

Intrinsic rewards guide visual resource allocation via reinforcement learning

Received: 21 May 2025

Accepted: 11 August 2026

Published online: 01 October 2026

 Check for updatesIvan Tomić^{1,2}✉, Rodrigo Raimundo¹ & Paul M. Bays¹✉

Humans and other animals prioritize visual processing of stimuli that signal rewards. While previous research has focused on tangible incentives (for example, money or food), the effects of intrinsic incentives—such as perceived competence—are less well understood. Here, across a series of visual estimation experiments, we manipulated observers' subjective sense of confidence in their judgements using either deceptive trial-by-trial feedback or real discrepancies in stimulus reliability. We found that observers prioritized encoding of stimuli associated with lower uncertainty or error, benefiting performance for stimuli already estimated accurately, while further impairing performance for those estimated poorly. These reward-driven biases, while potentially adaptive, impaired overall accuracy in the present tasks by causing resource allocation to deviate from the error-minimizing strategy. To account for these findings, we supplemented a normalization model of neural resource allocation with a simple reinforcement learning rule. Intrinsic and external rewards cumulatively shaped the values assigned to different stimuli by the model, and the resulting discrepancies biased resource allocation and thereby estimation error, quantitatively matching the data. These findings reveal how intrinsic reward signals can shape resource allocation in ways that are both adaptive and counterproductive, offering a computational basis for the motivational biases underlying cognitive performance.

To support adaptive behaviour and ensure survival, the brain has evolved to prioritize environmental cues that signal potential rewards^{1,2}. Selectively attending to reward-predicting stimuli facilitates efficient navigation of complex environments, helping organisms move towards more rewarding states^{3,4}. This selection process is crucial given the brain's limited processing capacity, as it enhances internal representations of valuable stimuli and facilitates the formation of stimulus–reward associations⁵. Whereas the bias towards processing stimuli associated with tangible rewards is well established, the influence of intrinsic rewards—positive motivational states associated with feelings of satisfaction and competence⁶—on sensory processing remains less understood.

Experiments using points-based and monetary incentives have found that associating stimuli with a higher probability, or greater

magnitude, of external reward facilitates voluntary, or top-down, attention^{7–9}. In addition, in visual search tasks, which primarily engage bottom-up processes, search times are faster for pop-out targets associated with higher rewards than stimuli predicting less or no reward¹⁰. Notably, the prioritization of reward-associated stimuli persists in subsequent tasks even when reward contingencies are removed, and previously rewarded features cease to be salient or task relevant^{11–13}. Consistent with this, studies have shown that eye movements are biased towards objects and spatial locations previously associated with rewards^{14–16}. This continued prioritization of previously rewarded stimuli, even when it no longer aligns with immediate task goals, suggests that reward learning creates a lasting effect that can involuntarily bias attention towards these stimuli^{17,18}.

¹Department of Psychology, University of Cambridge, Cambridge, UK. ²Department of Psychology, Faculty of Humanities and Social Sciences, University of Zagreb, Zagreb, Croatia. ✉e-mail: ivn.tomic@gmail.com; pmb20@cam.ac.uk




Fig. 1 | External reward manipulation in Experiment 1. **a**, Schematic of the task. **b**, Illustration of the experimental manipulation. Responses within 50° of the target motion direction were rewarded with either 15 or 5 points, depending on the colour of the cued object. **c**, Distribution of response errors and corresponding fits of the neural resource model. Histograms represent the data, while coloured curves and shaded areas depict model predictions ($M \pm$ standard error (s.e.)). **d**, Corresponding MAD from experimental data (circles with error bars) and from the neural resource model (lines with error patches). Error bars and patches indicate the mean and 95% CI. Dashed line indicates chance level performance. **e**, Observed (that is, freely estimated) resource allocation compared with the optimal allocation aimed at maximizing the total points in the task. For visualization purposes, allocation towards the low-reward item is shown. Dashed line indicates equal allocation. Allocation smaller than 1 indicates that more resource was allocated towards the high-reward item (1:0.695 for high-reward versus low-reward item). Error bars indicate the mean and 95% CI. The inset shows individual reward functions relating resource allocation to expected point totals, with each curve's peak indicating the allocation that maximizes expected reward. For ease of visualization, only a subset of observers is shown, and all curves are normalized to the same total reward.

$N = 30$. **f**, Normalization model of neural resource allocation²². Stimulus features are encoded in the activity of a population of idealized direction-selective neurons. All neurons are assumed to have identical bell-shaped (von Mises) tuning curves with concentration κ , translated across the feature space to peak at each neuron's individual preferred feature value, providing dense uniform coverage of the feature space. Each stimulus is encoded by an independent subpopulation of M neurons, and parameter α represents a multiplicative gain factor related to attention allocated to each stimulus—and therefore to the activity of each subpopulation (Methods). Divisive normalization operates across the entire population, scaling summed activity to a fixed level determined by a gain constant. Neurons generate spikes according to a Poisson process, with mean firing rate determined by the normalized input of each neuron. The joint response of the population to a particular stimulus can be visualized as a hill of activity, centred on the true stimulus feature. The subsequent behavioural response was modelled as ML decoding of the spiking activity. Examples of error distributions predicted by the neural resource model are shown on the right. As the resources allocated to a stimulus (gain) decreases, the distribution of errors becomes increasingly wide and heavy tailed⁵⁶.

The influence of external rewards on behaviour extends to visual working memory (VWM)¹⁹, which is known for its ability to flexibly store and maintain features of multiple objects within a limited capacity^{20–28}. The precision of representations increases as a function of the associated reward, indicating that VWM allocation also tracks reward values when multiple objects provide different rewards^{29,30} (see Allen et al.³¹ for a review). Objects that were previously associated with reward are also better remembered, even when they are currently task irrelevant³². Crucially, however, total VWM capacity does not show flexibility with reward^{33,34}, which is further evidenced by findings that improved performance for high-reward items is accompanied by a corresponding decline in performance for low-reward items³⁵. These results demonstrate that stimuli can be strategically prioritized for encoding in VWM through selective attention, leading to flexible allocation of limited capacity between items based on their assigned subjective values^{36,37}.

Many everyday behaviours lack immediate external feedback, making internally generated signals—that is, intrinsic rewards—potentially equally relevant. One hypothesized source of such internal guidance is the subjective sense of accuracy or confidence^{38,39}. Although the neural bases and computational processes underlying intrinsic rewards⁶ and metacognition⁴⁰ are diverse, recent work suggests that confidence-related signals constitute an important intrinsic driver of learning. Specifically, reinforcement-like learning can emerge solely from an observer's internal estimate of confidence, used as a proxy for external feedback^{41,42}. The dopaminergic system—and the striatum in particular—which has long been implicated in encoding the motivational value of actions⁴³, has also been shown to encode intrinsically generated motivational signals⁴⁴, including confidence^{42,45,46}. Consistent with this, neural signatures typically associated with external feedback are also elicited when individuals experience an internal sense of being correct or incorrect^{47–50}. More broadly, the idea that

intrinsic signals can support reinforcement learning (RL) has a long history in the computational literature^{51–53}. The present work identifies confidence-related signals within this framework as a potential driver of attentional allocation.

In the present study, we combined psychophysical measurement and computational modelling to investigate how different intrinsic and external factors affect the competition between visual stimuli for processing resources. We used a modified analogue report task^{54,55} in which observers were instructed to reproduce the direction of one of a pair of motion stimuli that differed in their associated history of reward. Across a series of experiments, we found that performance was consistently better for stimuli previously associated with larger external reward, but also for those associated with lower uncertainty or with improved performance feedback. To provide a mechanistic explanation of the observed behaviour, we developed a computational model that relates accumulation of past rewards, both intrinsic and external, to allocation of neural resources between stimuli, which in turn influences estimation performance.

Results

Differential rewards bias resource allocation

Building on existing evidence that external rewards can bias information processing, we began our investigation by quantifying their effects on representational fidelity in a motion reproduction task. In Experiment 1, observers viewed two coloured motion stimuli, and after a brief delay and the presentation of a colour cue, they were asked to reproduce the motion direction of the cued stimulus (Fig. 1a,b). Critically, in this experiment, we associated the colours of the stimuli with different external rewards by awarding accurate recall ($<50^\circ$ absolute error) with 15 points when items of 1 colour were tested versus 5 points for the other colour. Accumulated points were converted into a bonus payment to the observer. At the end of the experiment, all observers correctly identified which stimulus had provided the larger rewards. To determine whether the difference in external rewards influenced reproduction precision, we compared the mean absolute deviation (MAD) of responses between stimuli of different colours. We found strong evidence that response errors were smaller for items of the colour associated with the larger reward (Bayes factor (BF_{10}) = 18.7; median of the posterior over effect size δ = 0.575; 95% credible interval (CI) = (0.195, 0.966)) (Fig. 1c,d).

Neural resource allocation. The results of Experiment 1 indicate that observers prioritized encoding the stimulus associated with the larger reward. Importantly, recall for the low-reward item remained reliably better than chance, suggesting that prioritization was graded rather than all-or-none. To quantify the share of resources allocated to each item, we applied a normalization-based population coding model^{22,56} to the data from Experiment 1 (Fig. 1f). In this model, neural firing rate takes the role of a limited resource, which, in the simplest scenario, would be equally distributed between stimuli. Here we extended this model by freely fitting a gain modulation parameter, α , which increased the activity encoding one of the two stimuli while keeping the total gain (that is, mean activity) of the population constant (see the ‘Neural resource model’ section for more details). Consistent with the observed difference in response error, we found strong evidence for unequal resource allocation favouring items of the highly rewarded colour (low/high ratio 0.695; difference from equal allocation, BF_{10} = 23.8; δ = 0.594; 95% CI = (0.211, 0.987)) (Fig. 1e; maximum likelihood (ML) estimates for the other parameters are shown in the Supplementary Information).

We next investigated whether observers distributed resources in a way that would maximize the total number of collected points, which we considered an optimal allocation strategy for this task. To test this, we calculated the expected number of points awarded for a range of different allocation weights (see the ‘Optimal resource allocation’ section for more details). The values of α that maximized the reward

are shown in Fig. 1e (optimal allocation). Comparing the observed and optimal weights revealed strong evidence for a difference between the two (BF_{10} = 92.9; δ = 0.694; 95% CI = (0.297, 1.102)), with observers distributing resources more equally than would be required to maximize the total number of points (low/high ratio 0.34). Consequently, observers earned fewer points per trial than would be expected under the reward-maximizing allocation (observed = 7.69; optimal = 7.92; BF_{10} = 60.2). Finally, we found moderate evidence against a correlation of observed and optimal allocations (r = 0.21; BF_{10} = 0.26).

