at the L2′ bifurcation influenced the neuronal selection signal at the L2 bifurcation. Specifically, we selected the 50% of trials with highest or lowest difference between target and
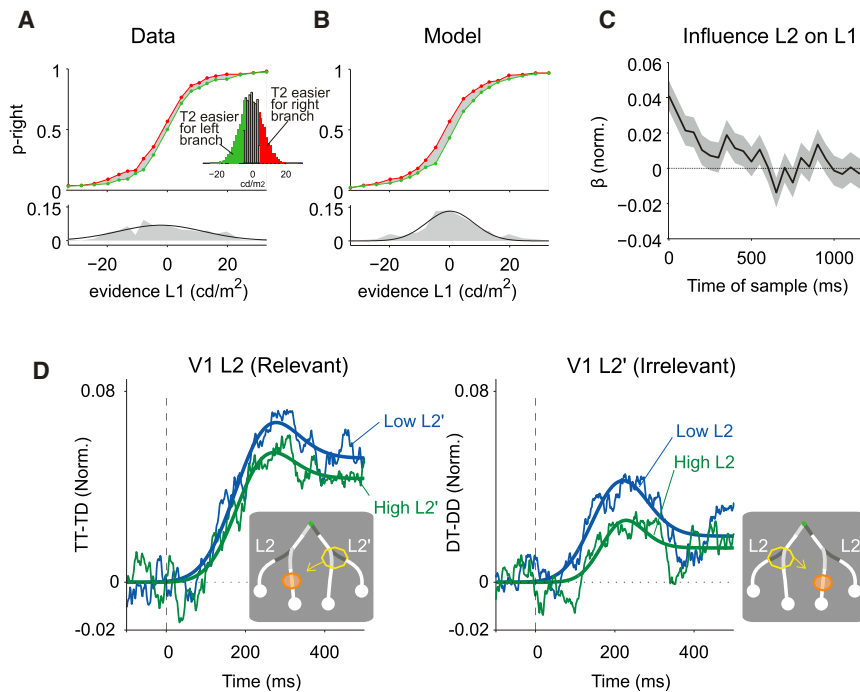
Figure 7. L1 Decisions Are Biased by L2 Difficulty

(A) We determined the difference in difficulty between the two L2 decisions in every trial (Equation 2) and selected 25% of trials for which the left (red) or right L2 decision (green) was easiest and calculated L1 psychometric functions for these trials. The bottom panel shows the difference between the psychometric functions (and a Gaussian fit).

(B) Same as (A) but for a model that optimally weighs the evidence at the three bifurcations to maximize reward.

(C) Time course of the backward influence of L2 difficulty on the L1 decision as estimated with logistic regression (black). The grey region indicates 95% confidence intervals for the regression coefficients.

(D) Neuronal correlates of the cross-talk between the two L2 decisions. Selection signals at L2 and L2′ are stronger when the sensory evidence conveyed by the first sample at the other L2 decision is weak (median split; blue curve, weaker evidence; green curve, stronger evidence at the other L2 decision).

distractor luminance (median split). A large luminance difference at the L2′ bifurcation indeed reduced the V1 selection signal at the L2 bifurcation (n = 95, p < 0.05, paired t test across recording sites). Vice versa, a large luminance difference at the L2 bifurcation weakened the V1 selection signal at the L2′ bifurcation (p < 0.05) (Figure 7D). This effect failed to reach significance in V4 (p > 0.2 for both, n = 31), possibly because of the smaller sample size. The backward influence of the L2 decisions onto the L1 decision is in accordance with the slightly shorter latency of the L2 selection signals (Figure 4) and provides additional evidence against a flat competition between four alternative shapes of the target curve (Figures 3A and 3B).

Neuronal Evidence against Mixed Strategies

The neuronal data do not only allow us to compare flat and hierarchical strategies. We also used the data to further examine the possibility of mixed strategies where the monkeys take the two decisions sequentially but in different orders across trials. If the monkeys first evaluate L1 in some of the trials but start with L2 in other trials, strong early neuronal selection signals at L1 should be associated with protracted evidence integration at L2. To test this possibility, we carried out two additional analyses. First, we repeated the behavioral reverse correlation analysis of Figure 2E separately for trials with strong and weak selection signals at L1 (median split of the neuronal response) in a time window of 200–300 ms. The psychophysical kernels of the two types of trials were largely overlapping (Figure S8A), which is not in accordance with a mixed strategy. Second, in trials in which monkeys evaluated L1 before L2, the activity at the L1 target branch should increase early but decrease later when they evaluate L2. The reversed order of events is predicted for trials in which the monkeys start with L2. Therefore, this mixed strategy predicts a negative correlation between activity in early and late time windows. When we computed the correlation between the neuronal responses in V1 and V4 between an early (200–300 ms) and late (400–500 ms) time window, we only observed positive correlations for L1 and L2 decisions (Figure S8B). These analyses support our conclusion based on the psychophysical kernels (Figure 2) that the monkeys' behavior is not in accordance with a mixed strategy.

