Neural correlates, computation and behavioural impact of decision confidence

Kepecs A., Uchida N., Zariwala H. and Mainen Z.F.

Confidence estimates in integrator models of decision-making

Computing decision confidence requires that in addition to a binary choice, the decision process yields a graded value measuring the reliability and consistency of the internal variables contributing to a decision. This can be achieved for other classes of models including models based on the integration of evidence that are able to also account for other features of behaviour, such as reaction times[1-3].

To demonstrate the generality of the model predictions (Fig. 4c,d) we simulated a version of the integrator model, the "race model". In race models, separate decision variables accumulate evidence for different options and the decision taken is determined by which decision variable reaches threshold first (Supplementary Fig. 6a). To Nsimulate this, in each trial the stimulus is a normally distributed random variable $s(t) \in N(\mu_{stim}, \sigma_{stim})$, where the sign of μ_{stim} sets the direction of correct choice and the signal-to-noise ratio μ_{stim}/σ_{stim} sets the difficulty of discrimination.

In the simplest version of the race model there are two independent decision variables that accumulate evidence for and against the hypothesis that $\mu_{stim} > 0$. Each decision variable, $e(t)$, accumulates evidence for one direction:

$$e^{+/-}(t) = \int_0^t s^{+/-}(t)dt \text{ where } s^{+}(t) = \begin{cases} s(t), & s(t) > 0 \\ 0, & s(t) \le 0 \end{cases} \text{ and } s^{-}(t) = \begin{cases} 0, & s(t) \ge 0 \\ -s(t) & s(t) < 0 \end{cases}$$

When one of the decision variables, $e^{+}(t)$ *or* $e^{-}(t)$, reaches a predetermined threshold, θ, the race is terminated and a decision is generated in favour of the decision variable crossing threshold first. Therefore at decision time, t_θ, $e^{+}(t_\theta) = \theta$ or $e^{-}(t_\theta) = \theta$. We simulated this model with the following parameters: $\mu_{stim} \in U(-0.2,0.2)$, $\sigma_{stim} = 1$, $\theta = 100$, *and* $dt = 1$. Supplementary Fig. 6c shows the fraction of choices in favour of the "+ hypothesis", $\mu_{stim} > 0$, as a function of stimulus, μ_{stim}. This psychometric curve is qualitatively similar to that of rats (compare to Fig. 1c).

An estimate of decision confidence can be computed in a race model by measuring the distance between the two decision variables at the time the race is terminated. This was originally proposed by Vickers[4], who termed it the "balance of evidence". To see that the distance between decision variables can provide a reasonable estimate of confidence, we plot choice accuracy as a function of this distance, $\Delta e = \left| e^{+}(t_\theta) - e^{-}(t_\theta) \right|$ (Supplementary Fig. 6d, dashed line). This distance, Δe, can be normalized, $\Delta e/\theta$, to yield the "balance of evidence" measure[4]. Here instead we sought to compute an estimate, δ, that is closer to the veridical confidence and reflects the actual outcome probability. For a perfectly calibrated or veridical confidence estimate δ would correspond to the probability of correct outcome, from chance level ($\delta = 0$) to perfect ($\delta = 1$) performance. The consideration of the theoretically appropriate calibration method is beyond our scope. For a given signal to noise ratio the correct calibration function may be derived by considering the error rate as a function of the decision threshold[1-3]. Here we used an approximation, $\delta = f(\Delta e) = 2/(1+e^{c(\Delta e/\theta)})$, with $c = 1/3$, which provides excellent performance across multiple stimuli (see Supplementary Fig. 6b). The role of calibration is illustrated in

Supplementary Fig 6d, showing that accuracy is nearly a linear function of δ (solid line) in contrast to the normalized distance, $\Delta e/\theta$ (dashed line). Finally, when this confidence estimate, δ, is plotted as function of stimulus and outcome it shows the same pattern as the model in main text (cf. Fig. 4d and Supplementary Fig. 6e). Note that for a given stimulus range (signal-to-noise ratio, μ_{stim}/σ_{stim}), there is only a single free parameter (threshold, θ) that determines the slope of the psychometric function. This qualitative pattern—the opposing V-shaped curves for correct and error choices—was robust to different choices of parameters. This example demonstrates that confidence estimates can be readily computed in other classes of decision models, and yield qualitatively similar predictions.

