LETTERS

Changes of mind in decision-making

Arbora Resulaj^{1,2}, Roozbeh Kiani³, Daniel M. Wolpert¹ & Michael N. Shadlen³

A decision is a commitment to a proposition or plan of action based on evidence and the expected costs and benefits associated with the outcome. Progress in a variety of fields has led to a quantitative understanding of the mechanisms that evaluate evidence and reach a decision¹⁻³. Several formalisms propose that a representation of noisy evidence is evaluated against a criterion to produce a decision⁴⁻⁸. Without additional evidence, however, these formalisms fail to explain why a decision-maker would change their mind. Here we extend a model, developed to account for both the timing and the accuracy of the initial decision⁹, to explain subsequent changes of mind. Subjects made decisions about a noisy visual stimulus, which they indicated by moving a handle. Although they received no additional information after initiating their movement, their hand trajectories betrayed a change of mind in some trials. We propose that noisy evidence is accumulated over time until it reaches a criterion level, or bound, which determines the initial decision, and that the brain exploits information that is in the processing pipeline when the initial decision is made to subsequently either reverse or reaffirm the initial decision. The model explains both the frequency of changes of mind as well as their dependence on both task difficulty and whether the initial decision was accurate or erroneous. The theoretical and experimental findings advance the understanding of decision-making to the highly flexible and cognitive acts of vacillation and self-correction.

Decision-making spans a vast range of types and complexity, from choosing your partner or deciding whether to dive left or right to save a goal to simply deciding when to lift your finger. Studies of simple perceptual decisions have provided insight into the neurobiological mechanisms responsible for decision-making in both monkeys and humans (for reviews, see refs 1-3, 10). These studies often require a binary choice between two possible stimulus categories, such as leftward or rightward motion. Psychophysical and neural data¹ support models, termed drift-diffusion⁶, random walk^{5,7} and race⁸, in which a decision is made when the accumulated noisy evidence (decision variable) reaches a criterion level, termed a decision bound. Such an accumulation process explains both the accuracy of decisions over a range of difficulty levels as well as the time required to make the decisions9. These models are naturally viewed as an extension of signal detection theory and Bayesian inference to streams of data over time4,11. One important limitation of the models is that they fail to explain why a decision-maker might change their mind after an initial decision has been taken. In some instances, such changes can lead to the correction of an initial error^{12,13}. Here we develop a task in which we can monitor changes of mind. We then extend the bounded-diffusion framework to explain both the frequency and the pattern of changes of mind.

Three naive participants observed a moving random-dot stimulus and made decisions about the direction of motion (leftward or rightward), which they indicated by moving a handle to either a leftward or rightward target (Fig. 1a). Critically, the moving dots were extinguished as soon as the subjects initiated their movement (Fig. 1b) and, hence, subjects could not acquire new evidence during their movement. The choice at initiation (initial hand trajectory) and reaction times as a function of task difficulty (coherence of dot motion) were explained by a model of bounded drift–diffusion (Fig. 2, black curves) consistent with previous studies in humans and monkeys^{1,9,14}. According to this model, evidence is accumulated until it reaches one of two bounds (corresponding to leftward and rightward decisions), which determines the choice and decision time.

Although no further visual information was available after movement initiation, the hand trajectories (Fig. 1c) gave a clear indication that in some trials observers changed their minds. That is, subjects generated a curved hand path that initially was on course to reach one target, but changed direction during the movement to finish at the other target. Although some changes of mind resulted in errors, the majority corrected an initial error. Changes of mind reliably improved accuracy (Fig. 2, top row: black and red circles correspond to the initial and final choices, respectively) for all three subjects by improving sensitivity to motion (P < 0.006 for each subject).



Figure 1 | Experimental set-up. a, Schematic of the visual display (rectangle). Subjects held the handle of a robotic interface (filled circle, shown here in the 'home' position) and moved to either a leftward or a rightward circular target depending on the perceived motion direction of a central random-dot display. A mirror system prevented subjects from seeing their arm. b, The time course of events that make up a trial. Each trial started when the subject's hand was in the home position. After a random delay, the dots became visible and the subject could view the moving dot stimulus for as long as they needed (up to 2s). Subjects indicated the direction of dot motion by moving to the leftward or rightward target. As soon as the subjects moved from the home position, the motion stimulus vanished. The trial ended when the subject reached one of the two targets. c, Sample hand trajectories from one subject. Most trajectories extend directly from the home position (bottom circle) to one of the choice targets. In a fraction of trials, the trajectories change course during the movement, indicating a change of mind.

¹Computational and Biological Learning Laboratory, Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK. ²Howard Hughes Medical Institute, Janelia Farm Research Campus, 19700 Helix Drive, Ashburn, Virginia 20147, USA. ³Howard Hughes Medical Institute, National Primate Research Center and Department of Physiology and Biophysics, University of Washington, Seattle, Washington 98195, USA.



Figure 2 | **Accuracy improves through changes of mind.** Data are from three subjects (S, A and E). The top row shows that the probability of a correct decision at initiation (black) is lower than at termination (red) for almost all motion strengths. The bottom row shows that reaction times are longer for weaker motion strengths. Solid curves are fits to the data of the bounded-accumulation model (fraction of variance explained by the model fit, R^2 , for subjects S, A and E are respectively 0.96, 0.95 and 0.98 for initial decision, 0.98, 0.96 and 0.99 for final decision, and 0.92, 0.74 and 0.87 for reaction time). In this model, processing after initial commitment leads to an improvement in performance during the post-initiation phase. Error bars, s.e.m.

The observation is seemingly paradoxical. If there is information available to make a better decision, it might be expected to influence the initial decision. Every normative, 'ideal-observer'-based theory of decision-making would posit the decision as an inference made on the available evidence. The paradox is resolved if the decision-maker does not use all of the available evidence to make the initial choice but can tap into further information in the period between commitment to the initial response and termination of the movement.