DISCUSSION

The computational power of the primate brain is thought to derive from its massively parallel architecture and not from its speed because elementary cognitive operations are relatively slow (Anderson and Lebiere, 1998; Newell, 1990; Roelfsema, 2005). A key cognitive process is making a sensory decision. Previous studies have explored the mechanisms for single perceptual decisions and found that they involve computations in distributed brain regions. Sensory cortices represent the evidence bearing on alternative actions (Fetsch et al., 2014; Salzman et al., 1990), and areas of the parietal and frontal cortices integrate the evidence as persistent activity during motor planning (Churchland et al., 2011; Kim and Shadlen, 1999; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). The decision process terminates when the integral of the evidence reaches a threshold, in accordance with normative models of decision formation in simple tasks (Drugowitsch et al., 2012; Rao, 2010; Smith and Ratcliff, 2004; Wald and Wolfowitz, 1948).

In this study, we went beyond single decisions and investigated the mechanisms for hierarchical decisions. We obtained a number of convergent results, indicating that the exploration of a small decision tree can initially take place in parallel. First,

Figure 8. Summary of the Results
After the initial visual response (elicited in V1 at a latency of 40 ms, white curve), there is an early phase of parallel decision making (orange shading). After 225 ms, neurons start to reflect the interactions between the three decisions, which are required for the optimal strategy, so that the selected curve is highlighted with increased activity (yellow shading) in the visual cortex. Bottom: the neural activity elicited by the stimulus (white curve), whereas the orange and yellow curves show the time courses of the initial parallel decision-making phase and the final strategy, respectively.

the reverse correlation analysis revealed that information for all decisions was integrated at the same time, and we observed the simultaneous development of neuronal selection signals for multiple target branches in the visual cortex. Second, the neurons initially also selected the target branch of the irrelevant bifurcation, which was eventually discarded. Third, the decision at the initial bifurcation was biased by sensory evidence at deeper bifurcations to yield an overall strategy that optimized the total amount of reward.

Information about the individual branches was eventually combined into a strategy through interactions between neuronal selection signals for the three decisions. The L2 and L2′ decisions influenced each other as well as the L1 decision, and, vice versa, the L1 decision amplified the neuronal selection signal for the relevant L2 decision (Figure 8). As a result, the reward income was higher than what can be achieved by schemes that make the L1 decision before considering L2. Our results therefore demonstrate how the primate brain can navigate efficiently through problem spaces. It is equipped with a parallel and context-sensitive process that integrates hierarchical decisions into an optimal strategy, thereby compensating for the protracted time course of individual decisions.

Alternative Strategies for Hierarchical Decisions

Classical methods in computer science explore a decision tree by considering one decision at a time. In contrast, the monkeys accumulated evidence for the decisions at the three branching points in parallel. Furthermore, evidence deeper in the decision tree influenced their first decision, particularly when the evidence was ambiguous.

We also considered alternative decision strategies. The first is a serial, mixed strategy where the monkeys start with one decision in some trials but with another decision in others. We ob-

tained three sources of evidence against mixed strategies: the shape of the psychophysical kernel in a task with a single decision did not differ from that in the hierarchical decision task (Figure 2); the strength of the L2 psychophysical kernel was similar in trials with strong and weak early L1 selection signals; and there was a positive correlation between selection signals in early and late time windows for each of the branches (Figure S8). These results, taken together, imply that it is extremely implausible that the overlapping kernels result from mixing trials in which L1 and L2 decisions are performed sequentially but in different orders.

The second alternative strategy is a flat four-alternatives choice that disregards the hierarchical nature of the decision tree. The data did not support such a flat strategy either, for a number of reasons. First, the accuracy of the L2 decision did not depend on the strength of evidence for the L1 decision, as predicted by a flat model (Figure 3). Second, an early phase of parallel evidence accumulation was followed by a later integration phase when the selection signal for the L2 decision was amplified relative to that for the L2′ decision, which is not compatible with a flat competition model. Third, the relative difficulty of the L2 decisions exerted a backward influence on the L1 decision, and evidence at one L2 decision influenced the selection signal at the other one, again not in accordance with a flat model. Taken together, the results revealed a genuine hierarchical decision-making process characterized by a first parallel decision-making phase followed by an integration phase enabling an optimal eye movement decision.

Does the Visual Cortex Contribute to Decision Making?