The confidence estimate examined is not unique, however, and may be generalized to other measures of decision uncertainty. For instance, the variance or entropy of the decision variables could provide a basis for confidence estimation. For the race model, the variance of the decision variables is simply the square of the estimate introduced above: $\delta^V=(e^+(t_\theta) - e^-(t_\theta))^2$. The entropy of decision variables p_i, is $\delta^S = -\sum_i p_i \log(p_i)$. For the race model this can be expressed as $\delta^S = -p^+ \log(p^+) - p^- \log(p^-)$ with $p^{+/-} = e^{+/-}(t_\theta)/(e^+(t_\theta) + e^-(t_\theta))$. The limiting cases of maximum and minimum uncertainty are instructive and easy to calculate. Uncertainty is highest when the two decision variables reach threshold at the same time, which would make $\delta^S = 1$ bit, while uncertainty is lowest when all the evidence accrues in favour of one decision variable, which would make $\delta^S = 0$ bit. Such measures are likely to be useful when considering confidence estimates based on a population of neurons[5]. Note, however, that δ^S needs to be calibrated in order to provide veridical outcome predictions. Therefore measures like δ^S could serve as sufficient statistics for computing confidence that need to be calibrated to yield an instantaneous estimate of outcome likelihood. The appropriate calibration function in turn can be found using reinforcement learning.

Intuition for the confidence model

To provide intuition into the unexpected pattern of uncertainty as a function of stimulus and outcome (Fig. 4c) we need to examine how different stimulus and memory sample configurations lead to a choice and an estimate of confidence (Fig. 4a). Note that the model (or a subject) has access only to a stimulus sample and not the stimulus type (e.g. 56/44) as plotted. First, consider that error choices occur when on a given trial the stimulus, s_i, and memory, b_i samples are reversed compared to the mean of their respective distributions. This can only occur within a region where the two distributions overlap, whereas correct ordering may occur over the entire range of stimulus values. Therefore the size of the overlap region will place a limit on the maximal confidence values that can be attained for error choices. Since the region of overlap is smaller than the entire range, the maximal distance between samples will be smaller for errors than correct choices, resulting in lower confidence estimates (higher uncertainty) in the choice on average. Moreover, the further away the stimulus is from the boundary, the smaller the possible region of overlap between their distributions, and therefore the smaller the maximum possible distance between the two samples in an error trial. Consequently, for easy stimuli errors are rare because the overlap is small, and in those few incorrect trials the samples cannot possibly be far from each other resulting in low confidence estimates (high uncertainty).

Decision confidence or risk?

Decision confidence is a form of subjective uncertainty, but it is important to distinguish it from other forms of uncertainty. The term "uncertainty" is often used synonymously with reward or outcome "risk"[6-10]. "Decision uncertainty" and "outcome uncertainty" are similar in that both are based on calculating the variance or entropy across a set of variables. The critical distinction is the set of variables over which the measure is calculated. Decision uncertainty is measured across variables observed in a single trial. Outcome uncertainty, in contrast, is measured across outcomes observed over multiple trials.

In free choice tasks it is possible to independently manipulate outcome probability and outcome variance. In contrast, in two-alternative psychophysical decision tasks outcome probability ranges from 0.5 to 1 and therefore covaries with outcome variance. Hence our data are consistent with either a representation of outcome probability or outcome uncertainty signals. Either way, the computation of such signal must incorporate an estimate of decision uncertainty.

References

1. 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, 700-65 (2006)/
2. Mazurek, M. E., Roitman, J. D., Ditterich, J. & Shadlen, M. N. A role for neural integrators in perceptual decision making. Cereb Cortex 13, 1257-69 (2003).
3. Ratcliff, R. & Smith, P. L. A comparison of sequential sampling models for two-choice reaction time. Psychol Rev 111, 333-67 (2004).
4. Vickers, D. Evidence for an accumulator model of psychophysical discrimination. Ergonomics 13, 37-58 (1970).
5. Zemel, R. S., Dayan, P. & Pouget, A. Probabilistic interpretation of population codes. Neural Comput 10, 403-30 (1998).
6. Critchley, H. D., Mathias, C. J. & Dolan, R. J. Neural activity in the human brain relating to uncertainty and arousal during anticipation. Neuron 29, 537-45 (2001).
7. Hsu, M., Bhatt, M., Adolphs, R., Tranel, D. & Camerer, C. F. Neural systems responding to degrees of uncertainty in human decision-making. Science 310, 1680-3 (2005).
8. Fiorillo, C. D., Tobler, P. N. & Schultz, W. Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299, 1898-902 (2003).
9. McCoy, A. N. & Platt, M. L. Risk-sensitive neurons in macaque posterior cingulate cortex. Nat Neurosci 8, 1220-7 (2005).
10. Tobler, P. N., O'Doherty, J. P., Dolan, R. J. & Schultz, W. Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. Journal of neurophysiology 97, 1621-32 (2007).