Although the stimulus vanishes upon movement initiation, there is information in the processing pipeline that is potentially available to the decision-maker after movement initiation. Sensory- and motor-processing latencies ensure that not all of the information available from stimulus onset to movement initiation contributes to the decision. The sum of these latencies, termed the non-decision time (t_{nd}), was estimated to be 300–400 ms in our experiments (Supplementary Table 1 and Methods). Single-unit recordings from the lateral intraparietal area of the macaque in eye-movement versions of this task suggest that the non-decision time includes sensory and motor delays of around 220 ms and 80 ms, respectively^{15,16}. We proposed that the unused information could be processed after the brain has committed to an initial choice, thereby requiring an extension of the bounded-diffusion mechanism that includes post-initiation processing.

An analysis of the motion evidence leading to the subjects' choices supports this hypothesis. Each stimulus was a noisy sequence of random dots, which led to rapid fluctuations in the motion evidence, as quantified by 'motion energy'^{16,17} favouring left or right. For each trial, we removed the average motion energy associated with that motion strength and direction, leaving only the moment-to-moment fluctuations about the mean. We then averaged those residuals to look for evidence in the stimulus in support of the subjects' initial choice. The stimulus fluctuations immediately after stimulus onset supported the initial choice (Fig. 3a, left-hand blue curve: average over first 150 ms is positive; P < 0.0001), whereas the fluctuations in the final few hundred milliseconds had little bearing on the choice. For each subject, we identified the time point at which the average

came within 1 s.e. of zero (arrows), thus providing an empirical estimate of non-decision time. The motion-energy filtering induces a delay of 50–150 ms (Fig. 3a, inset). Taking this into account, the initial choices depend on the earliest information in the stimulus, but ignore an epoch on the order of $t_{\rm nd}$.

The pattern was different for the subset of trials in which there is a change of mind. The early information from the stimulus provided weaker support for the initial choice (Fig. 3a, left-hand red trace) and exhibited a negative trend near the time of initiation (Fig. 3a, righthand red trace), in support of the final, changed decision. The motion energy in this later epoch was significantly more negative relative to that in the remaining trials (P < 0.0001). The observation that motion energy supports both the initial and final choices provides evidence against two main alternatives to post-initiation processing: (1) change of decision based on recall and/or reconsideration of evidence acquired before initiation¹⁸, and (2) correction of an initial motor error perhaps due to confusion about the stimulus-response mapping¹². The analysis instead supports a non-decision time in which information from the stimulus arrives too late to affect an initial decision but is present to refine it after the brain has committed to a particular response and action.

We next considered how this extended processing could explain the pattern of changes of mind in the data. In particular, we wished to explain the proportion of changes to correct and to erroneous choices as a function of motion strength (Fig. 3b, red and black symbols, respectively). A seemingly optimal solution to the problem is to suppose the subject wishes to use changes of mind to maximize the percentage of correct final choices. Then the subject ought to continue to accumulate evidence about direction until there is no more to be had (that is, until time t_{nd}) and to decide in favour of the more likely direction. This formulation holds regardless of the trade-off between speed and accuracy underlying the initial choice. This idea fails to explain our findings: it predicts too many changes and it would defer them to the end of the evidence stream, which is clearly not the case (for example, early changes of mind, Fig. 1c). Because the subject must complete a hand movement, the optimal solution is likely to incorporate motor costs (energy) associated with larger corrections nearer the end of the movement. This idea can be realized by incorporating new bounds in the post-decision period to change or reaffirm an initial decision based on some criterion, thereby allowing changes to occur earlier in the movement. We considered a variety of models (Methods). The most parsimonious of these is illustrated in Fig. 3c. In this model, once the initial bound has been reached and a decision made, evidence continues to accumulate until it either reaches a new 'change-of-mind' bound or a time deadline terminates post-initiation processing. The decision rule is to change only if the accumulated evidence reaches the change-of-mind bound and to reaffirm otherwise. The offsets of the new bound and the deadline (two parameters) were fitted to account for the changes of mind as a function of coherence (Fig. 3b, curves).

For all three subjects, the model fits imply that upon termination of the initial decision, the subjects set a new bound at a level that would necessitate a reversal of the sign of the accumulated evidence. The amount of evidence required for a subject to change their mind $(B_{\Delta},$ Supplementary Table 1) differed by ~30% across subjects, which explains the variation in the pattern of their changes. In all cases, the existence of this change-of-mind bound led to a significant improvement in the fits, in comparison with using all the available information (that is, no bound and choice based on the sign of the decision variable after t_{nd} ; P < 0.003 for all subjects, likelihood-ratio test). The deadline produced by the fit suggests that subjects avail themselves of most of the information in the processing pipeline. The model captures the complex dependence of post-initiation changes on both the motion strength and the initial decision ($R^2 = 0.63-0.85$ and 0.76-0.99 for changes to correct and incorrect choices, respectively). Changes of mind were most frequent at intermediate motion



Figure 3 | A bounded-accumulation model of decision-making with postinitiation processing explains changes of mind. a, Influence of motionenergy fluctuations on initial and final decisions. Data are shown for all the trials (blue) and the subset of trials with a change of mind (red) aligned at stimulus onset (left) and movement onset (right). Motion-energy fluctuations were obtained by applying a filter to the sequence of random dots shown in each trial and subtracting the mean for all trials sharing the same motion strength and direction (Methods). The residual fluctuations are designated positive if they support the direction of the initial decision. Shading indicates s.e.m. Arrows indicate the time preceding movement initiation at which the average motion-energy fluctuations for each subject falls to within 1 s.e. of zero. Inset, impulse response for the filter used to calculate motion energy. a.u., arbitrary units. b, The model explains the probability of changes of mind from incorrect to correct choices (model, red curves; data, red symbols) and changes of mind from correct to incorrect choices (model, black curves; data, black symbols) as a function of stimulus

strengths when the initial choice was erroneous. The model offers an intuitive explanation for this. Viewed as a decision process beginning at the initial decision bound, there is a higher probability of reaffirming the initial choice, because the accumulated evidence is far from the change-of-mind bound. A change of mind therefore requires strong evidence in the short time available for post-initiation processing to move the accumulated evidence to the change-of-mind bound. Such strong evidence ought to arrive when the initial choice is an error and when the motion is strong. However, if the motion is very strong, initial errors are rare.