Most previous studies on sensory decision making examined the activity of neurons in the parietal and frontal cortex (Churchland et al., 2011; Ding and Gold, 2012; Gold and Shadlen, 2000; Kiani and Shadlen, 2009; Kiani et al., 2008; Kim and Shadlen, 1999; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). This study differs because we recorded neuronal activity in areas V1 and V4, where neurons have smaller receptive fields, so that we could measure the selection of individual branches of the decision tree. Previous studies with the curve-tracing task suggested that the response modulation in the visual cortex is caused by shifts of object-based attention (Pooresmaeili et al., 2014; Roelfsema et al., 1998). The relationship between attentional selection signals and the neuronal correlates of decision making have, to our knowledge, not been explored extensively, although it is conceivable that the selection signals in the parietal and frontal cortex in previous decision-making tasks were also associated with shifts of spatial-, feature-, or object-based attention.

We do not know if these selection signals are generated in the visual cortex or in higher cortical areas. One possible scheme is that they formed within the visual cortex because they depended on luminance information. However, our results are also compatible with a scheme where the visual cortex provides the evidence but where the decision-making process takes place in higher brain regions (Law and Gold, 2008). This second scheme would imply a sophisticated feedback process, with early feedback from the decision areas selecting three branches in parallel and later feedback signals focusing on the chosen path

(Figure 8). Although our results cannot distinguish between these alternatives, they do provide an upper bound on the timing of the decisions. If the decision process occurs in higher cortical areas, then it needs to occur early enough so that there is sufficient time for the feedback signal to reach the visual cortex.

It is therefore remarkable that the first selection signals in V1 emerged early, after ~140 ms, during a phase when evidence accumulation was ongoing according to the behavioral results. Furthermore the V1 selection signals did not lag those in V4, unlike in other tasks with a clear backward progression of selection signals from higher to lower visual areas (Buffalo et al., 2010; Mehta et al., 2000) but in line with a recent curve-tracing study (Pooresmaeili et al., 2014), showing that selection in the frontal eye fields and V1 occurs at approximately the same time. These early selection signals are compatible with an active participation of the visual cortex in the decision. An advantage of such a participation is that the visual cortex provides an accurate map of the decision tree and can select one of a number nearby branches at a high spatial resolution (Mumford, 1992). Future studies could elucidate the contribution of the parietal and frontal cortex to hierarchical decision making, although this may require a paradigm that is better suited for neurons with large receptive fields.

Limits to Parallel Decision Making

Our data show, for the first time, that evidence for multiple decisions in a small, hierarchical problem space can initially form in parallel, before a second phase when the strategy is formed. Previous studies have indicated that multiple decisions are not always taken simultaneously. The brain switches to a serial decision-making strategy when individual decisions are more demanding (Moro et al., 2010), when the number of them exceeds a limit (Zylberberg et al., 2012), or when the sampling of evidence for different decisions requires eye movements (Chen et al., 2013).

Some of the limits to parallel decision making could be related to the maximum number of items that can be held in working memory (Luck and Vogel, 1997) because a sensory decision implies a neuronal memory element for integrating evidence. Indeed, the areas in the frontal and parietal cortex that have been implicated in perceptual decision making are also involved in working memory (Funahashi et al., 1989; Gnadt and Andersen, 1988), and it has been proposed that the neuronal substrates for decision making and working memory overlap (Soltani and Wang, 2010). Testing the limits of parallel decision making and identifying the relation between working memory and decision making may prove to be interesting topics for future research.

EXPERIMENTAL PROCEDURES

All experimental procedures complied with the NIH Guide for Care and Use of Laboratory Animals and were approved by the Institutional Animal Care and Use Committee of the Royal Netherlands Academy of Arts and Sciences. Three male macaque monkeys (*Macaca mulatta*; D, J, and A) participated in this study. In a first surgery, a head holder was implanted for head fixation during the experiments. In a separate surgery, microelectrode arrays (Blackrock Microsystems "Utah probes" with 4 × 4, 4 × 5, or 5 × 5 electrodes with a length of 1 or 1.5 mm) were implanted chronically in areas V1 and V4 under aseptic conditions and general anesthesia. Details of the surgical procedures

and postoperative care have been described elsewhere (Roelfsema et al., 1998, 2007).

Task and Stimuli

The monkeys were seated at a distance of 84 cm from a 21-inch cathode ray tube (CRT) monitor with a frame rate of 60 Hz. Trials started with a blank screen with a fixation point (green, diameter = 0.2°) in the center of the display (Figure 2A; Movie S1). The stimuli appeared when the monkey kept his gaze on the fixation point for 300 ms. The monkeys maintained the gaze within a fixation window of ~1° diameter, and trials were aborted when the animal broke fixation. The monkeys had to mentally trace a curve that bifurcated two times. The average eccentricities of the first and second bifurcation were 0.36° and 3.8°, respectively. When the stimulus had been in view for 500 ms, the fixation point disappeared, cueing the animal to make an eye movement to one of the four markers. If the animal maintained fixation after fixation point offset, then the stimuli remained on the screen for maximally 1,500 ms. However, in most trials, the animals responded within a few hundred milliseconds, and they were rewarded with a drop of juice when they chose the target marker (TT, Figure 2A). In case of an error, we added 500 ms to the inter-trial interval.