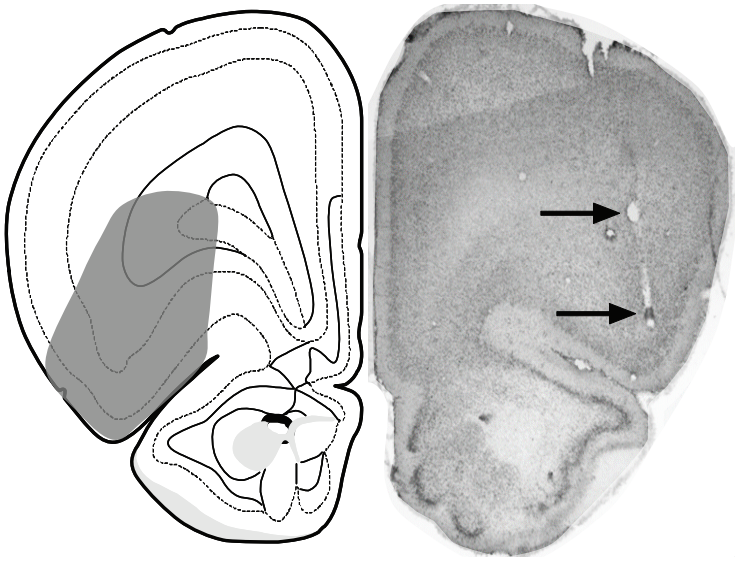

Supplementary Figure 1. Anatomical location of recording sites

Nissl-stained coronal section of rat frontal cortex showing the lesion sites from one tetrode. Left shows the estimated area of the recordings from 3 rats overlaid on a section 3.60 mm anterior to bregma. Recording locations ranged from +3.0 mm to +4.2 mm.

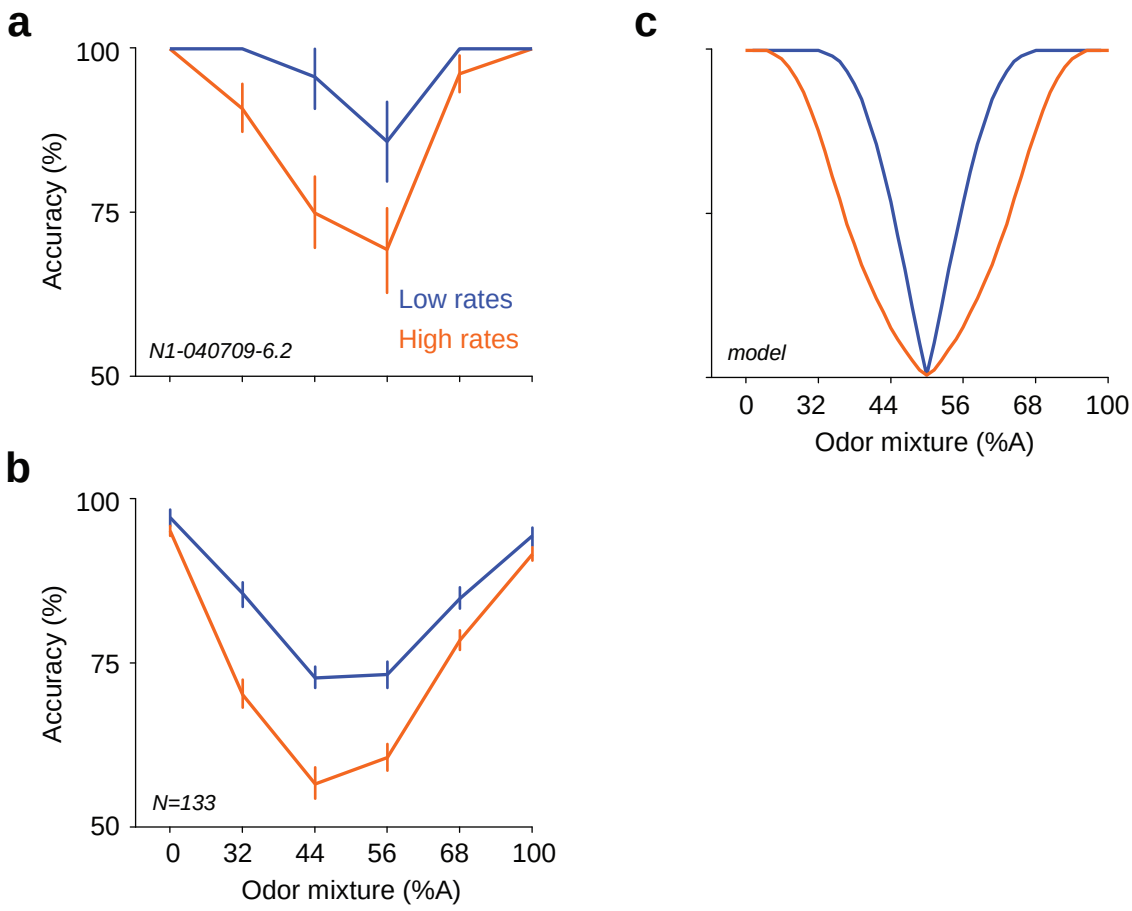

Supplementary Figure 2. Neural activity predicts behavioural accuracy beyond stimulus information

a, Behavioural accuracy as a function of stimulus and rate for a single neuron (same as in Fig 4.e,f). Mean accuracy +/- s.e.m is plotted as a function of stimulus and neural firing rate. Trials with at least one spike were assigned to low or high firing rates according to whether the spike count was above or below the median. **b**, Average behavioural accuracy as a function of stimulus and rate for the negative outcome selective neuron population (Fig 4.g,h). Error bars represent s.e.m. across neurons. **c**, Accuracy as a function of stimulus and decision uncertainty in the model. Trials are divided into low and high uncertainty groups (below and above median uncertainty levels).