Our central finding is that the same data stream may be sampled at different moments to support different decisions and, hence, a change of mind. As a further test of this idea, we placed the timing of the initial decision under experimental control. This allowed us to isolate changes of mind from the strategies governing the trade-off of speed and accuracy of initial decisions in the reaction-time experiment. Instead of responding when ready, subjects were trained to time the initiation of their movement so that it coincided with an expected auditory beep. The stimulus motion began at a random time 200-2,000 ms (mean, 440 ms) before the beep and ended at the beep or at movement initiation, whichever occurred first (Methods). This experiment therefore tested whether our suggested framework generalizes to a situation in which the time of the initial choice is determined by an exogenous cue. The results of this experiment, which are summarized in Supplementary Figs 1-3, confirm the finding that subjects base their initial choice on early evidence but can avail themselves of additional evidence in the processing pipeline to revise this choice. These data also conform to a variant of the bounded-accumulation mechanism with post-initiation processing (Methods and Supplementary Figs 2 and 3).



coherence. Error bars, s.e.m. c, Information flow diagram showing visual stimulus and neural events leading to a decision and a possible change of mind. The example illustrates a rightward motion stimulus that gives rise to an initial incorrect leftward choice with reaction time around 500 ms. The visual stimulus gives rise to a decision variable (blue trace) that is the accumulation of noisy evidence. This governs the initial choice and decision time. The initial decision is complete when a 'Right' or 'Left' bound is crossed (that is, $\pm B$ of evidence has accumulated). Data from neural recordings^{15,16} suggest that the delay from motion onset to the beginning of the accumulation (t_s) is around 200 ms, and the delay from the initial decision to movement initiation (t_m) is around 80 ms. The time of the termination is around the mean decision time for the three subjects. Further accumulation takes place on the evidence still in the processing pipeline; if the accumulated evidence reaches the opposite change-of-mind bound then the decision is reversed (red), and if the deadline is reached then the decision is confirmed (green).

We expect the change-of-mind mechanism to apply under a wide variety of conditions if there is time pressure to respond. When two of our subjects were instructed to perform the reaction-time experiment more slowly, their initial decisions were more accurate and there were fewer changes of mind (data not shown). The pattern was explained by the same model with higher initiation bounds⁹. Also, because in our study the subject must complete an arm movement, the optimal solution is likely to trade off accuracy against motor costs (energy) associated with larger corrections nearer the end of the movement. Determining the optimal bounds for such a trade-off will require the coupling of concepts derived from theories of optimal feedback control¹⁹ and decision-making models. We suspect that more complex situations, for example in which movements must be timed more precisely or when a correction is more costly, might necessitate both a reaffirmation bound and bounds whose heights vary over time.

Our proposed mechanism cannot explain all changes of mind. For example, it cannot explain corrections of initial errors that arise from confusion about stimulus–response associations¹². Furthermore, a change that depends on retrieval of information from memory or incorporation of a new decision policy (for example values) would require elaboration of the model. Presumably these types of vacillations could be based on more complex processes that involve memory retrieval or application of a new criterion on a stored decision variable.

Advances in understanding the neurobiology of decision-making have benefited from simple perceptual tasks^{18,20,21}, but the same principles appear to underlie decisions related to foraging², gambling²², social selection²³ and probabilistic reasoning²⁴. The common principle is that the representation of information bearing on choice is imperfect, thus inviting the application of some criterion against which to judge the evidence. The class of bounded-diffusion models^{5–7,25,26} extends this theory of signal classification⁴ to data streams and thus incorporates time costs as well^{27,28}. An unexpected virtue of such models demonstrated by our experiment is that a part of the data stream that is not used to make the decision can none-theless support revision after a response is initiated.

This formalism provides a view of decision-making in which subjects can exploit the expectation that late-arriving information may or may not be useful to refine a decision or action. We suspect that when a change of decision is costly, energetically or otherwise, subjects will naturally tend to shun this strategy and opt for longer initial decision times. A change is precluded when an action is ballistic, for instance when a subject makes an eye movement to a choice target^{9,15}. In these instances, a change of mind can only lead to a post-decision regret²⁹ or possibly a learning signal even in the absence of overt feedback. On the other hand, a variety of complex motor sequences might benefit from early initiation premised on the expectation of additional information that is in the pipeline. It is well known that the initiation and final specifications of a movement can be dissociated in time³⁰. What we have shown here is that when these processes act on the same data stream, they can lead to a change in a decision. We speculate that a common neural mechanism explains refinement of a movement after initiation and what we experience cognitively as a change of mind about a proposition.

METHODS SUMMARY

Three naive subjects performed the main experiment. The local ethics committee approved the protocol. Subjects moved a handle in the horizontal plane. A mirror overlaid virtual images from a computer monitor onto the plane of the movement. The hand position was displayed as a small blue circle. After a random delay, a dynamic random-dot stimulus appeared (Fig. 1). In each trial, the direction of motion was randomly chosen to be leftward or rightward. Task difficulty was varied randomly by controlling the fraction of motion as quickly and as accurately as they could, and to move the handle to either a leftward or rightward target. The motion stimulus was extinguished when the movement was initiated. The trial ended when the subject reached one of the targets. Subjects performed an initial training session of at least 500 trials followed by 1,500 test trials.

We recorded the hand trajectories at 1,000 Hz. For each trial, we measured the reaction time and the final target selection. Normally hand movements for easy trials (high coherence) were straight to the target. A change of mind was reflected in a trajectory that initially travelled towards one target but ended at the other. We calculated the area between the hand path and the line from the starting position to the midpoint between the two targets. A change of mind was detected if the area swept out by the hand on the side opposite the final chosen target exceeded 0.1 cm². This criterion was based on a control experiment using 100% coherent motion. We were therefore able to determine for each trial the choice at both initiation and termination of the movement.