We stochastically varied the luminance of the two line segments beyond the bifurcation points every 50 ms (Figure 2B), adding luminance sampled from a Gaussian distribution ($\sigma = 10$ cd/m^2) to the base luminance. The base luminance of the distractor segment was equal to that of the gray background (15 cd/m^2). The base luminance of the target branch was higher and varied with difficulty. In easy trials, the base luminance of the target branch was 45 cd/m^2. It was 28 or 35 cd/m^2 in intermediate trials, and 22 or 28 cd/m^2 in difficult trials (the precise value for intermediate and difficult trials was adjusted across days to the monkeys' accuracy). The non-varying contours and target circles (with a diameter of 0.7°) were white. We varied the position and size of the stimuli across recording sessions to ensure that the center of the receptive fields fell on the static segments in the recording sessions that were included in the analysis. The location of the static segments was constant within sessions so that a receptive field could fall either on an L1 or an L2 segment during the entire session. Monkeys J and A also performed the task with only the L1 bifurcation (Figure 2F). We included a total of 176,185 trials in the analysis of the main task and 17,933 trials in the task with a single bifurcation.

Recording and Analysis of Multi-unit Activity in Areas V1 and V4

MUA was recorded in areas V1 and V4 as in previous studies (Cohen and Maunsell, 2009; Legatt et al., 1980; Logothetis et al., 2001; Xing et al., 2009) as the envelope of the signal filtered between 750 and 5,000 Hz. In short, the MUA signals were amplified, band-pass-filtered (750–5,000 Hz), full wave-rectified, low-pass-filtered at 500 Hz, and sampled at a rate of 763 Hz. The MUA signal represents the pooled activity of a number of neurons in the vicinity of the electrode tip. The population responses obtained with MUA recording are identical to those obtained by pooling across single units (Cohen and Maunsell, 2009; Supèr and Roelfsema, 2005).

We only included recording sites with a clear visual response. We required that the trial-averaged difference between the maximum visual response (window from 0–300 ms after stimulus onset) and the spontaneous activity (−150–0 ms from stimulus onset) in V1 (V4) was at least 150% (75%) of the SD of the trial-to-trial variation in spontaneous activity. In V1, MUA was recorded at a total of 92 recording sites with receptive fields beyond the L1 bifurcation (n = 59, 19, and 14 for D, J, and A, respectively) and from 95 sites with receptive fields beyond the L2 bifurcation (n = 62, 14, and 19 for D, J, and A, respectively). In V4, MUA was recorded from 27 sites beyond L1 (n = 14 and 13 in D and J, respectively) and 31 sites beyond L2 (n = 13 and 18 in D and J, respectively). The average eccentricities of receptive fields on the segments beyond the L1 bifurcation were 4.7° (V1) and 5.8° (V4), and the average eccentricity of the receptive fields beyond the L2 bifurcation were 4.9° (V1) and 5.6° (V4). The V1 and V4 recordings in monkeys D and J were obtained in interleaved sessions.

The monkeys performed between 800 and 3,000 trials per recording session (mean, 1,299; SD, 380; sessions with <800 trials were excluded). The neuronal responses were normalized to the average peak response after subtracting spontaneous activity. Population responses were computed by averaging across the normalized responses at individual recording sites. Data from a recording site were used only once, first averaging across days for that

recording site before averaging across recording sites. The MUA analysis for the L1 decision included all trials with a correct L1 decision irrespective of the L2 decision (TT and TD, Figure 4A). The MUA analysis for the L2 decision included only correct trials (TT, Figure 4B), except for the analysis of Figure 4C, which included trials with a correct L1 decision irrespective of the L2 decision.

The statistical significance of the selection signal was tested with a Wilcoxon signed-rank test in a window between 200–500 ms after stimulus onset. We described the method to determine the latency of the selection signals in V1 and V4 in the Supplemental Experimental Procedures.