A reward-maximization strategy would come at the cost of higher error for the low-reward item. This could suggest that, in addition to maximizing external rewards, observers may also aim to achieve a certain level of accuracy on the task across all items, potentially because they find accuracy intrinsically rewarding (see also ref. 25).

Perceived accuracy biases resource allocation

Having confirmed that external rewards modulated allocation in the motion reproduction task, we next investigated effects of perceived accuracy, a possible form of intrinsic reward, on the same task. In Experiment 2a we presented manipulated feedback at the end of each trial to influence observers’ perception of their reproduction accuracy. Studies using false-feedback manipulations have shown that such manipulations can modulate subjective confidence⁵⁷, providing a principled basis for our novel trial-by-trial feedback manipulation. Observers were again presented with two coloured stimuli, and reproduced one indicated by a colour cue. We magnified the error presented at feedback when one colour was cued and minimized the error at feedback for the other colour (Fig. 2a,b). A post-experimental questionnaire revealed that 84% of observers judged stimuli of the colour associated with error-magnified feedback as more difficult to remember, indicating that we successfully associated stimulus identity (that is, colour) with perceived difficulty.

To assess the effects of perceived difficulty on response precision, we compared MAD between the response and the true target direction (rather than the one shown as feedback) for stimuli of the two colours (Fig. 2c,d). We found responses to be more precise for the stimulus with reduced feedback error, that is, the one perceived as easier to remember (BF_{10} = 29.4; δ = 0.672; 95% CI = (0.24, 1.12)). This finding indicates that the perception of better performance for stimuli of one colour, induced by feedback, led to improved actual performance for those stimuli.

The observed effect could be attributed to either capture of visual attention by the ‘easier’ item (that is, competition for visual processing resources) or the mnemonic prioritization of that item (that is, competition for memory resources). To differentiate between these possibilities, we conducted a follow-up experiment. Experiment 2b replicated the conditions of Experiment 2a but with stimuli presented sequentially to reduce encoding competition between the two objects and minimize the influence of attentional selection on resource allocation. Similar to Experiment 2a, 89% of observers judged the colour associated with magnified feedback errors as more difficult to remember. However, in contrast to Experiment 2a, comparing response errors across the two stimuli (Supplementary Fig. 1) revealed that the observed data were 9 times more likely under the null hypothesis, providing moderate evidence for a lack of difference in response precision between the two colours (BF_{10} = 0.11; δ = 0.009; 95% CI = (−0.183, 0.202)). This finding suggests that the effect observed in Experiment 2a is likely due to attentional competition during encoding. When that competition is mitigated, observers do not show preferential encoding based on perceived difficulty.

In addition, we verified whether the results of Experiment 2a could be explained by one object appearing easier to reproduce, the other appearing more difficult or their combined effect. To this end, we conducted a follow-up, Experiment 2c, which was identical to Experiment 2a except that observers took part in two sessions. In both sessions,




Fig. 2 | Perceived accuracy manipulation in Experiment 2a. **a**, Schematic of the task. **b**, Experimental manipulation illustration. Feedback error was magnified for one stimulus and minified for the other, based on the colour of the cued object. **c**, Distribution of response errors and corresponding fits of the neural resource model. Histograms represent the data, while coloured curves and shaded areas depict model predictions ($M \pm \text{s.e.}$). **d**, Corresponding MAD from experimental data (circles with error bars) and from the neural resource model (lines with error patches). Error bars and patches indicate the mean and 95% CI. Dashed line indicates chance level performance. **e**, Observed resource allocation

and optimal allocation aiming to minimize overall feedback variance in the task. Dashed line indicates equal allocation. Allocation larger than 1 indicates that more resource was allocated towards the error-minified item (minified versus magnified: 1.18:1). Error bars indicate the mean and 95% CI. The inset illustrates individual variability in feedback error as a function of resource allocation, with each curve's trough indicating the allocation level that minimizes feedback error. For ease of visualization, only a subset of observers is shown, and all curves are normalized to the same range of feedback variance. $N = 25$.

feedback for one object was veridical, while feedback for the other object was minified in one session and magnified in the other session. Different colour pairs were used to prevent transfer of learned associations across sessions. Testing the directional hypotheses based on Experiment 2a, we found moderate evidence that the minified item was recalled more accurately than the veridical item ($\text{BF}_{10} = 7.9$; $\delta = 0.557$; 95% CI = (0.12, 1.046)) (Fig. 3), but no evidence that the magnified item was recalled less accurately than the veridical item ($\text{BF}_{10} = 0.64$; $\delta = 0.253$; 95% CI = (0.017, 0.659)). In addition, we found strong evidence that the minified item was recalled more precisely than the magnified item ($\text{BF}_{10} = 22.5$; $\delta = 0.677$; 95% CI = (0.197, 1.046)), replicating the results from Experiment 2a. These results indicate that the error minifying manipulation alone was sufficient to bias resource allocation, whereas the error magnifying manipulation was not.

Neural resource allocation. The results of Experiment 2a show that observers prioritized encoding of the error-minified stimulus, that is, the one signalling better performance. Crucially, the error-magnified stimulus was still recalled with above-chance precision, consistent with a graded rather than all-or-none allocation of resources. To quantify resource distribution between the two objects, we again applied the neural resource model to the data, with results illustrated in Fig. 2c,e. We found that, on average, observers allocated 1.18 times more resources towards the error-minified stimulus (difference from equal allocation, $\text{BF}_{10} = 2.63$; $\delta = 0.45$; 95% CI = (0.05, 0.86); 19 out of 25 observers had $\alpha_{\text{observed}} > 1$) (Fig. 2d), consistent with the observed difference in response error between the stimuli of two colours.

Next, we investigated whether the observed allocation matched the predictions of an ideal observer who optimally weights neural activity to minimize overall feedback error in the task. We calculated the expected variance of feedback error across both items for a range of different allocation weights, and Fig. 2e shows optimal allocation weights that minimize this variance. The optimal strategy would require shifting

twice as many resources towards the error-magnified item ($\alpha_{\text{optimal}} = 0.52$). Importantly, we found extremely strong evidence that this was inconsistent with the observed allocation, which favoured the error-minified item ($\text{BF}_{10} = 3.77 \times 10^6$; $\delta = 1.72$; 95% CI = (1.09, 2.39)). Overall, these results indicate that observers did not adopt an allocation strategy that would minimize their feedback error variability ($\alpha = 0.52$), but instead did the opposite, allocating more neural resources to the item for which we systematically minified the error in feedback. Accordingly, observers' allocation resulted in higher overall feedback error than predicted by the optimal strategy ($\text{MAD}_{\text{observed}} = 0.535$; $\text{MAD}_{\text{optimal}} = 0.476$; $\text{BF}_{10} = 9.92 \times 10^3$). We also found moderate evidence against a correlation of observed and optimal allocations ($r = 0.23$; $\text{BF}_{10} = 0.29$).

In Experiment 2b, fitting the same neural resource model to the data revealed that the observed allocation parameter was numerically close to 1 ($\alpha_{\text{mean}} = 1.07$; $\text{BF}_{10} = 0.62$; $\delta = 0.18$; 95% CI = (−0.01, 0.38)), which aligns with the observed similarity in reproduction precision between the two stimuli. This further supports the conclusion that the effect observed in Experiment 2a depended on attentional competition during encoding. In Experiment 2c, more resource was allocated to the error-minified stimulus compared with the stimulus with veridical feedback ($\alpha_{\text{mean}} = 1.21$; $\text{BF}_{10} = 2.1$; $\delta = 0.47$; 95% CI = (0.022, 0.951)), while resource allocation was found to be comparable between the error-magnified stimulus and the veridical feedback stimulus ($\alpha_{\text{mean}} = 0.96$; $\text{BF}_{10} = 0.27$; $\delta = 0.11$; 95% CI = (−0.305, 0.535)), consistent with the observed pattern of reproduction errors.

Estimation difficulty biases resource allocation

Following Experiment 2, we aimed to determine whether preferential allocation and encoding would persist when varying objective stimulus difficulty rather than perceived performance. Drawing on previous findings showing a positive correlation in humans between subjective confidence and the motion strength of random dot kinematograms




Fig. 3 | Perceived accuracy manipulation in Experiment 2c. a, Distribution of response errors for the two tasks: in one task, feedback for one object was minified and veridical for the other object (top row); in the other task, feedback for one object was magnified and veridical for the other object (bottom row). Histograms represent the data, while coloured curves and shaded areas depict

model predictions ($M \pm \text{s.e.}$). **b**, Corresponding MAD from experimental data (circles with error bars) and from the neural resource model (lines with error patches). Error bars and patches indicate the mean and 95% CI. Dashed line indicates chance level performance. $N = 19$.

(RDK) stimuli⁵⁸ (see also ref. 40), we hypothesized that variations in the objective difficulty of stimuli would modulate internally generated confidence signals, driving the prioritization of specific stimuli as in Experiment 2. In Experiment 3a, we presented two coloured RDK stimuli with different coherence levels on the majority of trials to create differences in objective difficulty and associated confidence. We then assessed response precision on the remaining trials, during which both stimuli were presented with equal coherence (that is, equal difficulty) (Fig. 4a,b).

In Experiment 3a, all observers reported that stimuli of the colour associated with low coherence were more difficult to remember, confirming that the coherence manipulation produced a clear difference in perceived difficulty, despite the absence of performance feedback in this experiment. Results of Experiment 3a are shown in Fig. 4c–e. As expected, observers were more precise in reproducing the motion direction of the high-coherence stimulus on trials where the stimuli objectively differed in difficulty ($\text{BF}_{10} = 1.83 \times 10^5$; $\delta = 1.6$; 95% CI = (0.95, 2.29)). More importantly, on trials where the stimuli had equal coherence, reproduction was also more precise for the stimulus with the colour associated with high coherence (that is, the ‘easier’ colour) ($\text{BF}_{10} = 11.4$; $\delta = 0.63$; 95% CI = (0.18, 1.10)). This finding suggests that observers associated colour with difficulty when the two objects were presented with different levels of coherence, and subsequently allocated more resources to the stimulus they had learned was easier.