Full Methods and any associated references are available in the online version of the paper at www.nature.com/nature.

Received 25 April; accepted 10 July 2009. Published online 19 August 2009.

- Gold, J. I. & Shadlen, M. N. The neural basis of decision making. Annu. Rev. Neurosci. 30, 535–574 (2007).
- Sugrue, L. P., Corrado, G. S. & Newsome, W. T. Choosing the greater of two goods: neural currencies for valuation and decision making. *Nature Rev. Neurosci.* 6, 363–375 (2005).
- Schall, J. D. Neural basis of deciding, choosing and acting. Nature Rev. Neurosci. 2, 33–42 (2001).
- Green, D. M. & Swets, J. A. Signal Detection Theory and Psychophysics (Wiley, 1966).

- 5. Laming, D. R. J. Information Theory of Choice-Reaction Times (Wiley, 1968).
- Ratcliff, R. & Rouder, J. N. Modelling response times for two-choice decisions. Psychol. Sci. 9, 347–356 (1998).
- Link, S. W. The relative judgment theory of two choice response time. J. Math. Psychol. 12, 114–135 (1975).
- Smith, P. L. & Vickers, D. The accumulator model of two-choice discrimination. J. Math. Psychol. 32, 135–168 (1988).
- Palmer, J., Huk, A. C. & Shadlen, M. N. The effect of stimulus strength on the speed and accuracy of a perceptual decision. J. Vis. 5, 376–404 (2005).
- Heekeren, H. R., Marrett, S. & Ungerleider, L. G. The neural systems that mediate human perceptual decision making. *Nature Rev. Neurosci.* 9, 467–479 (2008).
- Beck, J. M. et al. Probabilistic population codes for Bayesian decision making. Neuron 60, 1142–1152 (2008).
- Rabbitt, P. & Vyas, S. Processing a display even after you make a response to it. How perceptual errors can be corrected. Q. J. Exp. Psychol. A 33, 223–239 (1981).
- Rabbitt, P. M. Error correction time without external error signals. *Nature* 212, 438 (1966).
- Smith, P. L. & Ratcliff, R. Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168 (2004).
- Roitman, J. D. & Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. 22, 9475–9489 (2002).
- Kiani, R., Hanks, T. D. & Shadlen, M. N. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* 28, 3017–3029 (2008).
- Adelson, E. H. & Bergen, J. R. Spatiotemporal energy models for the perception of motion. J. Opt. Soc. Am. A 2, 284–299 (1985).
- Romo, R., Hernandez, A., Zainos, A., Lemus, L. & Brody, C. D. Neuronal correlates of decision-making in secondary somatosensory cortex. *Nature Neurosci.* 5, 1217–1225 (2002).
- Todorov, E. Optimality principles in sensorimotor control. Nature Neurosci. 7, 907–915 (2004).
- Parker, A. J. & Newsome, W. T. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* 21, 227–277 (1998).
- Uchida, N., Kepecs, A. & Mainen, Z. F. Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nature Rev. Neurosci.* 7, 485–491 (2006).
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* 441, 876–879 (2006).
- Deaner, R. O., Khera, A. V. & Platt, M. L. Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Curr. Biol.* 15, 543–548 (2005).
- Yang, T. & Shadlen, M. N. Probabilistic reasoning by neurons. Nature 447, 1075–1080 (2007).
- Usher, M. & McClelland, J. L. The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592 (2001).
- Wong, K. F. & Wang, X. J. A recurrent network mechanism of time integration in perceptual decisions. J. Neurosci. 26, 1314–1328 (2006).
- Gold, J. I. & Shadlen, M. N. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308 (2002).
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765 (2006).
- Stuphorn, V., Taylor, T. L. & Schall, J. D. Performance monitoring by the supplementary eye field. *Nature* 408, 857–860 (2000).
- Ghez, C., Hening, W. & Favilla, M. Gradual specification of response amplitude in human tracking performance. *Brain Behav. Evol.* 33, 69–74 (1989).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements This work was supported by the Wellcome Trust, the European grant SENSOPAC IST-2005-028056, Howard Hughes Medical Institute and US National Eye Institute grant EY11378. We thank A. Faisal, H. Vincent, I. Howard and J. Ingram for their assistance. M.N.S. thanks Trinity College, Cambridge, for support.

Author Contributions D.M.W. and M.N.S. planned the experiments. A.R. performed the experiments. All authors analysed and interpreted results, and all authors wrote the paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. Correspondence and requests for materials should be addressed to M.N.S. (shadlen@uw.edu).

METHODS

Behavioural task. Four naive subjects (three male and one female) provided informed consent and participated in the experiment. The local ethics committee approved the protocol. Three subjects performed each of the reaction-time and cued-movement experiments (two subjects, S and E, performed both, with the reaction-time experiment first). Subjects were seated and used their preferred hand to hold the handle of a vBOT manipulandum³¹ that was free to move in the horizontal plane (Fig. 1a). Subjects were prevented from seeing their arm by a mirror that was used to overlay virtual images of a video display (updated at 75 Hz) onto the plane of the movement. A chin- and headrest ensured a viewing distance of 40 cm. The hand position was displayed as a small blue circle (radius, 0.5 cm).

The time course of a trial in the reaction-time experiment is shown in Fig. 1b. A trial began when the subject's hand was in the home position (circle of radius 1 cm; Fig. 1a). After a random delay, sampled from a truncated exponential distribution (range, 0.7-1.0 s; mean, 0.82 s), a dynamic random-dot stimulus appeared at the centre of the screen within a circular aperture subtending 5° of visual angle. The motion stimulus is described in detail in previous studies¹⁵. In each trial, the direction of motion was randomly chosen to be leftward or rightward. The stimulus density was 15.6 dots $deg^{-2} s^{-1}$. Dots were displayed for one video frame and then either replaced at a random position or displaced to the left or right three video frames (40 ms) later. This displacement would produce a speed of 7.1° s⁻¹. Thus the positions of the dots in frame four, say, were correlated only with the displaced dots in frames one and/or seven but with none of the dots in frames two, three, five and six. The probability that each dot would be displaced as opposed to randomly replaced, termed the per cent coherence, determined the task difficulty and was selected randomly from the set (0%, 3.2%, 6.4%, 12.8%, 25.6%, 51.2%).