Estimation of the Influence of Luminance Fluctuations on Behavior
We used a logistic regression model to measure the influence of luminance fluctuations on the decisions made at each bifurcation. We modeled the accuracy of the decision at the first bifurcation L1 as

$$Y = \left[1 + \exp\left(-\left(\beta_0 + \beta_{1,t}F_{T,t} + \beta_{2,t}F_{D,t} + \beta_{3,t}F_{dif,t} + \beta_4 L + \beta_{5,k}M_k\right)\right)\right]^{-1},$$
(Equation 3)

where β represents the fitted coefficients, Y indexes the accuracy of the decision, L is the base luminance of the target segment (candelas per square meter), M_k is a dummy variable identifying the monkey, $F_{T,t}$ and $F_{D,t}$ represent the luminance fluctuations of the target and distractor segments at time t, and $F_{dif,t}$ measures the ease of the relevant L2 decision relative to the ease of the irrelevant L2′ decision; i.e., $F_{dif,t} = (F_{T2,t} - F_{D2,t}) - (F_{T2',t} - F_{D2',t})$. To test the significance of the influence of L2 on L1 (Figure 7C), we performed a likelihood ratio test for nested logistic regression models with and without the β_3 term. We used the same model for the analysis of the single-decision task but without the β_3 term.

The regression model for the decision at the second bifurcation L2 was

$$Y = \left[1 + \exp\left(-\left(\beta_0 + \beta_{1,t}F_{T,t} + \beta_{2,t}F_{D,t} + \beta_{3,t}F_{first,t} + \beta_4 L + \beta_{5,k}M_k\right)\right)\right]^{-1}.$$
(Equation 4)

$F_{first,t} = (F_{T1,t} - F_{D1,t})$ modeled the possible influence of the luminance difference at the first bifurcation on the L2 decision, but we found that this term was negligible. In Figures 2E and 2F and 7C, the regression coefficients for the luminance fluctuations are shown relative to the base luminance (coefficients are divided by β_4).

Psychometric Curves
We constructed the psychometric curves (Figures 2D and 7A) by collapsing the time-varying luminance information to a single value per trial and branch. The weight of the luminance information presented at time t was measured by a regression analysis. For branch x (L1 or L2), the regression model was

$$Y_x = \left[1 + \exp\left(-\left(\beta_0 + \beta_{1,t}(F_{Tx,t} - F_{Dx,t}) + \beta_4 L + \beta_{5,k}M_k\right)\right)\right]^{-1}.$$
(Equation 5)

β represents the fitted coefficients, Y_x indexes the accuracy of the decision made for branch x, L is the base luminance of the target segment (candelas per square meter), M_k is a dummy variable identifying the monkey, and $F_{T,t}$ and $F_{D,t}$ represent the luminance (candelas per square meter) of the target and distractor segments at time t of branch x. Because we found that the effects of target and distractor luminance were of equal magnitude but opposite polarity (Figures 2E and 2F), Equation 5 evaluated the luminance difference $F_{T,t}-F_{D,t}$ (unlike Equations 3 and 4). We included the first ten samples (500 ms) in the regression. The weights $w_{norm}(t)$ in Equation 1 were obtained from the regression coefficients $\beta_{1,t}$, dividing each coefficient by the sum over the first ten samples:

$$w_{norm}(t) = \beta_{1,t} \Big/ \sum_{t=1}^{10} \beta_{1,t}.$$
(Equation 6)

We used the normalized weights $w_{norm}(t)$ to transform the time series of luminance values into a single value measuring the strength of the evidence (Ev) at the decision point Lx:

$$Ev(tr, Lx) = \sum_t w_{norm}(t)(Lum_{T,Lx}(tr,t) - Lum_{D,Lx}(tr,t)),$$
(Equation 7)

where $Lum_{T,Lx}(tr,t)$ and $Lum_{D,Lx}(tr,t)$ are target and distractor luminance at time t in trial tr.

To fit the psychometric curves of Figure 2D (and model optimal decisions, see below), we assumed that the binary decision at each bifurcation was made by computing the sign of an internal representation of the external evidence $Ev_{int}(tr) = f(Ev(tr))$. According to this model, the internal representation $Ev_{int}(tr)$ depends stochastically on the external evidence $Ev(tr)$, as in the Weber-Fechner law (Krueger, 1989). For simplicity, the model holds that the noise in $Ev_{int}(tr)$ has a Gaussian distribution with an SD that scales linearly with $Ev(tr)$:

$$Ev_{int}(tr) \sim N(Ev(tr) + bias, \sigma)$$

$$\sigma = a + b.Ev(tr),$$
(Equation 8)

where $N(\mu, \sigma)$ is a normal distribution with mean μ and standard deviation σ. We fitted the a, b, and $bias$ parameters to the psychometric curves in Figure 2D. The fits of the model were excellent because the fraction of unexplained variance was lower than 0.2% for both L1 and L2. Simpler models (e.g., if σ does not depend on $Ev(tr)$) provided much poorer fits to the psychometric curves.