This result was replicated in the laboratory setting of Experiment 3b, where observers were required to maintain eye fixation at the centre of the screen during stimulus presentation (Fig. 4g–i). Despite preventing observers from overtly shifting their attention towards 1 stimulus during encoding, 74% of observers correctly identified 1 colour as more difficult. As expected, responses were more precise for the high-coherence stimulus on trials where stimuli differed in coherence ($\text{BF}_{10} = 4,414$; $\delta = 1.366$; 95% CI = (0.718, 2.046)). In addition, responses were more precise for the colour associated with high coherence on trials where both stimuli were presented with equal coherence ($\text{BF}_{10} = 4.2$; $\delta = 0.56$; 95% CI = (0.092, 1.052)). However, the observed difference could again be explained by attentional demands at encoding. When objects were presented sequentially (Supplementary Fig. 2), response precision was comparable across colours when both stimuli had the same coherence ($\text{BF}_{10} = 0.51$; $\delta = 0.267$; 95% CI = (−0.157, 0.709)). This was despite observers being able to judge which item was more difficult

(84%) and a noticeable precision advantage for the high-coherence stimulus on trials when coherence levels varied between objects ($\text{BF}_{10} = 9.22 \times 10^4$; $\delta = 1.76$; 95% CI = (1.015, 2.550)). Compared with other experiments, response distributions in this experiment exhibit more pronounced peaks around the direction opposite to the target (that is, elevated tail ends), which also produces some exaggeration of MAD compared with the model. The tendency of our sensory system to encode orientation of a motion path (that is, the line on which movement occurs) partly independently of direction is well documented^{59,60} and may be especially pronounced when motion stimuli are presented in the periphery rather than at fixation.

Neural resource allocation. Consistent with the findings from Experiment 2, Experiment 3 demonstrated that observers, when presented with objects associated with different levels of performance, prioritized the encoding of the stimuli perceived as easier. Also consistent with previous experiments, observers performed above chance for the more difficult item, supporting the interpretation that resource allocation was graded rather than all-or-none. To quantify the distribution of resources across the two items, we again applied our population coding model to the data.

In Experiment 3a, the allocation estimates from the model indicated that observers allocated nearly twice as much resource (1.76:1) to the stimulus colour associated with higher coherence (Fig. 4f), and this allocation deviated from equal allocation ($\text{BF}_{10} = 2.68 \times 10^4$; $\delta = 1.4$; 95% CI = (0.792, 2.031)). Interestingly, in this experiment—but not in others—we found evidence that observers with a higher gain parameter on average allocated their resources more equally ($r_{y,\alpha} = -0.58$; $\text{BF}_{10} = 10.53$; 95% CI = (−0.78, −0.18)). We next investigated whether the observed allocation was consistent with an optimal allocation strategy aimed at minimizing overall response variance in the task. To this end, we simulated performance on the variable coherence trials using a range of different allocation weights, and found that the optimal strategy for most observers was equal allocation (Fig. 4f). Comparing the observed and optimal weights revealed strong evidence that the observed weights were, on average, larger than the optimal weights ($\text{BF}_{10} = 9,580$; $\delta = 1.341$; 95% CI = (0.733, 1.977)). This allocation resulted in higher overall response error than predicted by the optimal strategy ($\text{MAD}_{\text{observed}} = 0.627$; $\text{MAD}_{\text{optimal}} = 0.569$; $\text{BF}_{10} = 147.8$).




Fig. 4 | Estimation difficulty manipulation in Experiment 3a and 3b

(simultaneous presentation). **a**, Schematic of the task. **b**, Experimental manipulation illustration. In most trials, the two colours were associated with different levels of motion estimation difficulty (that is, variable coherence); in the remaining trials, both objects had the same level of difficulty (that is, equal coherence). Motion with different coherence levels produces varying degrees of perceptual noise, with higher coherence reducing noise. This perceptual noise was incorporated into the neural resource model as an additive component, alongside memory noise. **c, d**, Distribution of response errors and corresponding fits of the neural resource model. Histograms represent the data, while coloured curves and shaded areas depict model predictions ($M \pm s.e.$). Panel **c** depicts variable coherence trials, and **d** depicts equal coherence trials. **e**, Corresponding MAD from experimental data (circles with error bars) and from the neural resource model (lines with error patches). Error bars and patches indicate the mean and 95% CI. Dashed line indicates chance level performance. Inter. refers to

the intermediate (that is, 65%) coherence level. **f**, Observed resource allocation (estimated across all conditions) and optimal allocation (estimated across variable coherence trials) aiming to minimize overall recall variance in the task. Dashed line indicates equal allocation. Allocation larger than 1 indicates that more resource was allocated towards the easier item (high versus low coherence: 1.76:1). Panels **g–j** are the same as **c–f**, but for the simultaneous presentation condition of Experiment 3b. Panels **g** and **h** depict variable and equal coherence trials, respectively; panel **i** depicts the corresponding MAD values, while panel **j** shows observed and optimal resource allocation. **j**, Allocation larger than 1 indicates that more resource was allocated towards the easier item (high versus low coherence: 1.61:1). The insets illustrate individual variability in response variance as a function of resource allocation, with each curve's trough indicating the allocation level that minimizes overall response error. Error bars indicate the mean and 95% CI. For ease of visualization, all curves are normalized to the same range of recall variance. $N = 22$ in Experiment 3a; $N = 19$ in Experiment 3b.

Interestingly, our simulation shows almost no variation in estimated optimal allocations (Fig. 4f,j). This shows that allocating resources to the low-coherence item—at least in the current experimental setting—cannot effectively mitigate the effects of perceptual noise introduced by manipulating coherence, without introducing substantial error for the high-coherence item. The only exception to near-equal allocation occurs when the effect of perceptual noise for the low-coherence item is so great as to render estimates effectively random. In this case, resources can be reallocated to improve precision for the high-coherence item without worsening performance for the low-coherence item.

These findings were replicated in Experiment 3b. When objects were presented simultaneously, the model estimated that observers allocated resources at a ratio of 1.61:1 in favour of the high-coherence stimulus (Fig. 4j). This allocation was again different from equal allocation ($BF_{10} = 830.6$; $\delta = 1.166$; 95% CI = (0.566, 1.794)) and from optimal, which was again close to equal allocation (mean $\alpha_{\text{optimal}} = 1.03$; $BF_{10} = 789$; $\delta = 1.16$; 95% CI = (0.561, 1.786)). Finally, fitting a free allocation parameter to the data from the equal coherence condition with sequential presentation (Experiment 3b) revealed a ratio of 1.2:1 in favour of the colour associated with high coherence; however, we did not find evidence that this was different from equal allocation ($BF_{10} = 0.82$; $\delta = 0.343$; 95% CI = (-0.09, 0.797)).

Interim conclusion

In Experiment 1, we observed a clear effect of external rewards on representational fidelity in a motion reproduction task, with observers allocating more cognitive resources to high-reward items. In Experiments 2 and 3, we found a similar effect using novel manipulations, where observers allocated more resources to the stimulus that was perceived as easier, either based on manipulated error feedback (Experiment 2) or internal confidence in estimation (Experiment 3). In both of these experiments, the observed allocation deviated from predictions made by an optimal strategy aimed at minimizing overall feedback or response error. In addition, we demonstrated that differences in estimation performance were abolished when competition during encoding was removed by presenting stimuli sequentially, suggesting they arise from unequal allocation of attentional resources at the encoding stage.

Building on these results and existing literature^{50,61,62}, we argue that observers in our study found higher accuracy and confidence in their performance intrinsically rewarding, and learned to associate this intrinsic reward with a stimulus feature (that is, one of the two colours). This association biased resource allocation towards subsequent stimuli with the same feature. Importantly, although reward driven, this biased allocation was not a strategy that would maximize intrinsic reward on these tasks because observers had no influence over which stimulus was cued for report on a given trial. Indeed, the direction of the biases induced by implicit rewards in Experiments 2 and 3 meant that they were counterproductive: increasing overall error variability relative to a strategy of equal allocation. Therefore, instead of evaluating these data from the perspective of optimal performance, we propose a neural model inspired by RL to elucidate these findings.

RL model of resource allocation

To further explore the dynamics of resource allocation, we developed a computational model that integrates principles of neural coding and RL. The proposed RL account of resource allocation extends the neural resource model^{22,56} by incorporating a value-updating mechanism that allows external and intrinsic rewards to influence the future distribution of neural resources (Fig. 5). A key contribution of our model is the concept that rewards—both intrinsic and external—obtained from reproduction of a stimulus become associated with the identifying features of that stimulus, affecting their subjective value and biasing allocation of resources in subsequent encounters. We found that

this approach accurately predicted resource allocations estimated by freely fitted allocation weights, indicating that behavioural estimation performance could be successfully inferred from an analysis of accumulated rewards.

External reward. In Experiment 1, the stimulus colour associated with a high reward (15 points) was expected to accumulate greater value relative to the colour associated with a low reward (5 points). On average, observers earned points on 80% of trials when the high-reward stimulus was probed and 66.5% of trials when the low-reward stimulus was probed, leading to an average accumulation of 602 and 166 points, respectively. To apply the proposed RL model to each observer's data, we combined individual trial-by-trial external rewards with estimates of internal confidence (equation (9)).

Figure 6a shows the average trajectory of resource allocation across trials (see Supplementary Fig. 5a for individual trajectories). This trajectory shows an early shift in resource allocation towards the preferred item, followed by a stable plateau. For ease of visualization, trajectories are presented as directed towards the preferred object, defined as the object receiving a greater average resource allocation across all trials.

Crucially, since our RL account is grounded in the same neural resource model previously utilized to fit the psychophysical data and quantify resource allocation (Fig. 1a,c), we can directly compare estimates across the two models. Here we focus on the comparison of estimated resource allocations, while ML estimates and comparisons for the other parameters are shown in Supplementary Fig. 6a. Importantly, the freely estimated resource allocation (observed allocation in Fig. 1c) is based on behavioural errors only, with no information about rewards, and so can serve as a benchmark for evaluating performance of the RL model. As shown in Fig. 6b, we observed a strong positive correlation between the freely estimated allocation parameter and the mean allocation derived from the history of accumulated rewards ($r = 0.976$; 95% CI = (0.941, 0.988); $BF_{10} = 3.34 \times 10^{16}$). The close alignment of these two distinct methods suggests that the history of accumulated rewards can effectively account for resource allocation in this task.

Intrinsic reward: perceived accuracy. In Experiment 2 we manipulated the response error presented in feedback to influence the perceived difficulty of reproducing stimuli of each colour. This manipulation resulted in participants experiencing systematically larger feedback errors for stimuli of one colour (magnified feedback MAD = 0.837) than the other (minified feedback MAD = 0.233). To model this data within our RL account, we assume that the feedback on each trial provided an intrinsic reward that was associated with the corresponding stimulus colour. This assumption is supported by evidence that observers' subjective evaluations tend to favour smaller feedback errors over larger ones because they suggest higher accuracy⁶³.

In the model, rewards derived from feedback were integrated with those derived from internal confidence. Figure 6d illustrates the mean trajectory of resource allocation across trials (Supplementary Fig. 5b shows individual trajectories for example observers). The model fits again indicate that resources were unequally allocated between stimuli, although the bias is smaller than observed in the experiment with external rewards.