The subjects were instructed to judge the direction of the moving random dots as quickly and as accurately as they could, and to reach to a corresponding circular target (one on the left and one on the right; radius, 1.5 cm; 20 cm from the starting position and 28° from the midline; Fig. 1a). Critically, when the movement was initiated—that is, the hand crossed the boundary of the home-position circle—the random-dot stimulus was extinguished. Subjects were required to reach the target with a movement duration of 500 ± 200 ms. The trial ended when the subject reached one of the targets. Subjects were provided with visual feedback of whether they had made the correct choice (for the 0% coherence trials, half of the trials were randomly designated 'correct'). Subjects were of the dot aperture—the targets were large enough that they could be easily reached using peripheral vision. Subjects performed an initial training session of at least 500 trials followed by 1,500 test trials.

In the cued-movement task, subjects heard five beeps equally spaced in time (500-ms spacing) and were required to initiate movement on the fourth beep and reach the target on the fifth beep (Supplementary Fig. 1a). Random-dot motion began at a random interval before the fourth beep (truncated exponential distribution: range, 0.2-2 s; mean, 0.44 s). The motion display was extinguished on the fourth beep or at the time of movement initiation if the subject slightly anticipated the beep. Feedback was provided to maintain movement initiation and termination within $\pm 100 \,\mathrm{ms}$ of the fourth and fifth beeps, respectively. Again, subjects were given feedback of whether they had made the correct choice. Subjects performed an initial training session of 500 trials followed by 2,000 test trials. Data analysis. We recorded the hand trajectories at 1,000 Hz. For each trial, we quantified the reaction time (time to movement initiation from start of motion stimulus) and the final target selection. In addition, we developed a measure, based on the hand trajectories, of whether subjects had changed their decision during the movement. Normally hand movements for easy trials (high coherence) were straight to the target (Fig. 1c). A change of mind was reflected in a trajectory that initially travelled towards one target but ended at the other. We calculated the area between the hand path and the line from the starting position to the bisector of the two targets. A change of mind was deemed to have occurred if the area swept out by the hand on the side opposite the final chosen target exceeded 0.1 cm² and the point of maximum horizontal deviation was outside the home position. This criterion was chosen on the basis of a control experiment with two of our subjects using the reaction-time condition but with 100%-coherent motion stimuli. We expected to see few, if any, changes of mind under this condition and in fact observed two change-of-mind trials out of 400, both of which were obvious lapses with swept areas at least three times larger than the criterion, suggesting that our method of determining changes of mind is conservative. We were therefore able to determine for each trial the choice at both initiation and termination of the movement.

Modelling. For the reaction-time experiment (Figs 2 and 3), we adapted a bounded-accumulation model (Fig. 3c) to explain the initial- and final-choice

frequencies (Fig. 3b). We first explain the model for the initial choices and then expand it to explain changes of mind.

For the initial choices, the model posits that evidence accumulates from a starting point, y_0 , until it reaches an upper or lower bound $(\pm B)$, which determines the initial choice and decision time. The increments of evidence are idealized as normally distributed random variables with unit variance per second and mean $\mu = kC + \mu_0$, where *C* is signed motion strength (a positive value corresponding to rightward motion and negative value corresponding to leftward motion); *k*, *B*, y_0 and μ_0 are free parameters. The parameters *B* and *k* explain the trade-off between the speed and the accuracy of the initial choices; μ_0 and y_0 are respectively drift and starting-point offsets, which explain bias for one of the choices. The bias terms were not necessary for all subjects (Supplementary Table 1).

This formulation leads to the following simplification³², which may help to provide an intuition for the effect of motion strength on initial choice and reaction time. If $y_0 = 0$, the probability of a rightward initial choice is

$$P_{\text{right}} = \left[1 + \exp\left(-2\mu B\right)\right]^{-1}$$

and the mean decision time is

$$t_{\rm d} = \frac{B}{\mu} \tanh\left(\mu B\right)$$

The reaction time incorporates additional latencies from stimulus onset to the beginning of the bounded-accumulation process and from the termination of the process to the beginning of the motor response. The sum of these latencies, the non-decision time t_{nd} , is an additional parameter of the model such that the measured reaction time is $t_d + t_{nd}$, which we set for each direction choice.

Because the stimulus duration in each trial equals the reaction time, there is additional evidence from the stimulus that is potentially available for processing after the brain has committed to an initial choice. The model incorporates this additional information as follows. When the initial decision ends, the accumulation continues (from $\pm B$) until either a second, post-initiation change-of-mind bound is crossed, in which case the decision is reversed, or a temporal deadline is exceeded, in which case the initial decision is reaffirmed (Fig. 3c). The height of this new bound was offset by B_{Δ} from the initiation bound. A value of $B_{\Delta} = B$ would imply that a change of mind occurs when the evidence changes sign, and a value of $B_{\Delta} = 2B$ would imply that a change requires an amount of net evidence represented by the initial bounds. The values for our subjects were between *B* and 2*B*.

The fits to the initial choices and reaction times provide the sensitivity parameter (k), initial bounds (B) and non-decision times (t_{nd}) used in the postinitiation analyses. We then considered a series of plausible models for the post-initiation phase. These models were intended to explain the observed initial and final choices (bivariate observations: left-left, left-right and so on) given fixed values for k, B and t_{nd} . The strategy ensures that all comparison models are on equal footing and that the number of parameters for post-initiation is small. We compared an 'optimal' model using all available evidence (no additional degrees of freedom (d.f.)), a single flat change-of-mind bound (d.f. = 1), a flat change-of-mind bound with a deadline (as described above; d.f. = 2), flat bounds for change of mind and for reaffirmation (d.f. = 2), and variants of these models with quadratic collapsing bounds (an extra 1-2 d.f. to parameterize the collapse). We used a likelihood-ratio test for nested models and supported these comparisons using the Bayes information criterion³³. On the basis of these comparisons, we adopted the simplest model that accounted for all the subjects' data (Fig. 3c): one with a single change-of-mind bound and a cut-off that would censor late information acquired during t_{nd} . The parameters for this model are shown in the final two rows of Supplementary Table 1. All fits were performed using maximum-likelihood methods. Model choice probabilities and reactiontime distributions were derived from numerical solutions of Fokker-Planck equations for the bounded-diffusion process³⁴.