Model of Reward Maximization
We used a model to estimate the optimal decision strategy, simulating the same number of trials as the monkeys performed while sampling from the probability distributions of $Ev_{int}(tr,L1)$, $Ev_{int}(tr,L2)$, and $Ev_{int}(tr,L2')$, as defined in Equation 8. The model used the relative difficulty of the two L2 decisions to make the L1 decision that maximized the reward probability. The probability of obtaining a reward R depends on the L1 decision as follows:

$$p(R, L1 = Left) = p_1(Ev_{int}(tr,L1))p_2(Ev_{int}(tr,L2_a))$$
$$p(R, L1 = Right) = (1 - p_1(Ev_{int}(tr,L1)))p_2(Ev_{int}(tr,L2_b)),$$
(Equation 9)

where L2$_a$ and L2$_b$ are the L2 bifurcations contingent on choosing left and right at the L1 bifurcation, respectively, and p_1 and p_2 are the psychometric functions for the L1 and L2 decisions, respectively. In each trial, the model selected the L1 decision associated with the higher reward probability in Equation 9.

Bounded Accumulation Models
To compare the flat decision model to the hierarchical model (Figure 3), we fitted both models to the psychometric functions of Figure 2D. The flat model instantiates a race between four accumulators, one for each choice, and the decision variable that reaches the bound first dictates the choice. The state of each accumulator at time τ is given by the sum of momentary evidence up to time τ:

$$dv_i(\tau) = \sum_{t=0}^{\tau} me_i(t).$$
(Equation 10)

The momentary evidence for each accumulator depends on the combinations of the luminance at the different branches as follows (Figure 3B):

$$
\begin{bmatrix} me_1(tr,t) \\ me_2(tr,t) \\ me_3(tr,t) \\ me_4(tr,t) \end{bmatrix} = \begin{bmatrix} +k_1 & +k_2 & 0 \\ +k_1 & -k_2 & 0 \\ -k_1 & 0 & +k_2 \\ -k_1 & 0 & -k_2 \end{bmatrix} \times \begin{bmatrix} I_{L1}^T(tr,t) - I_{L1}^D(tr,t) \\ I_{L2}^T(tr,t) - I_{L2}^D(tr,t) \\ I_{L2'}^T(tr,t) - I_{L2'}^D(tr,t) \end{bmatrix} + \begin{bmatrix} \mathcal{N}_1(0,\sigma) \\ \mathcal{N}_2(0,\sigma) \\ \mathcal{N}_3(0,\sigma) \\ \mathcal{N}_4(0,\sigma) \end{bmatrix},
$$
(Equation 11)

where me_{1-4} is the momentary evidence for the four choices, k_1 and k_2 are scaling parameters, and σ is the standard deviation of the Gaussian noise added independently to each decision variable per time step (Figures 3A and 3B). The equation in Figure 3A shows the state of the decision variable for one of the races, dv_1, with momentary evidence me_1 expanded according to Equation 11. The state of the accumulators is updated at 20 Hz, the rate of the external (known) noise. The process continues until one of the accumulators reaches a bound (set at an arbitrary value). We fitted k$_1$, k$_2$, and σ to the accuracy data (Figure 2D) by minimizing the root-mean-square error between the psychometric functions for data and model.

The hierarchical model differs from the flat model because the decisions for each bifurcation are made independently; i.e., the accuracy at one bifurcation is independent of the luminance at other bifurcations. The model implements a

race between two accumulators for every bifurcation. The two streams of momentary evidence that drive the two accumulators are given by

$$
\begin{aligned}
me_T(tr,t) &= k_x\big(I^T_{Lx}(tr,t) - I^D_{Lx}(tr,t)\big) + \mathrm{N}(0,\sigma_x)\\
me_D(tr,t) &= -k_x\big(I^T_{Lx}(tr,t) - I^D_{Lx}(tr,t)\big) + \mathrm{N}(0,\sigma_x),
\end{aligned}
\quad\text{(Equation 12)}
$$

where $me_{T/D}$ is the evidence for the local choice, x equals 1 or 2 for the L1 and L2 decisions, respectively, and the Gaussian noise was sampled independently for each decision variable. The state of the decision variable at time τ is given by the sum of momentary evidence up to that time. The equation in Figure 3D illustrates the state of a decision variable at bifurcation Lx at time τ. The final choice is determined at a second stage that combines the local decisions at each bifurcation. As for the flat model, we fit k_1, k_2, σ_1, and σ_2 to the psychometric curves of Figure 2D.

We used a permutation test to determine whether the strength of the evidence at L1 had a significant influence on the accuracy at L2 (Figure 3F; Figure S2). We computed the sum of differences in accuracy between the two psychometric curves in Figure 3F; i.e., $\sum_i(p^{low\ L1}_{correct}(i) - p^{high\ L1}_{correct}(i))$, where i indexes the evidence at L2 (x axis in Figure 3F). The significance of this difference was determined by a comparison of this statistic to a distribution obtained by shuffling the L1 difficulties 2,000 times. p Values correspond to the fraction of the shuffled distribution that is at least as extreme as the test statistic.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, eight figures, and one movie and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2015.08.015.