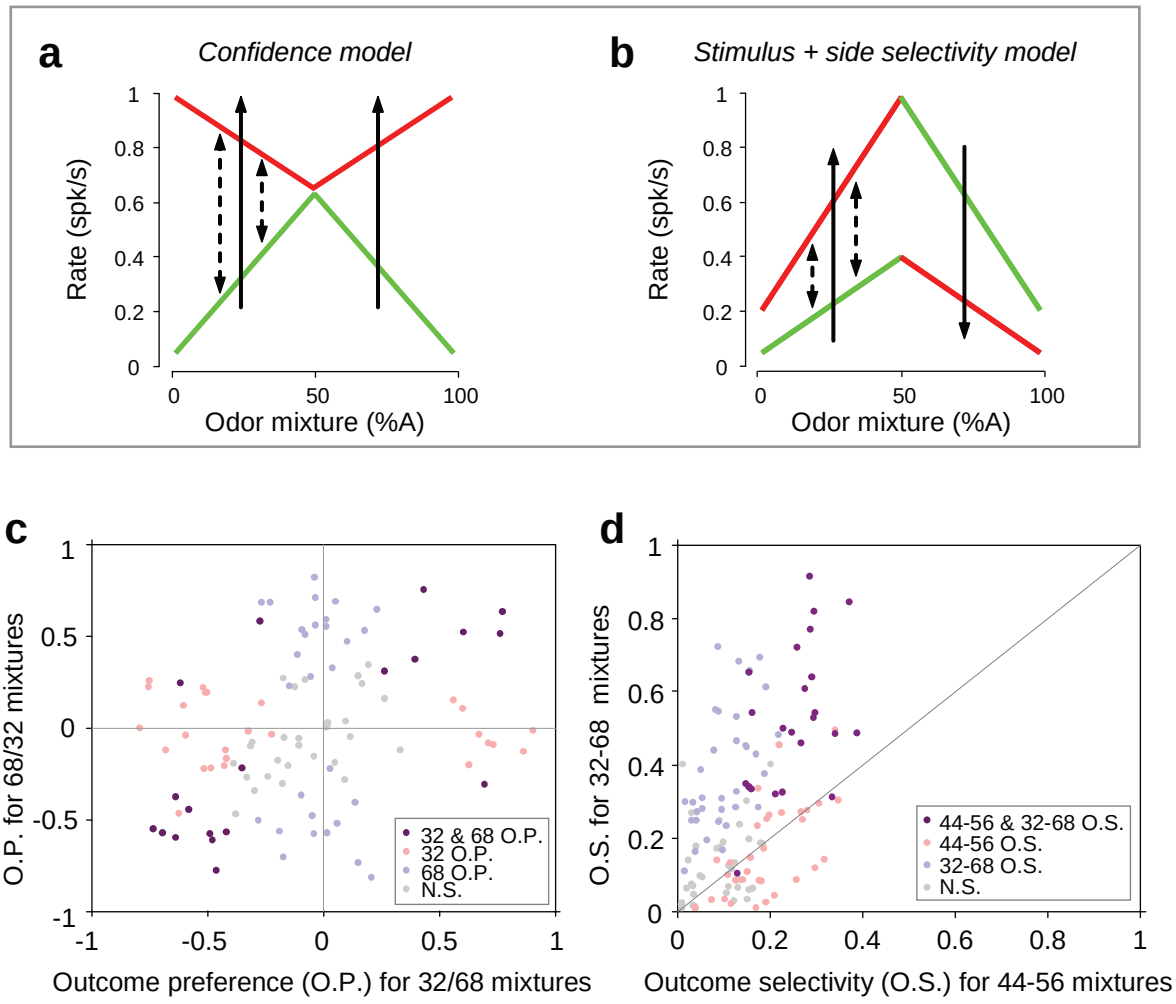

Supplementary Figure 3. Outcome selectivity follows predictions of the confidence model

a,b, Schematics illustrating the predictions of the confidence model (a), and the side selectivity model (b), with colours denoting the outcome of a choice (correct: green; error: red). Dashed arrows signify the distance between error and correct choices of a given difficulty. The direction of solid arrows signifies whether error or correct choices have higher rates for a given stimulus. **c,** Outcome preference index (O.P.) for 32/68 mixtures as a function of O.P. for 68/32 mixtures. All neurons included were significant when stimuli were pooled and colours show (see inset) whether these values were significantly outcome selective when considered separately ($P < 0.05$, permutation test). **d,** Outcome selectivity index (O.S.) for 32-68 mixtures (32/68 and 68/32 combined) as a function of O.S. for 44-56 mixtures (44/56 and 56/44 combined). 68/32 mixtures. All neurons included were significantly outcome selective when stimuli were pooled and colours show (see inset) whether these values were significant when considered separately ($P < 0.05$, permutation test).