Although it appears that a large number of parameters were used to model the initial and final choices, the strategy is conservative and intuitive. We used six parameters for the fits to the initial choices and reaction times to ensure that the estimates of parameters that affect the post-initiation phase (k, B and t_{nd}) were as accurate as possible. A model with just three parameters gives acceptable fits for the initial choices and reaction times for all three subjects, but the additional parameters explain the small biases in two of the subjects and the 4–10-ms difference in t_{nd} for leftward and rightward choices. Although several of these terms have negligible effects for one or more subjects (Supplementary Table 1), they produce more accurate estimates of k, B and t_{nd} . As noted above, the simple two-parameter model used to fit the post-initiation data was supported by an extensive model comparison. To perform this model comparison with as much power and sensitivity as possible, it was necessary to place all models on equal footing by supplying the best possible values for the inherited parameters (k, B

and $t_{\rm nd}$). In particular, we did not want to justify a more complicated model (for example one with collapsing bounds) simply because the additional degrees of freedom could explain residual error in *k* and *B*. Our strategy is conservative in that it tends to reduce the explanatory power of more complex models for changes of mind.

We also performed a cross-validation analysis to ensure that the large number of parameters in our fit to the reaction-time data did not lead to overfitting. We split each subject's data set into two equal halves (random permutation of trials at each motion strength) and fitted each separately. We used the fits from one half to predict the other half of the data. The cross-validation fits, goodness of fit and parameter estimates are shown in Supplementary Fig. 4 and Supplementary Tables 3 and 4. The similarity of the predictions and fits provides reassurance that the model is not overparameterized.

A simpler version of the model was used to fit data from the cued-movement experiment (Supplementary Figs 2 and 3). Here the non-decision time, t_{ndy} delimits the portion of the data stream available for the initial choice. In a trial in which the stimulus is displayed for a time t_{stim} , subjects can use $t_d = t_{stim} - t_{nd}$ of the data stream (or no information if the stimulus duration is shorter than the nondecision time) to determine their initial choice, and a further t_{nd} (or the stimulus duration if shorter than t_{nd}) to potentially revise their decision. Put simply, the initial choice is governed by the sign of the decision variable after t_d of diffusion, whether or not it has terminated. Post-initiation processing occurs on the remaining data stream until either the left or right choice bound is reached. The same symmetric bounds were used before and after initiation. A key difference from the reaction-time experiment is that once the accumulated evidence has reached a bound, the diffusion process terminates and there is no opportunity for a change of mind. Thus, only non-terminated decisions after time t_d are eligible for a change of mind. This seems sensible because, unlike in the reaction-time experiment, the subject does not choose the time of initiation. Termination of the process is tantamount to accepting that the level of evidence is sufficient for a choice. Model fits (for k, B and t_{nd} , see Supplementary Table 2) were obtained using maximum-likelihood methods. Because initiation was timed to coincide with an external beep in this experiment, the main effect of the bounds was to curtail the improvement in accuracy that would be expected for perfect integration for long times t_d (ref. 16). The initial- and final-choice probabilities were derived by numerical solution of Fokker-Planck equations for each trial, using the same stimulus durations as in the data set.

Statistical analysis. Unless otherwise stated, *P*values are based on *t* statistics constructed from parameter estimates and their associated standard errors. We calculated the standard errors by using the inverse Hessian from maximum-likelihood fits wherever possible, or a bootstrap procedure³⁵ when the numerical solution of the Fokker–Planck equation did not support accurate calculation of the Hessian. For the fraction of t_{nd} , we report the 95% confidence interval (method of fiducial limits³⁶, likelihood-ratio test) because this parameter is bounded by zero and one. The R^2 values accompanying the model fits were calculated as one minus the fraction of unexplained variance for the data points displayed in the graphs. To evaluate the differences between initial- and final-choice probabilities, we did not rely on the model in Fig. 3 but instead performed logistic regression. Accordingly, the probability of choosing right is given by

$P_{\text{right}} = [1 + \exp(-b_0 - b_1 C - b_2 I - b_3 IC)]^{-1}$

where *C* is the signed motion strength, *I* is an indicator variable (zero for initial choice and one for the final choice) and b_i are fitted coefficients. To test for improved sensitivity (accuracy) with changes of mind, we evaluated the null hypothesis { $H_0: b_3 \leq 0$ }. An alterative formulation—probability correct as a function of unsigned motion strength—confirmed the statistical significance of this analysis as well as the analysis of the cued-motion experiment.

For the motion-energy analyses, we extracted a time series from the sequence of random dots shown in each trial by applying a filter for rightward and leftward motion with passband centred at $1.0 \text{ cyc} \text{ deg}^{-1}$ and 7.1 Hz, thus matching the speed and dot displacement in our stimulus (for details, see refs 16, 17). The difference in these time series represents momentary evidence in favour of one or the other choice. To combine data across trials, we removed the average motion energy associated with each trial's motion strength and direction. We then applied a sign convention so that positive fluctuations are in the direction of the subject's initial choice. The graphs in Fig. 3a and Supplementary Fig. 2b show these averaged residuals, time-locked to either stimulus onset or movement initiation.