REFERENCES

Ahumada, A.J., Jr. (1996). Perceptual classification images from Vernier acuity masked by noise. Perception 25, 1831–1840.

Anderson, J.R., and Lebiere, C. (1998). The Atomic Components of Thought (Erlbaum).

Beierholm, U.R., Anen, C., Quartz, S., and Bossaerts, P. (2011). Separate encoding of model-based and model-free valuations in the human brain. Neuroimage 58, 955–962.

Bollimunta, A., and Ditterich, J. (2012). Local computation of decision-relevant net sensory evidence in parietal cortex. Cereb. Cortex 22, 903–917.

Buffalo, E.A., Fries, P., Landman, R., Liang, H., and Desimone, R. (2010). A backward progression of attentional effects in the ventral stream. Proc. Natl. Acad. Sci. USA 107, 361–365.

Chen, L., Meier, K.M., Blair, M.R., Watson, M.R., and Wood, M.J. (2013). Temporal characteristics of overt attentional behavior during category learning. Atten. Percept. Psychophys. 75, 244–256.

Churchland, A.K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A., and Shadlen, M.N. (2011). Variance as a signature of neural computations during decision making. Neuron 69, 818–831.

Cohen, M.R., and Maunsell, J.H.R. (2009). Attention improves performance primarily by reducing interneuronal correlations. Nat. Neurosci. 12, 1594–1600.

Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat. Neurosci. 8, 1704–1711.

Daw, N.D., Gershman, S.J., Seymour, B., Dayan, P., and Dolan, R.J. (2011). Model-based influences on humans' choices and striatal prediction errors. Neuron 69, 1204–1215.

Ding, L., and Gold, J.I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. Cereb. Cortex 22, 1052–1067.

Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. J. Neurosci. 32, 3612–3628.

Drugowitsch, J., DeAngelis, G.C., Klier, E.M., Angelaki, D.E., and Pouget, A. (2014). Optimal multisensory decision-making in a reaction-time task. eLife 3, 03005.

Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. Trends Cogn. Sci. 14, 172–179.

Fetsch, C.R., Kiani, R., Newsome, W.T., and Shadlen, M.N. (2014). Effects of cortical microstimulation on confidence in a perceptual decision. Neuron 83, 797–804.

Fleming, S.M., Maloney, L.T., and Daw, N.D. (2013). The irrationality of categorical perception. J. Neurosci. 33, 19060–19070.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J. Neurophysiol. 61, 331–349.

Ghose, G.M., and Maunsell, J.H. (2002). Attentional modulation in visual cortex depends on task timing. Nature 419, 616–620.

Gnadt, J.W., and Andersen, R.A. (1988). Memory related motor planning activity in posterior parietal cortex of macaque. Exp. Brain Res. 70, 216–220.

Gold, J.I., and Shadlen, M.N. (2000). Representation of a perceptual decision in developing oculomotor commands. Nature 404, 390–394.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. 30, 535–574.

Huang, Y., Hanks, T., Shadlen, M., Friesen, A.L., and Rao, R.P. (2012). How Prior Probability Influences Decision Making: A Unifying Probabilistic Model. In Advances in Neural Information Processing Systems 25, F, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, eds. (Pereira: Curran Associates, Inc.), pp. 1268–1276.

Khayat, P.S., Pooresmaeili, A., and Roelfsema, P.R. (2009). Time course of attentional modulation in the frontal eye field during curve tracing. J. Neurophysiol. 101, 1813–1822.

Kiani, R., and Shadlen, M.N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324, 759–764.

Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. J. Neurosci. 28, 3017–3029.

Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. 2, 176–185.

Krueger, L.E. (1989). Reconciling Fechner and Stevens: Toward a unified psychophysical law. Behav. Brain Sci. 12, 251–267.

Law, C.T., and Gold, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat. Neurosci. *11*, 505–513.

Lee, S.W., Shimojo, S., and O'Doherty, J.P. (2014). Neural computations underlying arbitration between model-based and model-free learning. Neuron *81*, 687–699.

Legatt, A.D., Arezzo, J., and Vaughan, H.G., Jr. (1980). Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. J. Neurosci. Methods *2*, 203–217.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. Nature *412*, 150–157.

Luck, S.J., and Vogel, E.K. (1997). The capacity of visual working memory for features and conjunctions. Nature *390*, 279–281.

Luck, S.J., Chelazzi, L., Hillyard, S.A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. J. Neurophysiol. *77*, 24–42.

Ludwig, C.J.H., Gilchrist, I.D., McSorley, E., and Baddeley, R.J. (2005). The temporal impulse response underlying saccadic decisions. J. Neurosci. *25*, 9907–9912.