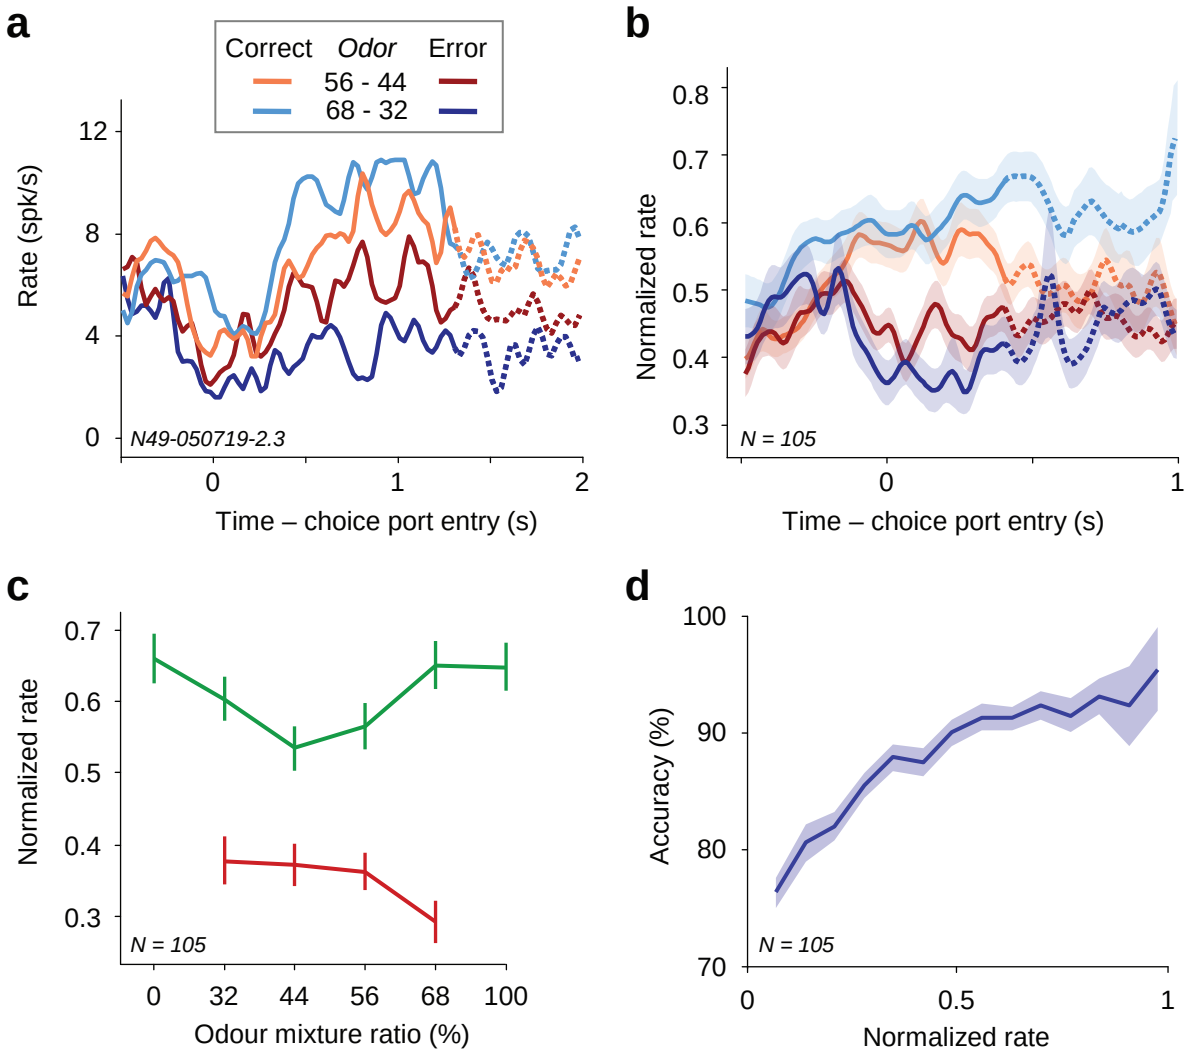

Supplementary Figure 4. Positive outcome selective neural population

a, Activity of an example neuron showing positive outcome selectivity. Firing rate of a single cell aligned to the time of entry into the choice port. Trials are grouped by stimulus difficulty (56/44 and 44/56, redish; 68/32 and 32/68 blueish colours) and trial outcome (correct/error, light/dark colours; see inset). Only activity occurring before the onset of water delivery or choice port exit is averaged into the PSTH. After the anticipation period (1.3 s in this session) the PSTH curves are dashed, signifying a time period during which in some trials rats experienced reward delivery. **b**, Mean normalized firing of the positive outcome selective population (those showing increased firing rate in error trials during the anticipation period). Only activity occurring while the rat was in the choice port and waiting for a reward is averaged. Dashed curves signify time at which reward was experienced in a fraction of sessions. Legend is same as for panel (a). **c,** Population tuning curves for firing rate during the initial 0.4 s outcome anticipation period as a function of stimulus type and outcome for the same population as in (b). Individual tuning curves were normalized and error bars represent s.e.m. across