Comparing the estimated allocation derived from the RL account with the freely fitted allocation parameter in the neural resource model (Fig. 6e), we found a strong positive correlation ($r = 0.911$; 95% CI = (0.777, 0.958); $BF_{10} = 3.44 \times 10^7$). Consistent with the findings from Experiment 1, the correspondence between these two distinct approaches indicates that the history of accumulated intrinsic rewards provides an explanation for resource allocation in the task with manipulated feedback. ML estimates and comparisons for the other parameters are shown in Supplementary Fig. 6b.




Fig. 5 | The neural resource allocation account applied to the motion estimation task. **a**, On each trial, motion directions of the two stimuli are encoded in the spiking activity of populations of neurons, with mean activity determined by the relative allocation of resources to stimuli. On the basis of the cue colour, one of the populations is decoded to yield an estimated direction with an associated uncertainty that varies from trial to trial. The uncertainty of the estimate, the accuracy feedback (if present) and any points awarded represent different forms of intrinsic and external reward, which are combined as a weighted sum into a composite reward (Δ_t). This composite reward is then used to update the relative value (v) associated with the stimulus colours. Finally, this relative value is transformed via an exponential mapping into a neural gain factor (α), which controls the fraction of resources allocated to each stimulus on the subsequent trial. In this framework, resource allocation is entirely driven by the history of accumulated rewards. **b**, Throughout the reported experiments, the

two colours of stimuli are systematically related to different intrinsic or external rewards, so the relative value assigned to each colour progressively diverges over the sequence of trials. **c**, Fraction of total resources allocated to the high-reward stimulus over trials, based on relative value shown in **b**. Note that the fraction of resources allocated to an item is determined by the relative value accumulated up to, but not including, that trial. The dashed line represents the mean allocation across all trials (~65%). The remaining resources (~35%) are allocated to the low-reward stimulus. **d**, Unequal resource allocation is reflected in differences in the mean absolute error across trials when the high-reward or low-reward stimulus is cued for report. **e**, Larger reward weight (unspecified reward c), with a constant leak factor, results in a stronger preference for one stimulus over the other in terms of resource allocation. **f**, Larger leak factor (y), with a constant reward weight, leads to a weaker preference.

Intrinsic reward: estimation difficulty. Experiment 3 investigated the role of objective difficulty in the representation of motion information. On most trials, two stimuli with different coherence levels (85% and 45%) were presented. We hypothesized that internal confidence in each item's motion direction, reflecting a metacognitive estimate of

accuracy, functions as an intrinsic reward which observers associate with each item's identity (that is, colour)⁴². To model the psychophysical data in the simultaneous presentation condition, we estimated internal confidence by exploiting the close coupling between uncertainty and trial-to-trial variability in error within the neural resource model.




Fig. 6 | Modelling results. **a**, Resource allocation across trials inferred by the RL account in the external reward experiment (Experiment 1). Circles represent the mean fraction of resources across observers allocated on each trial towards the overall preferred object. **b**, Correlation between mean allocations inferred by the RL account and freely estimated allocations. The red line shows the mean prediction of the fitted linear regression model, and the shaded area indicates the 95% CI. **c**, Difference in MAD between trials on which the probed item had

below-median and above-median resources allocated to it, as estimated by the RL account. On average, MAD was larger when less resource was allocated to the probed stimulus. Error bars show the mean and 95% CI. $N = 30$. **d–f**, Same as panels **a–c**, but for the perceived accuracy experiment (Experiment 2). $N = 25$. **g–i**, Online estimation difficulty experiment (Experiment 3a). $N = 22$. **j–l**, Lab-based estimation difficulty experiment (Experiment 3b, simultaneous condition). $N = 19$.

Informed by observed response error on each trial, we derived the posterior probability distribution of likelihood precision and used the most probable precision as a basis for intrinsic reward (equation (12)). While internal confidence was also incorporated in this way when modelling data from the previous two experiments, in Experiment 3, it was the sole source of reward influencing resource allocation.

The mean trajectories from Experiments 3a and 3b are shown in Fig. 6g,j, respectively. Again, we visualized the obtained individual trajectories in example participants (Supplementary Fig. 5c,d). In both experiments, all parameter estimates obtained with the neural resource model and the RL account strongly covaried (Fig. 6h,k and Supplementary Fig. 6c,d). Importantly, this was also true for estimates of resource allocation. Across the two experiments, we found very consistent and strong positive correlations between the freely estimated allocation parameter and the mean allocation derived from the history of accumulated rewards (Experiment 3a: $r = 0.833$; 95% CI = (0.589, 0.923); $BF_{10} = 1.29 \times 10^4$; Experiment 3b: $r = 0.843$; 95% CI = (0.575, 0.933); $BF_{10} = 3.67 \times 10^3$). Consistent with the findings from the first two experiments, this strong correspondence indicates that the history of accumulated intrinsic rewards based on internal confidence effectively accounts for resource allocation in this task.

Changes in resource allocation predict response precision. Our finding that freely estimated resource allocation strongly correlates across participants with resource allocation based on the history of rewards supports the conclusion that human resource allocation is guided by a reward-driven value assignment to objects in the visual environment. To further substantiate this claim, we investigated whether variability in resource allocation across trials within individual participants, derived from the RL model, predicts the magnitude of their response errors.

To investigate this, we examined response errors as a function of model-predicted resource allocation, with the hypothesis that trials with lower allocated resources to the probed item would show greater errors. To this end, we performed a median-split analysis for each observer based on the estimated fraction of allocated resources towards the item associated with larger reward (that is, individual trajectories similar to those shown in Fig. 6, left column). Specifically, we calculated the MAD of response errors for trials with above-median and below-median resource allocation, separately for trials where the high-reward or low-reward item was tested. We hypothesized that MAD would be greater on trials where the RL model indicated that below-average resource was allocated to the probed item, that is, below-median trials when the high-reward object was probed and above-median trials when the low-reward object was probed. To test this, we computed a composite score for each observer equal to the sum of the signed difference in MAD between below-median and above-median trials when the high-reward item was probed and the signed difference in MAD between above-median and below-median trials when the low-reward item was probed. In all four experiments, the composite scores indicated that lower predicted resource allocation, based on the history of rewards, corresponded on average to larger MAD of response errors (Fig. 6c,f,i,l). This was confirmed with one-sided t -tests against zero, which provided moderate to extreme evidence for a difference in the hypothesized direction: Experiment 1 ($BF_{10} = 5.55 \times 10^4$; $\delta = 1.096$; 95% CI = (0.635, 1.570)); Experiment 2 ($BF_{10} = 564$; $\delta = 0.866$; 95% CI = (0.402, 1.345)); Experiment 3a ($BF_{10} = 3.74$; $\delta = 0.441$; 95% CI = (0.073, 0.874)); and Experiment 3b ($BF_{10} = 192.4$; $\delta = 0.919$; 95% CI = (0.377, 1.487)).

Discussion

In the present study, we investigated how human observers represent stimuli associated with varying levels of external and intrinsic reward. Across three psychophysical experiments, we paired object identities with different rewards and found observers developed higher estimation accuracy for the items associated with larger rewards. In two

additional experiments, we demonstrated that this effect was driven by competition for attentional, rather than mnemonic, resources. To provide a mechanistic explanation of this behaviour, we developed a neural model incorporating an RL rule that directs resource allocation towards more rewarding stimuli. Our key finding—based on our experimental data and computational modelling—is that a resource allocation mechanism based solely on the history of accumulated rewards is sufficient to explain differences in estimation performance based on intrinsic as well as external rewards.

In the first experiment, we investigated the effects of external rewards on representational fidelity in a motion reproduction task. Both the psychophysical results and computational modelling provided compelling evidence that observers allocated more processing resources to objects associated with a higher reward, resulting in more precise reproduction of high-reward stimuli compared with low-reward ones. This finding aligns with a broad body of research demonstrating that external rewards, such as points or money, influence various aspects of information processing, including the allocation of attentional resources^{17,18} and working memory³¹, while also facilitating motor responses, such as hand movements and saccades, towards rewarding stimuli^{63–65}.

Although intrinsic motivation has long been recognized as a component of RL^{51–53}, its influence on representational fidelity in perceptual and memory tasks has received comparatively little empirical attention. Building on the premise that accuracy itself is rewarding^{60,61,62}, we conducted two experiments that manipulated perceived accuracy (via feedback) and objective estimation difficulty (via signal strength) in a motion reproduction task. We found converging evidence at both the behavioural and computational levels indicating that observers allocate more neural resources towards objects associated with better estimation performance—this, in turn, improves the internal precision of the corresponding object when encountered on the next trial, which further reinforces the association between object identity and reward. This was the case whether performance differences were induced by experimentally manipulated feedback (Experiment 2) or by objective differences in stimulus discriminability (Experiment 3).

We argue that observers derived intrinsic reward from confidence in their responses and feedback on their accuracy. In our tasks, the association of these rewards with the distinguishing feature of the presented objects (that is, colour) leads to a bias in resource allocation, favouring subsequent stimuli that share the same feature. This proposal aligns with the notion that perceptual features linked to rewards are prioritized in sensory processing owing to their incentive salience (for example, refs. 66,67). Moreover, neural evidence supports this notion by demonstrating that sensory representations are modulated by the history of rewards, underscoring the impact of reward associations on perceptual processing⁶⁸. To make our proposal concrete, we developed a mechanistic model grounded in the principles of population coding and RL. Specifically, our RL account operates by analysing accumulated rewards and allocating proportionally more resources to objects previously associated with higher rewards. We found that this model closely replicated resource allocation estimates obtained from freely fitted parameters, suggesting that the history of accumulated intrinsic and external rewards is sufficient to account for the observed patterns of resource allocation.

A key novel finding from the proposed model is that both internally generated and externally manipulated (via feedback) estimates of accuracy, when associated with an object's distinguishing feature, can bias the subsequent processing of objects that share that feature. While the role of external feedback in reinforcing behaviours more generally has been widely acknowledged (for example, refs. 69,70), recent research demonstrates that internal confidence can similarly reinforce behaviour even in the absence of explicit feedback. For instance, improvements in sensitivity in a perceptual learning task have been observed without external feedback⁷¹. Guggenmos et al.⁴² proposed that such learning is driven by confidence prediction errors—discrepancies between an individual's current confidence

and their expected confidence level. Notably, the neural substrate for these prediction errors has been identified in the striatum (see also ref. 72), a brain region traditionally linked to reward processing. Our findings contribute to a growing body of literature that highlights the importance of metacognition⁴⁰ and self-reinforcement³⁸ as critical processes in the pursuit of rewards.