Ludwig, C.J.H., Davies, J.R., and Eckstein, M.P. (2014). Foveal analysis and peripheral selection during active visual sampling. Proc. Natl. Acad. Sci. USA *111*, E291–E299.

Mehta, A.D., Ulbert, I., and Schroeder, C.E. (2000). Intermodal selective attention in monkeys. I: distribution and timing of effects across visual areas. Cereb. Cortex *10*, 343–358.

Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. Science *229*, 782–784.

Moro, S.I., Tolboom, M., Khayat, P.S., and Roelfsema, P.R. (2010). Neuronal activity in the visual cortex reveals the temporal order of cognitive operations. J. Neurosci. *30*, 16293–16303.

Motter, B.C. (1994). Neural correlates of attentive selection for color or luminance in extrastriate area V4. J. Neurosci. *14*, 2178–2189.

Mumford, D. (1992). On the computational architecture of the neocortex. II. The role of cortico-cortical loops. Biol. Cybern. *66*, 241–251.

Neri, P., and Heeger, D.J. (2002). Spatiotemporal mechanisms for detecting and identifying image features in human vision. Nat. Neurosci. *5*, 812–816.

Newell, A. (1990). Unified theories of cognition (Cambridge, MA: Harvard University Press).

Nienborg, H., and Roelfsema, P.R. (2015). Belief states as a framework to explain extra-retinal influences in visual cortex. Curr. Opin. Neurobiol. *32*, 45–52.

Peck, C.J., Jangraw, D.C., Suzuki, M., Efem, R., and Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex. J. Neurosci. *29*, 11182–11191.

Pooresmaeili, A., Poort, J., and Roelfsema, P.R. (2014). Simultaneous selection by object-based attention in visual and frontal cortex. Proc. Natl. Acad. Sci. USA *111*, 6467–6472.

Rao, R.P.N. (2010). Decision making under uncertainty: a neural model based on partially observable markov decision processes. Front. Comput. Neurosci. *4*, 146.

Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. *20*, 873–922.

Roelfsema, P.R. (2005). Elemental operations in vision. Trends Cogn. Sci. *9*, 226–233.

Roelfsema, P.R., Lamme, V.A.F., and Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. Nature *395*, 376–381.

Roelfsema, P.R., Khayat, P.S., and Spekreijse, H. (2003). Subtask sequencing in the primary visual cortex. Proc. Natl. Acad. Sci. USA *100*, 5467–5472.

Roelfsema, P.R., Tolboom, M., and Khayat, P.S. (2007). Different processing phases for features, figures, and selective attention in the primary visual cortex. Neuron *56*, 785–792.

Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. *22*, 9475–9489.

Romo, R., and Salinas, E. (2003). Flutter discrimination: neural codes, perception, memory and decision making. Nat. Rev. Neurosci. *4*, 203–218.

Russell, S.J., Norvig, P., Canny, J.F., Malik, J.M., and Edwards, D.D. (1995). Artificial Intelligence: A Modern Approach (Prentice Hall).

Salzman, C.D., Britten, K.H., and Newsome, W.T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. Nature *346*, 174–177.

Schall, J.D. (2001). Neural basis of deciding, choosing and acting. Nat. Rev. Neurosci. *2*, 33–42.

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

Smith, P.L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. Trends Neurosci. *27*, 161–168.

Smittenaar, P., FitzGerald, T.H.B., Romei, V., Wright, N.D., and Dolan, R.J. (2013). Disruption of dorsolateral prefrontal cortex decreases model-based in favor of model-free control in humans. Neuron *80*, 914–919.

Soltani, A., and Wang, X.-J. (2010). Synaptic computation underlying probabilistic inference. Nat. Neurosci. *13*, 112–119.

Supèr, H., and Roelfsema, P.R. (2005). Chronic multiunit recordings in behaving animals: advantages and limitations. Prog. Brain Res. *147*, 263–282.

Wald, A., and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. Ann. Math. Stat. *19*, 326–339.

Williams, P.E., Mechler, F., Gordon, J., Shapley, R., and Hawken, M.J. (2004). Entrainment to video displays in primary visual cortex of macaque and humans. J. Neurosci. *24*, 8278–8288.

Xing, D., Yeh, C.-I., and Shapley, R.M. (2009). Spatial spread of the local field potential and its laminar variation in visual cortex. J. Neurosci. *29*, 11540–11549.

Zylberberg, A., Dehaene, S., Roelfsema, P.R., and Sigman, M. (2011). The human Turing machine: a neural framework for mental programs. Trends Cogn. Sci. *15*, 293–300.

Zylberberg, A., Ouellette, B., Sigman, M., and Roelfsema, P.R. (2012). Decision making during the psychological refractory period. Curr. Biol. *22*, 1795–1799.