For the statistical analysis of the motion energy time-locked to movement initiation, we used the data from all trials (blue curves) to identify the point in time (for each subject) at which stimulus motion fluctuations no longer influence the initial choice, using an arbitrary value of 1 s.e. from zero. This procedure gives a model-free estimate of t_{nd} . We analysed the motion energy from the change-of-mind trials from this time until movement initiation. To test whether the total motion energy in an epoch differed significantly from zero, we applied a permutation test (randomization of the sign of motion energy in each trial)³⁵. To compare the motion energy in change-of-mind and reaffirmation trials, we applied a bootstrap procedure. We calculated the total motion energy in the change-of-mind trials using the epoch defined above and compared this with the distribution of values obtained in randomly resampled trials without change of mind over the identical epochs. This bootstrap comparison compensated for a lack of power due to there being relatively few change-of-mind trials (for example, neither of the trends in the left-hand red curves of Fig. 3a and Supplementary Fig. 2b are significantly different from zero).

- Howard, I. S., Ingram, J. N. & Wolpert, D. M. A modular planar robotic manipulandum with end-point torque control. J. Neurosci. Methods 181, 199–211 (2009).
- Shadlen, M., Hanks, T., Churchland, A., Kiani, R. & Yang, T. in Bayesian Brain: Probabilistic Approaches to Neural Coding (ed. Doya, K. et al.) 209–237 (MIT Press, 2006).
- Kass, R. E. & Wasserman, L. A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. J. Am. Stat. Assoc. 90, 928–934 (1995).
- 34. Risken, H. The Fokker–Planck Equation: Methods of Solution and Applications 2nd edn (Springer, 1989).
- Efron, B. The Jackknife, the Bootstrap and Other Resampling Plans (Society for Industrial and Applied Mathematics, 1982).
- 36. Wang, Y. H. Fiducial intervals: what are they? Am. Stat. 54, 105-111 (2000).

SUPPLEMENTARY INFORMATION



Supplementary Figure 1. Cued movement task. **a**, Behavioral task. As in the main experiment, subjects identified the direction of motion of dynamic random dot stimuli. Instead of responding when ready, subjects were trained to time the initiation of their movement so that it coincided with the 4th of a series of 5 beeps and to bring the cursor to the choice target on the 5th beep. The stimulus motion began at a random time 200-2000 ms (mean 440 ms) before the 4th beep and ended at the beep or at movement initiation, whichever occurred first (see Methods for additional details). **b**, Distribution of initial hand velocity in the cued condition for the 3 subjects. Histograms of lateral (horizontal) hand velocity when it had moved 2 cm outward from the home position. Hartigan's test of unimodality³⁷ shows that the distributions are bimodal (p<0.001 for all subjects). Note that the classification of initial choices was not based on these measurements but instead on analysis of the hand trajectories (see Methods). Subjects clearly indicated a choice at initiation by moving in the direction of one of the targets. On many trials, however, subjects reversed this initial choice during the movement (see Supplementary Figure 2).



Supplementary Figure 2. Accuracy improves through "changes of mind" in the cued movement task. **a**, Data from three subjects. The top row shows the probability of a correct decision at initiation (black) and at termination (red) for the trials with movement initiation times (with respect to stimulus onset) shorter than 360ms (1st tercile). The bottom row shows trials with later movement initiation times (360-510 ms; 2nd tercile). Smooth curves are model fits; error bars are s.e.m. **b**, Influence of motion energy fluctuations on initial and final decisions. The data are shown for all the trials (blue) and the subset of trials with a change of mind (red) aligned at stimulus onset (left) and movement onset (right). Same conventions as in Figure 3a (main text). Arrows indicate the time from movement onset that the average motion energy fluctuations for each subject falls to within 1 s.e. of zero.

Notes. For the shorter durations, the changes of mind led to an improvement in decision accuracy (p < 0.003 for all 3 subjects). In contrast to the previous model, the decision variable might fail to reach a termination bound before initiation, especially for the shorter duration stimuli, whereas it may reach a bound some time prior to initiation, especially for the longer duration stimuli. A small adjustment to the model accommodates these two situations (see Methods). For all subjects, the model explains the degree of improvement in accuracy associated with changes of mind (panel a, solid lines, $R^2 = 0.92$ to 0.95 across subjects) and the degree of improvement in both the initial and final choices as a function of viewing duration. Moreover, it accounts for the proportion of changes to correct and to erroneous choices as a function of motion strength and stimulus duration (Supplementary Figure 3). The motion energy analysis suggests that information acquired at the beginning of motion viewing affects the initial choice (panel b, left blue curve; over first 150 ms p<0.0001), whereas fluctuations near the end of motion viewing are unrelated (right blue curve). In contrast, these late fluctuations do influence the decision changes that occur after initiation. Near the time of initiation, the motion energy on change of mind trials was significantly less than the motion energy on the remaining trials (p<0.0001; bootstrap, see Methods). Compared to the reaction time experiment, the motion fluctuations appear to influence the initial choice later into the trial (panel b, arrows), consistent with approximately 100 ms shorter t_{nd} in the model fits (Supp. Table 2). We suspect that this may be a consequence of the external cue to initiation, resembling the conditions that are known to produce 'express' saccadic latencies³⁸.



Supplementary Figure 3. Probability of change of mind from incorrect to correct (left column) and from correct to incorrect (right column) in the cued movement task varies with motion strength and with the time of movement initiation. Data were divided into 3 quantiles based on the movement initiation time relative to motion onset (red, the shortest quantile; blue, the longest quantile). The curves are the expected values based on the model fits shown in Supplementary Fig. 2a. The model tends to underestimate the frequency of changes of mind because it does not incorporate initial biases or initial guesses. Despite this limitation, the fit captures qualitative differences between subjects and between correct and error vacillations.



Supplementary Figure 4. Cross validation for the reaction time experiment. The graphs correspond to those shown in Figures 2 and 3b of the paper. The symbols are means using a random half of the data for each subject. Solid curves are model predictions, based on the fit to the other half of each subject's data. Dashed and dot-dash curves show the fits to the entire data set and to the half of the data shown here. Error bars are s.e.m. The similarity of the predictions and fits provides reassurance that the model is not over parameterized (see also Supplementary Table 3).