In the present experiments, we did not measure confidence directly; instead, we relied on manipulations that have been shown to selectively alter subjective confidence^{57,58,73,74}, together with computational modelling that tracks internal uncertainty without requiring explicit metacognitive reports^{56,75}. These aspects of our design allowed us to probe confidence-related processes without imposing additional metacognitive demands. In support of this interpretation, false-feedback manipulations have been shown to reliably modulate subjective confidence while leaving accuracy unchanged^{57,73,74}. Given these findings, the effects we observe under our trial-by-trial feedback manipulation are unlikely to be explained solely by expected performance derived from observing feedback, and instead point to changes in internal confidence. Nevertheless, recent work has underscored the importance of more clearly dissociating expected accuracy or success from confidence when examining the mechanisms that guide (meta) cognitive performance. Monetary-incentive paradigms achieve this by manipulating reward valence (that is, gains versus losses)^{76,77}. Future work should build on these observations by explicitly orthogonalizing accuracy and confidence signals across other types of task scenario used here. In addition, although high accuracy in detecting which item provided greater reward prevents us from examining individual differences in allocation as a function of observers' beliefs about reward, future studies could more directly investigate how expectations about differential rewards, feedback and difficulty influence resource allocation. Finally, owing to their relatively low frequency, we did not explicitly model the 180° errors commonly observed in motion perception, whose origin has been discussed in previous works^{59,60}. An important open question is whether these errors in the present setting are indicative of attentional resource allocation and specifically whether they increase when attentional resources are shifted towards another item.

Using a range of tasks similar to the one used in this study, previous research has demonstrated that humans possess knowledge about the uncertainty with which individual items are reported (for example, refs. 58,78,79). Population coding models^{80,81} have been particularly effective in capturing subjective confidence⁷⁵, as well as proxies such as response latency⁸². Within the population coding framework, an ideal observer of spiking activity would derive their confidence estimate—whether internal or explicitly reported—based on the precision of the posterior distribution, which represents the probability of the stimulus value given the observed neural activity. In the neural resource model^{22,56}, the precision of the posterior (or likelihood, assuming a uniform prior over stimulus space) varies from trial to trial, as a result of stochastic variation in the number of spikes available for decoding. We calculated the most probable estimate of posterior precision on each trial to serve as an indicator of internal confidence. On this basis, the model successfully recreated freely estimated resource allocations based on our data.

We estimated the optimal allocation strategy that would maximize external reward and expected accuracy (or equivalently, minimize feedback and response error in Experiments 2 and 3, respectively) and simulated performance under this policy. Across all experiments, the observed allocations systematically deviated from this optimal pattern, and observers' achieved outcomes—number of points, feedback error and recall error—were consistently worse than those predicted under optimal allocation. These observations indicate that participants did not behave in a way that would maximize expected accuracy. For their behaviour to nevertheless be strategic, the resulting policy would have to coincidentally resemble that expected under an intrinsic, implicit reinforcement mechanism, while simultaneously being suboptimal—and even detrimental—with respect to accruing rewards.

In a similar vein, a recent study⁸³ provided theoretical and empirical evidence suggesting that sensory processing is optimized to maximize fitness (that is, rewards), rather than to ensure perceptual accuracy. Supporting this idea, neurophysiological studies have demonstrated that early sensory systems encode both sensory information about a stimulus and non-sensory information regarding the behavioural relevance of stimuli^{3,84}. Embedding stimulus–reward contingencies within the sensory representation of a stimulus facilitates the prioritization of behaviourally relevant information during encoding. These previous findings may help explain why our observers' allocation strategies were not optimized for accuracy in the task; however, they were also not optimized for maximizing rewards. In the experiment with external rewards, we found that observers allocated resource more equally between items than would be predicted by a reward-maximizing strategy. The RL model captured the observed allocation strategy based on a weighted combination of points-based external rewards and confidence-based intrinsic rewards—this combination of factors could lead observers to maintain a certain level of performance even for stimuli associated with low external reward. When considered across all experiments, our results point to a reward-driven allocation of resources that, while prioritizing reward-related stimuli, is not optimized to obtain rewards in the specific tasks we investigated.

In the broader literature on confidence and intrinsically motivated behaviour, it is well established that information seeking (that is, curiosity) tends to peak at intermediate levels of confidence⁸⁵. This reflects a general tendency for people to disengage from tasks they are either unlikely to succeed at or highly confident they can solve. Similar effects have been observed in attentional allocation: for example, infants preferentially attend to events with intermediate levels of information^{86,87}. Delayed reproduction of analogue features, such as motion direction in the present study, is a challenging task even with low-noise high-contrast stimuli. However, we cannot rule out the possibility that in a simpler task where the response associated with one item could be made sufficiently trivial, a similar disengagement of attention and reallocation to the more challenging item might be observed. An important question for future studies will be how intrinsic reward signals, including confidence, influence resource allocation within the broader range of intrinsically motivated behaviour.

Our results also contribute to prominent theories in neuroscience, psychology and economics^{88–91}, which consider how humans and other animals link the mental effort required for a task with the value of its outcome (that is, the reward). Behavioural studies demonstrate that, when faced with tasks offering equal rewards but varying in effort, humans tend to avoid those perceived as more difficult^{92,93}. On the basis of this, it has been argued that cognitive effort is experienced as carrying disutility, that is, acting as a discount factor on expected rewards^{91,94}. This hypothesis has been substantiated by the observation that cognitive effort reduces neural responses to rewards following an effortful task⁹⁵. In the present results, perceived (Experiment 2) or objective difficulty in estimation (Experiment 3) similarly appears to have discounted or reduced the subjective value of a stimulus, leading observers to prioritize 'easier'—and thus, in principle, more 'rewarding'—items for encoding. However, because observers had no control over which stimuli were selected for test, this allocation strategy did not result in more reward in our tasks and could even be counterproductive. This raises the wider question of whether humans may similarly allocate effort suboptimally, driven by intrinsic reward, in other situations where they have limited control over what information will subsequently become relevant.

Methods

Apparatus

In the online experiments, tasks were presented via web browsers on observers' personal computers and were coded in JavaScript and HTML Canvas. In the laboratory experiment, stimuli were displayed

on a 69 cm gamma-corrected LCD monitor with a refresh rate of 60 Hz. Observers were seated in a dark room and viewed the monitor from a distance of 60 cm, with their heads supported by a forehead and chin rest. Eye position was monitored online at 1,000 Hz using an infrared eye tracker (SR Research). Stimulus presentation and response registration were controlled by a script written in Psychtoolbox (v3.0.18.13)^{96,97} and executed in Matlab R2023b (The MathWorks). Responses were collected using a computer mouse.

Participants

A total of 215 naive observers (118 women, 91 men, 6 preferred not to say; $M_{\text{age}} = 27.7$, standard deviation (SD)_{age} = 5.0) took part in the study after giving written informed consent in accordance with the Declaration of Helsinki. All observers reported normal colour vision and normal or corrected-to-normal visual acuity. For the online experiments, observers were recruited using Prolific (<https://www.prolific.co>) and were remunerated £6 h⁻¹ for their participation. Procedures were approved by the University of Cambridge Psychology Research Ethics Committee. For the laboratory experiments, observers were recruited through the Cambridge Psychology research sign-up system and were remunerated £10 h⁻¹. Observers in Experiment 1 received a bonus payment proportional to the number of collected points. By contrast, observers in Experiments 2 and 3 did not receive any additional bonus payments. Our rationale was to avoid introducing external reward contingencies into experiments that were intended to be compared against Experiment 1.

For the online experiments, we used a Bayesian stopping rule to determine the sample size. The stopping rule guides when enough evidence has been gathered to support a decision, thus optimizing the sample size. In particular, we continued testing observers until we obtained strong evidence, as estimated by the Bayes factor, in favour of either H_0 ($\text{BF}_{10} \leq 0.1$, indicating evidence supporting no difference between the 2 conditions of interest) or H_1 ($\text{BF}_{10} \leq 10$, indicating evidence supporting a difference between the 2 conditions). If neither hypothesis was supported, data collection ceased after reaching 100 observers. In Experiment 1, we assessed differences in mean absolute reproduction error in the analogue report task between stimuli associated with high and low reward, which were the conditions of interest for the Bayesian stopping rule. In Experiment 2, we tested for differences in mean absolute reproduction error between error-minified and error-magnified stimuli (Experiments 2a and 2b), and additionally included veridical items in Experiment 2c. In Experiment 3, we compared mean absolute reproduction errors on trials where stimuli were presented in different colours but with equal coherence. For the laboratory experiment (Experiment 3b), we aimed to collect a number of participants similar to that in Experiment 3a. In total, 30 observers participated in Experiment 1 (18 women, 12 men, $M_{\text{age}} = 27$). Twenty-five observers participated in Experiment 2a (17 women, 7 men, 1 preferred not to say, $M_{\text{age}} = 27$), 100 participated in Experiment 2b (50 women, 45 men, 5 preferred not to say, $M_{\text{age}} = 28.5$), and 19 participated in Experiment 2c (8 women, 11 men, $M_{\text{age}} = 28.7$). Finally, 22 (12 women, 10 men, $M_{\text{age}} = 27.8$) and 19 (13 women, 6 men, $M_{\text{age}} = 24.2$) observers participated in Experiments 3a and 3b, respectively.

Stimuli

The stimuli in this study were RDK. On each trial, two RDK stimuli, each consisting of 40 dots, were presented within a circular aperture. A percentage of the dots (specified below) moved in a coherent direction, while the remaining dots moved in random but consistent directions within the aperture⁹⁸. When a dot exited the aperture, it was replaced by a new dot at the aperture's edge, maintaining a constant dot density. In all experiments, 1 stimulus was always green (RGB colour values; online: 47, 195, 129; lab: 0, 199, 128) and the other was always blue (online: 24, 199, 233; lab: 0, 187, 241). In Experiments 2c and 3b, the same observers completed two tasks (see below). In these experiments, stimuli were either green

and blue or orange (237, 154, 0) and magenta (255, 79, 208), balanced across observers and presentation conditions. Across all tasks, stimuli were presented against a mid-grey background. The motion direction for each stimulus was drawn randomly and independently on every trial.

For the online experiments, all measures in pixels are reported for a 1,920 × 1,080 resolution and 60 Hz refresh rate. When a different resolution or refresh rate was detected, all measurements of size, positioning and speed were automatically adjusted to maintain consistency in stimuli presentation across different display settings. The stimulus aperture was 105 pixels in diameter, and each dot had a radius of 3 pixels. Two apertures were positioned 220 pixels to the left or right of the screen centre. On each frame, the dots were shifted by 3 pixels in a specific direction. In the laboratory experiment, two apertures (1.4 dva radius) were presented horizontally aligned with the fixation annulus, positioned at 5 dva to the left and right. Each dot was 0.15 dva in diameter and travelled at 4 dva s⁻¹ speed.