neurons. **d,** Mean accuracy as a function of the firing rate for the same neural population as in b. Firing rates were binned for individual neurons and the mean accuracy was calculated for each range of firing rates. These curves were normalized to a maximal firing rate of 1 and averaged. Error bars represent s.e.m. across neurons.

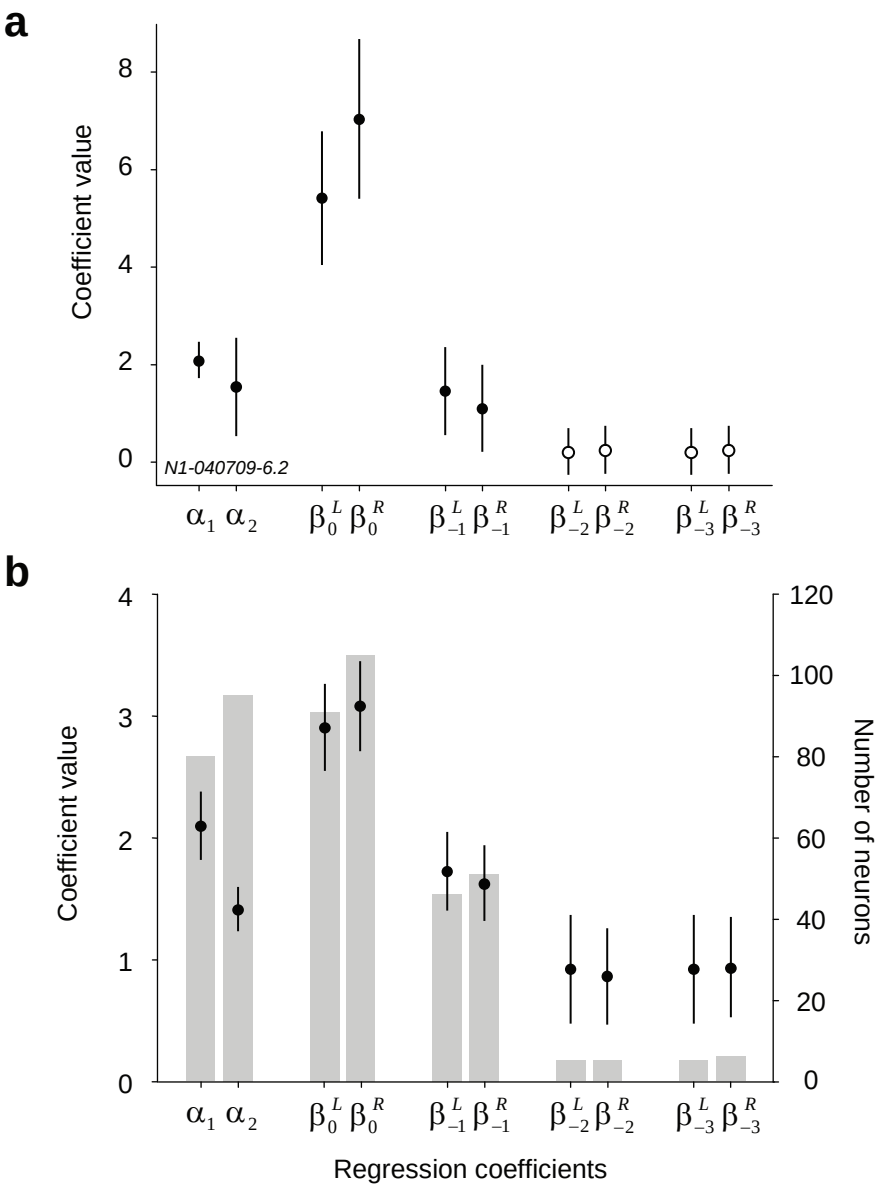

Supplementary Figure 5. Impact of outcome history on firing rates

a, Regression coefficients for the multiple regression analysis (see Methods) for the neuron presented in Figure 2a,b. Error bars mark the standard deviations of the coefficients estimated using leave-one-out-bootstrap validation. Filled circles indicate coefficients that are significantly different from 0 ($P < 0.05$, permutation test). $\beta_t^{L/R}$ are the coefficients for the outcome at trial t ($t = 0$ for the current trial) on either the left or right (L/R) choice sides. α_1 and α_2 measure the influence of the stimulus difficulty and the choice. **b**, Mean regression coefficients (left scale) for the population of negative outcome selective neurons (n = 133) and their standard error. Bar plots show the number of neurons (right scale) with significant regression coefficients ($P < 0.05$, permutation test). Only neurons with significant regression coefficients are included in the averages.