	Subject S	Subject E	Subject A
k	0.30±0.025	0.27±0.03	0.25±0.03
В	13.2±0.6	12.4±0.5	13.0±1.3
<i>t_{nd} right</i> (ms)	322±12	368±11	390±28
t_{nd} left (ms)	326±13	378±12	395±29
μ_0	0.006 ± 0.002	0.018±0.003	0.013±0.003
y ₀	0.0004±0.0002	0.0003±0.00005	-0.0005±0.0002
B_{Δ}	23.3±0.6	18.4±1.0	25.5±0.7
PIP duration as fraction of t_{nd}	1.0 (CI: >0.8)	0.68 (CI: 0.54 -0.7)	1.0 (CI: >0.95)

Supplementary Table 1. Fitted parameters of the accumulation-to-bound model with post-initiation processing in the reaction time experiment. The six parameters for initiation and two parameters for post-initiation processing are shown with standard errors for each subject (see Methods). The interval for post-initiation processing (PIP) is bounded by t_{nd} ; parentheses show the 95% confidence intervals (CI).

	Subject S	Subject E	Subject P
k	0.30±0.03	0.26±0.03	0.19±0.05
В	16.1±2.8	15.1±7.8	22.9±7.5
t_{nd} (ms)	260±18	273±12	273±12

Supplementary Table 2. Fitted parameters of the accumulation-to-bound model with post-initiation processing in the cued experiment. The three parameters are shown with standard errors for each subject.

	$S \over R_{all}^2, R_{half}^2, R_{pred}^2$	$E \\ R^2_{all}, R^2_{half}, R^2_{pred}$	$A onumber R_{all}^2, \mathcal{R}_{half}^2, \mathcal{R}_{pred}^2$
Prob Correct init (Fig 2, top black)	0.956, 0.938, 0.939	0.98, 0.98, 0.96	0.954, 0.976, 0.976
Prob Correct final (Fig 2, top red)	0.976, 0.962, 0.960	0.99, 0.97, 0.96	0.964, 0.938, 0.937
Initiation time (Fig 2, lower)	0.92, 0.69, 0.54	0.87, 0.89, 0.87	0.74, 0.61, 0.56
Prob Ch. to correct (Fig 3b, red)	0.63, 0.67, 0.47	0.85, 0.51, 0.43	0.72, 0.64, 0.62
Prob Ch. to error (Fig 3b, black)	0.76, 0.84, 0.82	0.99, 0.90, 0.89	0.88, 0.769, 0.767
mean	0.85, 0.83, 0.75	0.94, 0.90, 0.82	0.85, 0.79, 0.77
combined	0.91, 0.69, 0.55	0.87, 0.89, 0.88	0.74, 0.61, 0.56

Supplementary Table 3. Fraction of variance for the cross validation in the reaction time experiment. The R^2 terms describe the fraction of variance in data that is explained by the smooth curves (model) in Figures 2, 3b and Supplementary Figure 4. For each of the five predicted functions and for each subject, there are three R^2 terms. The R_{all}^2 are computed from the data and model fits shown in the main text figures, using all trials. The R_{half}^2 are computed from a random half of the data and model fit to these data (Supplementary Figure 4, dot-dash curve). The R_{pred}^2 are computed from the same random half of the data and the model fit to the other half of the data (Supplementary Figure 4, solid line). This last term furnishes the cross validation: a measure of how well the model derived from one set of observations *predicts* a different set of observations. The bottom two rows show (*i*) the means of the R^2 values from the five graphs and (*ii*) the R^2 calculated from the total sum of squares and total residual sum of squares from the lowered number of trials, the fits to the random half of the data are worse than the fits to the entire data set. Critically, the cross validation (prediction) is only slightly worse than the fit to the same data, confirming that the model does not over fit the data. Note that for each subject, one model accounts for all five comparisons.

	Subject S orig, 1^{st} , 2^{nd} (% Δ)	Subject E orig, 1^{st} , 2^{nd} (% Δ)	Subject A orig, 1^{st} , 2^{nd} (% Δ)
k	0.30, 0.30, 0.30 (0.6%)	0.27, 0.26, 0.32 (19.3%)	0.25, 0.26, 0.25 (1.5%)
В	13.2, 12.8, 13.6 (5.3%)	12.4, 12.5, 11.4 (8.8%)	13.0, 13.1, 12.8 (2.6%)
t_{nd} right (ms)	322, 325, 319 (1.9%)	368, 363, 391 (7.3%)	390, 393, 388 (1.3%)
t_{nd} left (ms)	326, 332, 320 (3.3%)	378, 375, 402 (6.5%)	395, 400, 393 (1.7%)
μ_{0}	0.006,0.006,0.006 (0.21%)	0.018, 0.019,0.020 (5%)	0.013, 0.014,0.013 (6.5%)
y 0	0.0004, 0.0004,0.0004 (2.2%)	0.0003,0.0004,0.0003(9.5%)	-0.0005,-0.0005,-0.0005 (5.5%)
B_{Δ}	23.3, 24, 22.3 (7.2%)	18.4, 18.3, 16.2 (11.7%)	25.5, 25.7, 25.0 (2.7%)
PIP duration as fraction of t_{nd}	1, 1, 1 (0%)	0.68, 0.69, 0.50 (27.7%)	1, 1, 1 (0%)

Supplementary Table 4. Model parameter estimates from cross validation for the reaction time experiment. Each column contains the original parameter estimate based on all data (same as Supplementary Table 1), and parameter estimates from fits to random halves of the data. The term in parentheses is the percentage difference between the latter two estimates: the absolute value of the difference, divided by the larger of the two estimates.

Supplementary Notes

- ³⁷ Hartigan, J.A. & Hartigan, P.M., The Dip Test of Unimodality. *Annals of Statistics* 13 (1), 70-84 (1985).
- ³⁸ Fischer, B. & Boch, R., Saccadic eye movements after extremely short reaction times in the monkey. *Brain. Res.* 260, 21-26 (1983).