Procedure and task

In all experiments, observers completed an analogue report task⁵⁵. Each trial began with the presentation of a central fixation annulus. In the laboratory experiment, gaze direction was monitored using an eye-tracking camera, and observers were required to maintain gaze fixation within a radius of 2° around the central annulus for 500 ms before the trial could proceed. After achieving stable fixation, the fixation annulus changed appearance (that is, became thinner) to signal that the memory array would be presented in 500 ms. In the online experiment, the appearance of the fixation annulus changed after a fixed interval of 500 ms. The sample array, consisting of two RDK stimuli, was then shown for 750 ms, followed by a 1,000 ms delay period. A centrally presented colour cue subsequently indicated which of the previously presented stimuli, distinguished by colour, was the target that the observers should recall and report the direction of.

Once observers were ready to give their response, they could begin moving the cursor with a mouse or trackpad, which triggered the appearance of a randomly oriented white arrow within the central annulus. Observers were instructed to align the direction of the arrow with the previously presented motion direction of the cued stimulus. In the online experiment, responses were confirmed by pressing the spacebar, while in the laboratory experiment, they were confirmed by pressing the right mouse button. In all experiments, trials from different experimental conditions were randomly interleaved.

Experiment 1: external reward. In Experiment 1, we investigated how external rewards influence motion reproduction precision. To this end, observers received 15 points for reporting a motion direction within 50° of the target direction when the target was of 1 colour (for example, green), and 5 points when it was of the other colour (for example, blue). Responses that were more than 50° from the target direction did not receive any points. The colour associated with high versus low reward was chosen randomly for each observer at the beginning of the experiment. Both stimuli were presented with the same coherence (85%), and no error feedback was provided. Accumulated points were converted to a bonus payment at the end of the experiment, and observers were informed of this at the beginning of the experiment. Overall, they could collect a maximum of 1,000 points, which was equivalent to a bonus payment of £1.50. Observers completed 20 practice trials and 100 experimental trials. The trials were divided into 2 equal blocks with a break of at least 1 min in between, and the complete testing session lasted approximately 15 min.

Experiment 2: perceived accuracy. Experiments 2a, 2b and 2c were designed to investigate the role of feedback on the precision of motion reproduction. In all experiments, at the end of each trial, following the response, we presented feedback showing the reported and target motion directions. Unbeknownst to participants, we manipulated

the feedback by artificially minifying or magnifying errors for one stimulus colour. All three experiments were identical except for the following differences. In Experiments 2a and 2c, two stimuli were presented simultaneously for 750 ms at two distinct locations, whereas in Experiment 2b, the stimuli were presented sequentially at the same two locations, each for 750 ms. In Experiment 2b, the order of presentation and the colour cues were balanced across conditions. In Experiments 2a and 2b, feedback for one item was minified and feedback for the other item was magnified. In Experiment 2c, observers completed two tasks on separate days: in one task, feedback for one item was minified while feedback for the other item was veridical; in the other task, feedback for one item was magnified while feedback for the other item was veridical. For Experiment 2c, the order of tasks was randomized, as were the colour pairs associated with the manipulations: one task used a green–blue pair, and the other used an orange–magenta pair.

Feedback magnification was done by shifting the presented target motion direction (θ^*) away from the reported direction ($\hat{\theta}$) and thereby inflating the presented response error for the designated ‘difficult’ item. This was done according to the following equation:

$$\theta^* = \theta \pm 50 \sin(\hat{\theta} - \theta), \quad (1)$$

where θ is the true motion direction and all angles are expressed in degrees. Similarly, we systematically minimized the error in the feedback for the other colour, designated as the ‘easy’ item. The magnification and minimization of errors were randomly assigned to one of the two colours (that is, green or blue, and orange or magenta) for each observer at the beginning of the experiment. The RDK stimuli were presented with 85% coherence. At the beginning of the experiment, during the instructions, we informed observers that individuals might vary in their ability to perceive the motion of stimuli of different colours. This was intended to make any perceived differences in difficulty appear plausible. At the end of the experiment, observers were debriefed and the true purpose of the study was revealed. In all 3 experiments, observers completed 12 practice trials and 100 experimental trials. The trials were divided into 2 equal blocks with a break of at least 1 min in between, and the complete testing session lasted approximately 15 min. At the very end of the experiment, observers were asked whether the feedback they received was consistent with their performance and, if not, to elaborate. No participants provided responses indicating that they suspected the feedback was manipulated.

Experiment 3: estimation difficulty. In Experiments 3a and 3b, we investigated the role of stimulus discriminability on the fidelity of visual representations. To achieve this, we presented two stimuli with different levels of coherence on 67% and 70% of all trials in Experiments 3a and 3b, respectively. Specifically, the stimulus of 1 colour was presented with 85% (high) and the stimulus of the other colour with 45% (low) coherence. These variable-coherence trials were randomly interleaved with trials where both stimuli had the same intermediate (65%) coherence. The assignment of low and high coherence to specific colours was randomized for each observer at the beginning of the experiments. No feedback was provided during these experiments.

Experiment 3a was conducted online, while Experiment 3b took place in the laboratory. In Experiment 3a, on all trials stimuli were presented simultaneously. In Experiment 3b, the same observers performed the task with both simultaneous and sequential presentations, with the order of these conditions counterbalanced across participants. To prevent transfer effects between conditions, we used different colour combinations: in one condition, stimuli were presented in green and blue, while in the other, they were presented in orange and magenta. The colour combinations were randomly assigned to each presentation condition.

In Experiment 3a, observers completed 20 practice trials and 150 experimental trials. The trials were divided into 2 blocks with a

mandatory break of at least 1 min in between, resulting in a total testing session duration of around 20 min. Experiment 3b (that is, the laboratory experiment) consisted of 400 trials, divided into 8 equal blocks. In half of the blocks, stimuli were presented simultaneously, while in the other half, they were presented sequentially. Half of the observers completed the simultaneous blocks first, followed by the sequential blocks, and vice versa for the other half. At the beginning of each block sequence (that is, simultaneous or sequential task), observers performed 20 practice trials to familiarize themselves with the task. In Experiment 3b, observers were required to maintain central fixation throughout the stimulus presentation. If gaze deviated by more than 2°, a warning message appeared on the screen, and the trial was aborted and restarted with newly randomized stimuli. Completing Experiment 3b took approximately 90 min.

Analysis

All stimulus values were analysed and are reported with respect to the circular parameter space of possible motion directions, $[-\pi, \pi]$ radians. Response error for each trial was measured as the angular difference between the reported and target motion directions. To quantify the dispersion of response errors, we calculated the MAD across trials for each condition and observer. Higher MAD values indicate greater average reproduction error. We report all analyses with MAD, but note that we found the same pattern of differences across experiments with other measures of response variability (circular standard deviation and variance, and cosine dissimilarity).

To compare differences in performance across conditions, we used Bayesian hypothesis tests, implemented in JASP (version 0.18.3)⁹⁹ with the default Jeffreys–Zellner–Siow prior on effect sizes¹⁰⁰. We confirmed that our conclusions were robust across a range of plausible prior specifications. We report Bayes factors which compare the relative predictive adequacy of two competing hypotheses (for example, alternative and null) and quantify the change in belief that the data bring about for the hypotheses under consideration¹⁰¹. For example, $BF_{10} = 10$ indicates that the data are 10 times more likely to occur under the alternative hypothesis (that is, there is a difference) than under the null hypothesis (that is, there is no difference). Evidence for the null hypothesis is indicated by $BF_{10} < 1$, in which case the strength of evidence is indicated by $1/BF_{10}$. Evidence assessed via the Bayes factor is best understood as a ratio-scaled value ranging from 0 to infinity. For clarity in communication, we also use an interpretative framework for Bayes factor values, following the classification scheme outlined by Lee and Wagenmakers¹⁰²: $BF = 1$ as no evidence; $1 < BF < 3$ as weak or anecdotal evidence; $3 \leq BF < 10$ as moderate evidence; $10 \leq BF < 30$ as strong evidence; $30 \leq BF < 100$ as very strong evidence; and $BF \geq 100$ as extreme evidence. It is critical to note that while we utilize these discrete categories, they are arbitrary and should serve only as rough guidelines. Along with the Bayes factor, we report the median of the posterior distribution over the effect size (δ) and the accompanying 95% CI.

Neural resource model

We analysed observers’ response errors with an established model based on the principles of population coding^{22,56,103,104} (Fig. 1f). In this framework, a visual stimulus (θ) is encoded by an idealized population of neurons whose activity is determined by their individual tuning functions. All neurons are assumed to share the same bell-shaped von Mises tuning function,

$$f_i(\theta) = \exp(\kappa(\cos(\theta \ominus \varphi_i) - 1)), \quad (2)$$

where κ determines the tuning concentration and $\ominus$ is subtraction on a circle. These tuning functions are translated through the feature space to peak at each neuron’s preferred value (φ_i), such that they provide dense uniform coverage of the entire feature space. In a population of

M neurons encoding N visual stimuli, the average response of the i th neuron in response to a stimulus value θ is

$$\bar{n}_i(\theta) = \frac{\gamma}{MN} f_i(\theta), \quad (3)$$

where γ is the population's mean total firing activity. This can be simplified by defining $\xi = \frac{\gamma}{N}$ as the mean firing activity encoding a specific stimulus θ . If activity associated with multiple stimuli is combined or normalized¹⁰⁵ at a population level γ , equation (3) implements a form of limited resource²². The spike count produced by each neuron is drawn from a Poisson distribution

$$n_i(\theta) \sim \text{Poiss}(\bar{n}_i(\theta)), \quad (4)$$

and the decoded motion direction estimate is obtained by ML estimation of the population spiking activity $\mathbf{n}$

$$\hat{\theta} = \arg \max_{\theta} p(\mathbf{n}|\theta). \quad (5)$$

The resulting distribution of decoding errors, for a given total number of spikes $m = \sum_i n_i \sim \text{Poiss}(\xi)$, is described as a mixture of von Mises (ϕ) distributions,

$$p(\hat{\theta}|\theta, m) = \int p(r|m, \kappa) \phi(\hat{\theta}; \theta, r\kappa) dr, \quad (6)$$

with

$$p(r|m, \kappa) = \frac{I_0(\kappa r)}{(I_0(\kappa))^m} r \psi_m(r), \quad (7)$$

where $r\psi_m(r)$ is the probability density function for resultant length r of a uniform random walk of m steps. The full distribution of response errors predicted by the model is a mixture of probability distributions $p(\hat{\theta}|\theta, m)$ weighted with the probability of obtaining m spikes. For a complete derivation of the distribution of response errors, see refs. 22,75.