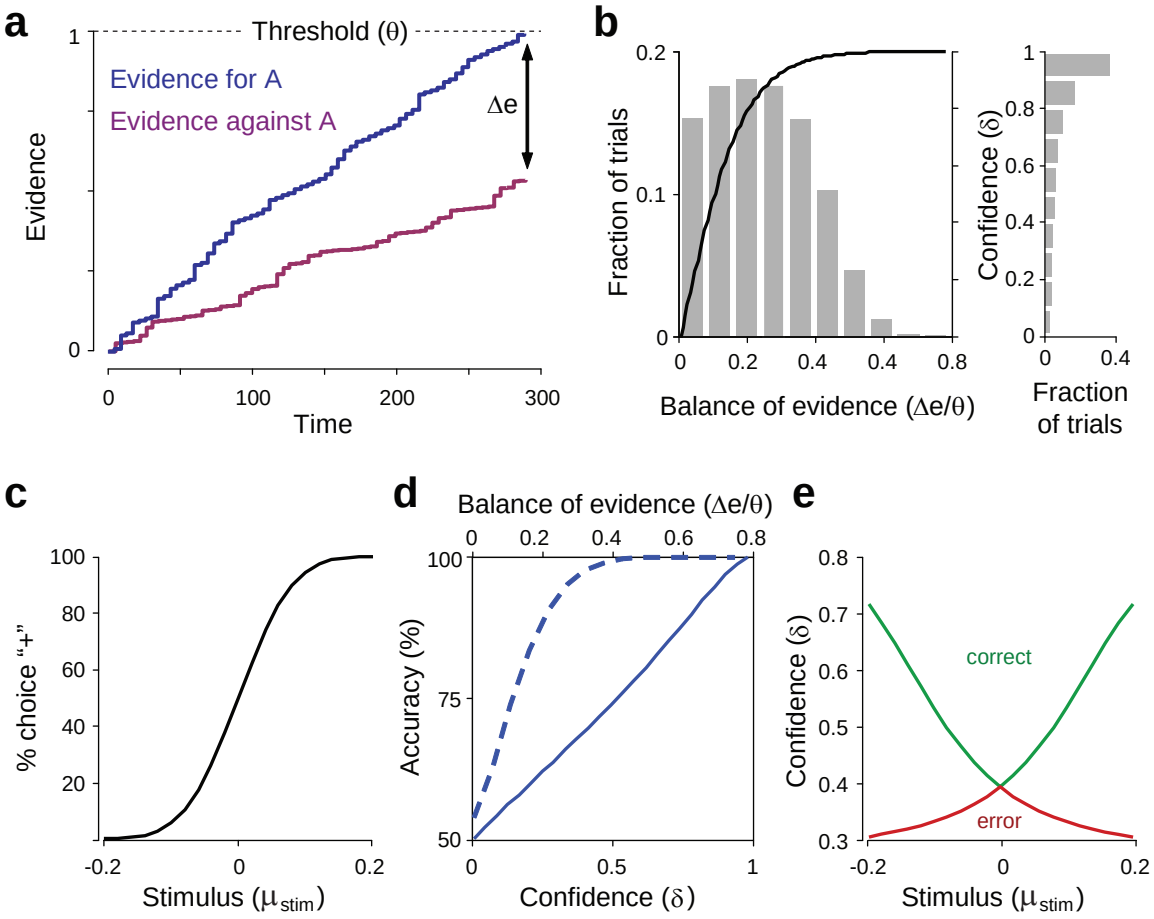

Supplementary Figure 6. Computing confidence in a "race model" of decision making

a, Schematic of a trial for the race model of decision making. Evidence for and against an alternative "A" accumulates separately over time. When the accumulated evidence crosses a threshold level a decision is generated its favour. Confidence about the decision may be estimated based on the distance between the decision variables, Δe. **b**, Calibration of Δe provides the estimate of confidence, δ. Left histogram shows the distribution of Δe across trials with the calibration function overlaid. The resulting distribution of δ is displayed on the right. **c,** Percentage of choices towards the "+ hypothesis", $\mu_{stim} > 0$, as a function of stimulus. **d,** Choice accuracy as a function of confidence, δ (solid line), and balance of evidence, $\Delta e/\Theta$ (dashed line). **e,** Mean confidence estimates generated by the model conditioned on stimulus and trial outcome.