While the neural resource model is rooted in principles of neural coding, its predictions for behaviour have been shown to be mathematically equivalent to estimation based on stochastic sampling⁵⁶. According to this interpretation, an estimate is generated by averaging a set of noisy samples of a stimulus; the number of samples varies randomly around a mean determined by the fraction of resources allocated to it (see ref. 56 for further details).

The model has two free parameters: the population's mean total firing activity (γ) and the concentration of the tuning function (κ). In scenarios when multiple objects (N) need to be represented, the total resource γ is typically divided equally among objects (that is, $\xi = \gamma/N$). Here we extend this basic approach by incorporating an allocation parameter, or gain factor α , which controls the neural activity allocated to one object (see also Bays²²). Without loss of generality, we fixed the gain factor for one object at 1, while treating the gain factor for object j (see below for details of each experiment) as a free parameter when fitting the model to the data. The neural activity allocated to object j can be expressed as $\xi_j = p_\alpha \gamma$, where $p_\alpha = \alpha/(1 + \alpha)$ represents the proportion of total neural activity. The remaining activity (proportion $1 - p_\alpha$) is allocated to the other object.

In Experiment 1, which involved the manipulation of external reward, the allocation weight for the high-reward item was fixed at 1, while the allocation weight for the low-reward item was freely estimated. In Experiment 2, in which we manipulated perceived accuracy, the allocation weight for the error-magnified item was fixed at 1, and the allocation weight for the error-minified item was freely estimated. In Experiment 3, which involved estimation difficulty manipulation,

we simultaneously fitted responses on variable-coherence and equal-coherence trials. Building on our previous work¹⁰⁶, we assumed that the strength of the motion signal is controlled by the coherence level of RDK stimuli, such that the value encoded into the neural population is given by

$$\bar{\theta} \sim \mathcal{WN}(\theta, \sigma_{\text{coherence}}), \quad (8)$$

where $\mathcal{WN}$ is a wrapped normal with mean θ and variance $\sigma_{\text{coherence}}$ accounting for additive Gaussian noise. For simplicity, we considered 85% coherence (high coherence) as perceptually noiseless and further assumed that $\sigma_{45\%} > \sigma_{65\%}$, where 45% was the low-coherence level and 65% was the intermediate-coherence level used in the equal-coherence trials. In addition, the allocation weight for the low-coherence colour (45% and half of 65% stimuli) was fixed at 1, while the allocation weight for the high-coherence colour (85% and half of 65% stimuli) was freely estimated across variable-coherence and equal-coherence trials. In other words, on equal-coherence trials, differences in response precision were explained solely by the allocation weight. By contrast, on variable-coherence trials, the allocation weight and perceptual noise jointly accounted for variations in response precision.

Optimal resource allocation. To identify the optimal levels of resource allocation, we conducted a simulation study. For each observer, we simulated model predictions using the best-fitting parameters of the neural resource model, specifically the population's mean total firing activity (γ), and the concentration of the tuning function (κ), along with a grid of potential allocation weights.

For Experiment 1, we analytically determined the number of points based on the model-predicted response distributions under different allocation weights. The allocation weights were tested across a grid ranging from 0.001 to 2 in increments of 0.01, resulting in 200 distinct values. The optimal allocation weight was identified as the value that maximized the total reward across both high-reward and low-reward items.

For Experiment 2, we numerically simulated the variance (that is, squared circular SD) of feedback errors using the same grid of allocation weights used in Experiment 1. This analysis was based on 10^7 simulated trials drawn from the error distribution predicted by the model. The optimal allocation weight was determined as the value that minimized the total variance of feedback errors across both error-minified and error-magnified items.

For Experiment 3, we analytically modelled the response variance on the variable coherence trials (that is, for the high-coherence and low-coherence items) across a range of allocation weights. We used a grid of 200 values, spanning from 0.01 to 6 in increments of 0.03. The optimal allocation weight was identified as the value that minimized the total response variance for both high-coherence and low-coherence items (that is, 85% and 45% coherence, respectively). In Experiment 3a, the simulation yielded values around $\alpha_{\text{optimal}} = 1$ for all but 1 outlier observer, for whom the estimate reached the endpoint of the examined grid ($\alpha_{\text{optimal}} = 6$). This occurred owing to the model estimating high levels of perceptual noise for medium-coherence and low-coherence stimuli, suggesting that minimizing overall error would be achieved by allocating all resources to the high-coherence object. We exclude this data point in Fig. 4d, and the comparison of observed and optimal allocations is based on the remaining observers. Including this observer's data and performing a non-parametric test did not change our conclusions.

Optimal allocation in this case depends on two additional features of the experiment: first, that the relative frequency of conditions is equal (that is, an equal number of trials across conditions), and second, that observers cannot predict which item will be probed for reproduction (that is, trials were randomly interleaved and probed across conditions).

RL account of resource allocation

We developed a quantitative model to describe how the history of accumulated rewards from multiple objects influences subsequent resource allocation towards those objects. The proposed model extends the neural resource model by incorporating a simple RL rule, which directs behaviour towards more rewarding stimuli. Importantly, our model applies the same RL rule to both external and intrinsic rewards. In the standard RL framework, analysis typically focuses on external rewards, such as points or money, which are provided by the environment as a direct response to the agent's actions. Our model broadens this scope to include intrinsic rewards—those that are inherently pleasurable and drive behaviour—such as the sense of being accurate in a task.

Drawing on the conducted experiments and the motion direction reproduction task, the general overview of this account is as follows: on a particular trial, the received points or money (external reward), perceived accuracy owing to feedback (intrinsic reward), and an individual's internal confidence-based estimate of precision (intrinsic reward) collectively update the value (v) of a particular object associated with these rewards. This computed value influences the allocation of cognitive resources to that object in subsequent encounters, thereby modulating the precision with which the object is represented.

Formally, in the simplest scenario involving only two objects, rather than defining the accumulated reward for each object separately, we can define the relative accumulated reward on trial t as

$$v_t = (1 - \gamma)v_{t-1} + \Delta_t, \quad (9)$$

$$\Delta_t = I_{\pm}(c_1 r_{\text{ext}_t} + c_2 \exp(-|\epsilon_{\text{fb}_t}|) + c_3 \kappa_{r_t}), \quad (10)$$

where γ is a leak component, r_{ext_t} is the number of received points, ϵ_{fb_t} is feedback error, κ_{r_t} is an estimate of internal confidence, and $\{c_1, c_2, c_3\}$ are respective weights accounting for different scales of rewards and the types of reward prioritized by observers. The variable $I_{\pm}$ takes the value of +1 or -1 depending on the object identity, that is, reproduction of the green or blue item, with the assignment of conditions being arbitrary. Positive values of v indicate a higher relative value for 1 item ($I = +1$), while negative values of v indicate a higher relative value for the other item ($I = -1$). To account for the fact that v can range from $-\infty$ to $+\infty$, we transform it into gain parameter α using the following equation:

$$\alpha_{t+1} = e^{2v_t}. \quad (11)$$

This allows us to compute the proportion of spiking activity that will be allocated on a subsequent trial ($t + 1$) to the item identified as $\text{sign}(I) = +1$, given by $\alpha/(1 + \alpha)$, with the remaining spiking activity allocated to the other item. When $v = 0$, such as at the beginning of the task, both items are perceived as having equal value, resulting in an equal distribution of neural resources between them.

The leak component (γ) functions as a temporal filter, modulating the influence of past rewards on resource allocation. When $\gamma = 1$, the system entirely ignores accumulated past values, making the value of an object—and thus resource allocation in the next trial—rely exclusively on the reward from the most recent trial. Conversely, when $\gamma = 0$, the accumulated value is fully retained and integrated with the most recent reward. The necessity of the leak component becomes particularly evident in scenarios where rewards are discontinued: a non-zero leak will gradually equalize the relative value and resource allocation across objects, returning them to a state of equilibrium.

The first reward component of equation (9), r_{ext_t} , reflects the experimental manipulation of Experiment 1. In this experiment, observers received 15 points for responses with an error of less than 50° for high-reward objects and 5 points for low-reward objects. Responses with an error greater than 50° received no points. When applying this model to the data, we used values of $r_{\text{ext}} = \{0.15, 0.05, 0\}$ to represent the rewards for high-reward, low-reward and no-reward trials,

respectively. Note that the utility of rewards, and especially money, is often considered a nonlinear function rather than linear, as assumed here^{107,108}. We used a linear utility function for simplicity and mathematical convenience. Incorporating a parametric non-linear utility function (for example, a power law with exponent β) would introduce an additional free parameter without changing the qualitative pattern of predicted results or the main conclusion: that resource allocation is biased towards items providing higher rewards. Future studies could explore conditions in which non-linear reward functions, reflecting the concave utility of money, might prove critical for explaining behaviour in similar psychophysical tasks.

The feedback component of the model (ϵ_{fb}) addresses the experimental manipulation of Experiment 2. In this experiment, we systematically manipulated feedback error by reducing it for one stimulus and increasing it for another. We hypothesized that feedback serves as an intrinsic reward, with stimuli receiving minified feedback errors being perceived as more rewarding than those with magnified feedback errors. In modelling this relationship, feedback error was assumed to be exponentially related to the object's value v , with smaller feedback errors corresponding to higher rewards, leading to a greater increase in the object's value. This exponential relationship reflects diminishing sensitivity to large feedback errors, such that a wide range of larger errors yields relatively minimal and similar rewards, whereas a narrow range of smaller errors results in considerably higher but more variable rewards.

The final component of our model is the estimate of internal confidence (κ_r). While internal confidence can be assessed through self-reported or metacognitive measures, our approach leverages the inherent mechanism of the neural resource model to quantify uncertainty in the decoded (that is, reported) value. Our approach relies on the principle that the width of the likelihood function reflects the uncertainty of the estimate. The likelihood function evaluates how well various stimulus values align with the observed neural activity: a broad likelihood function is compatible with many different feature values, suggesting lower precision in the ML estimate (the peak of the likelihood function), whereas a narrow likelihood function implies a more precise estimate. Owing to the probabilistic generation of spikes across retrievals (equation (4)), the likelihood has the form of a von Mises with concentration κ_r proportional to the resultant vector length of the preferred values associated with each of the emitted spikes (m), with higher spike counts producing a narrower likelihood function on average³⁶. This formulation has previously been shown to quantitatively reproduce findings from studies in which participants were asked to rate their subjective confidence in each estimate^{75,78}. Consequently, the precision of the likelihood function emerges as a natural candidate for a computational estimate of the observer's internal confidence.