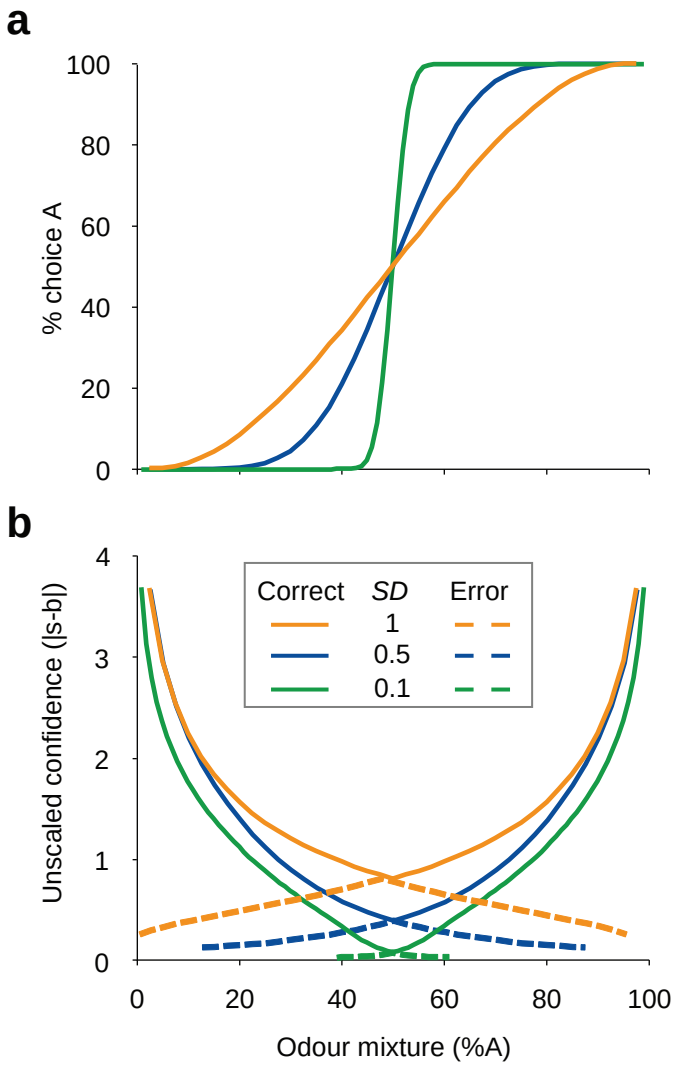

Supplementary Figure 7. Parameter robustness of decision confidence patterns

a, Psychometric functions generated by the categorization model presented in the main text. The total noise, $\sigma_{noise} = (\sigma^2_{bound} + \sigma^2_{stim})^{1/2}$, was varied to produce psychometric functions with different slopes. b, Mean unscaled confidence estimates, $d_i = |s_i -- b_i|$, generated by the same model and plotted as a function of stimulus and trial outcome.

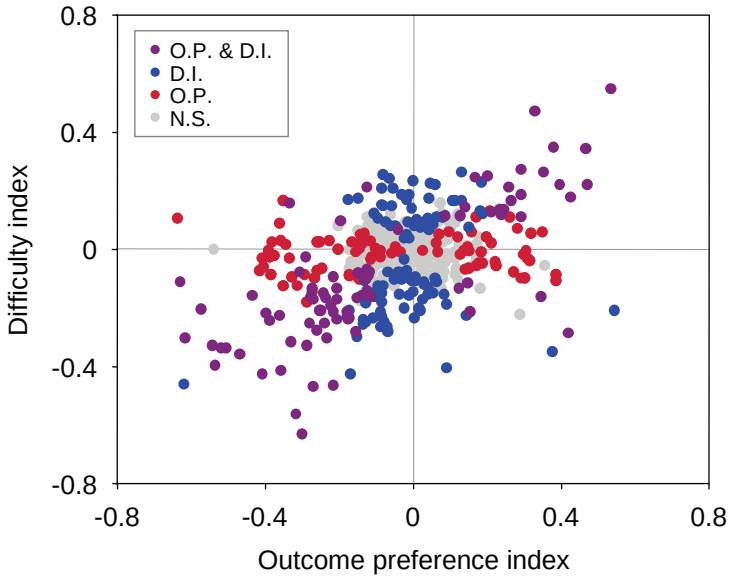

Supplementary Figure 8. Stimulus and outcome selectivity are correlated

Difficulty index (D.I.) as a function of outcome preference (O.P.) as described in the Methods. Each dot is one neuron. Colours show (see inset) whether these values were significant ($P < 0.05$, permutation test).