To measure internal confidence associated with each response, we determined the most probable resultant vector length given the individual response errors and the probabilistic distribution of spike values, which was fully characterized by ξ and κ . Specifically, for each trial, we used Bayes' rule to find the posterior probability of resultant vector length r given the error on that trial ϵ , marginalizing over total spike count m ,

$$p(r|\epsilon, \kappa, \xi) = \frac{p(\epsilon|r, \kappa)p(r|\kappa, \xi)}{\int p(\epsilon|r, \kappa)p(r|\kappa, \xi)dr} \quad (12)$$

$$p(\epsilon|r, \kappa) = \phi(\epsilon; 0, \kappa r) \quad (13)$$

$$p(r|m, \xi) = \int p(r|m, \kappa)p(m|\xi)dm, \quad (14)$$

where $p(r|m, \kappa)$ is given by equation (7) and $p(m|\xi)$ is the Poisson probability mass function with mean ξ . Applying MAP estimation to this posterior distribution returns the most probable estimate of resultant length for a given response error,

$$\hat{r} = \arg \max_r p(r|\epsilon, \kappa, \xi). \quad (15)$$

We then use $\kappa_{\hat{r}} = \hat{r}\kappa$ as a measure of internal confidence on the given trial. For experiments with variable estimation difficulty (Experiment 3), the probability density function $p(\epsilon|r, \kappa)$ was obtained by convolving the von Mises probability density function (the right-hand side of equation (13)) with a wrapped normal distribution with freely fitted standard deviation. Finally, $\kappa_{\hat{r}}$ was then calculated as the concentration parameter of a von Mises distribution whose resultant length matches that of $p(\epsilon|\hat{r}, \kappa)$.

Model fitting

To model the observed allocation within the neural resource model^{22,56}, which has two free parameters—the mean population activity (y) and the precision of the tuning functions (κ)—we introduced an additional parameter, the gain modulation α (ref. 22), resulting in a total of three free parameters in Experiments 1 and 2. In Experiment 3, which involved an estimation difficulty manipulation, the neural resource model was extended by two additional parameters ($\sigma_{45\%}$ and $\sigma_{65\%}$) to capture the effects of variable sensory noise introduced by different coherence levels. This brought the total number of free parameters in the neural resource model for Experiment 3 to five.

The RL account retained all parameters of the neural resource model except the gain modulation parameter α , while introducing four new parameters, namely the leak parameter (y) and reward weight parameters (c_1, c_2, c_3). In all three experiments, we modelled the leak parameter (y) and the effect of internal confidence (c_3) on resource allocation (equation (9)); in addition, we modelled the effect of external reward (c_1) only in Experiment 1 while setting it to zero in all other experiments, and feedback error (c_2) only in Experiment 2 while also setting it to zero in all other experiments. This resulted in the estimation of five free parameters in Experiments 1 and 2, and six in Experiment 3. When fitting the model to the data, the leak parameter was constrained between 0 and 1, and all three weight parameters were limited to a range of -1 to 1 .

For all models, we obtained a separate ML fit for each individual observer. These fits were derived using the Nelder–Mead simplex method (via the *fminsearch* function in MATLAB). A MATLAB toolbox implementing the neural resource model is available for download from <https://bayslab.com/toolbox>.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

All data are publicly available at <https://doi.org/10.17605/OSF.IO/EHKQA>.

Code availability

The analysis code is publicly available at <https://doi.org/10.17605/OSF.IO/EHKQA>.

References

- Berridge, K. C., Robinson, T. E. & Aldridge, J. W. Dissecting components of reward: ‘liking’, ‘wanting’, and learning. *Curr. Opin. Pharmacol.* **9**, 65–73 <https://doi.org/10.1016/j.coph.2008.12.014> (2009).
- Berridge, K. C. & Robinson, T. E. Parsing reward. *Trends Neurosci.* **26**, 507–513 [https://doi.org/10.1016/S0166-2236\(03\)00233-9](https://doi.org/10.1016/S0166-2236(03)00233-9) (2003).
- Stănişor, L., Van Der Togt, C., Pennartz, C. M. A. & Roelfsema, P. R. A unified selection signal for attention and reward in primary visual cortex. *Proc. Natl Acad. Sci. USA* **110**, 9136–9141 <https://doi.org/10.1073/pnas.1300117110> (2013).
- Schultz, W. Behavioral theories and the neurophysiology of reward. *Annu. Rev. Psychol.* **57**, 87–115 <https://doi.org/10.1146/annurev.psych.56.091103.070229> (2006).
- Maunsell, J. H. R. Neuronal representations of cognitive state: reward or attention? *Trends Cogn. Sci.* **8**, 261–265 <https://doi.org/10.1016/j.tics.2004.04.003> (2004).
- Blain, B. & Sharot, T. Intrinsic reward: potential cognitive and neural mechanisms. *Curr. Opin. Behav. Sci.* **39**, 113–118 <https://doi.org/10.1016/j.cobeha.2021.03.008> (2021).
- Navalpakkam, V., Koch, C., Rangel, A. & Perona, P. Optimal reward harvesting in complex perceptual environments. *Proc. Natl Acad. Sci. USA* **107**, 5232–5237 <https://doi.org/10.1073/pnas.091972107> (2010).
- Lee, J. & Shomstein, S. The differential effects of reward on space- and object-based attentional allocation. *J. Neurosci.* **33**, 10625–10633 <https://doi.org/10.1523/JNEUROSCI.5575-12.2013> (2013).
- Libera, C. D. & Chelazzi, L. Visual selective attention and the effects of monetary rewards. *Psychol. Sci.* **17**, 222–227 <https://doi.org/10.1111/j.1467-9280.2006.01689.x> (2006).
- Kristjansson, A., Sigurjonsdottir, O. & Driver, J. Fortune and reversals of fortune in visual search: reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Atten. Percept. Psychophys.* **72**, 1229–1236 <https://doi.org/10.3758/APP.72.5.1229> (2010).
- Anderson, B. A., Laurent, P. A. & Yantis, S. Value-driven attentional capture. *Proc. Natl Acad. Sci. USA* **108**, 10367–10371 <https://doi.org/10.1073/pnas.1104047108> (2011).
- Libera, C. D. & Chelazzi, L. Learning to attend and to ignore is a matter of gains and losses. *Psychol. Sci.* **20**, 778–784 <https://doi.org/10.1111/j.1467-9280.2009.02360.x> (2009).
- Peck, C. J., Jangraw, D. C., Suzuki, M., Efem, R. & Gottlieb, J. Reward modulates attention independently of action value in posterior parietal cortex. *J. Neurosci.* **29**, 11182–11191 <https://doi.org/10.1523/JNEUROSCI.1929-09.2009> (2009).
- Theeuwes, J. & Belopolsky, A. V. Reward grabs the eye: oculomotor capture by rewarding stimuli. *Vis. Res.* **74**, 80–85 <https://doi.org/10.1016/j.visres.2012.07.024> (2012).
- Anderson, B. A. & Kim, H. Mechanisms of value-learning in the guidance of spatial attention. *Cognition* **178**, 26–36 <https://doi.org/10.1016/j.cognition.2018.05.005> (2018).
- Anderson, B. A. & Kim, H. On the representational nature of value-driven spatial attentional biases. *J. Neurophysiol.* **120**, 2654–2658 <https://doi.org/10.1152/jn.00489.2018> (2018).
- Awh, E., Belopolsky, A. V. & Theeuwes, J. Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends Cogn. Sci.* **16**, 437–443 <https://doi.org/10.1016/j.tics.2012.06.010> (2012).
- Anderson, B. A. et al. The past, present, and future of selection history. *Neurosci. Biobehav. Rev.* **130**, 326–350 <https://doi.org/10.1016/j.neubiorev.2021.09.004> (2021).
- Bays, P. M., Schneegans, S., Ma, W. J. & Brady, T. F. Representation and computation in visual working memory. *Nat. Hum. Behav.* **8**, 1016–1034 <https://doi.org/10.1038/s41562-024-01871-2> (2024).
- Bays, P. M., Gorgoraptis, N., Wee, N., Marshall, L. & Husain, M. Temporal dynamics of encoding, storage, and reallocation of visual working memory. *J. Vis.* **11**, 6 <https://doi.org/10.1167/11.10.6> (2011).
- Gorgoraptis, N., Catalao, R. F. G., Bays, P. M. & Husain, M. Dynamic updating of working memory resources for visual objects. *J. Neurosci.* **31**, 8502–8511 <https://doi.org/10.1523/JNEUROSCI.0208-11.2011> (2011).
- Bays, P. M. Noise in neural populations accounts for errors in working memory. *J. Neurosci.* **34**, 3632–3645 <https://doi.org/10.1523/JNEUROSCI.3204-13.2014> (2014).

23. Emrich, S. M., Lockhart, H. A. & Al-Aidroos, N. Attention mediates the flexible allocation of visual working memory resources. *J. Exp. Psychol. Hum. Percept. Perform.* **43**, 1454–1465 <https://doi.org/10.1037/xhp0000398> (2017).
24. Sprague, T. C., Itthipuripat, S., Vo, V. Y. A. & Serences, J. T. Dissociable signatures of visual salience and behavioral relevance across attentional priority maps in human cortex. *J. Neurophysiol.* **119**, 2153–2165 <https://doi.org/10.1152/jn.00059.2018> (2018).
25. Yoo, A. H., Klyszejko, Z., Curtis, C. E. & Ma, W. J. Strategic allocation of working memory resource. *Sci. Rep.* **8**, 16162 <https://doi.org/10.1101/329870> (2018).
26. Taylor, R., Tomić, I., Aagten-Murphy, D. & Bays, P. M. Working memory is updated by reallocation of resources from obsolete to new items. *Atten. Percept. Psychophys.* **85**, 1437–1451 <https://doi.org/10.3758/s13414-022-02584-2> (2023).
27. Griffin, I. C. & Nobre, A. C. Orienting attention to locations in internal representations. *J. Cogn. Neurosci.* **15**, 1176–1194 <https://doi.org/10.1162/089892903322598139> (2003).
28. Oberauer, K. Control of the contents of working memory—a comparison of two paradigms and two age groups. *J. Exp. Psychol. Learn. Mem. Cogn.* **31**, 714–728 <https://doi.org/10.1037/0278-7393.31.4.714> (2005).
29. Klyszejko, Z., Rahmati, M. & Curtis, C. E. Attentional priority determines working memory precision. *Vis. Res.* **105**, 70–76 <https://doi.org/10.1016/j.visres.2014.09.002> (2014).
30. Atkinson, A. L., Oberauer, K., Allen, R. J. & Souza, A. S. Why does the probe value effect emerge in working memory? Examining the biased attentional refreshing account. *Psychon. Bull. Rev.* **29**, 891–900 <https://doi.org/10.3758/s13423-022-02056-6> (2022).
31. Allen, R. J., Atkinson, A. & Hitch, G. J. Getting value out of working memory through strategic prioritisation: implications for storage and control. *Q. J. Exp. Psychol.* **78**, 405–424 <https://doi.org/10.1177/17470218241258102> (2024).
32. Gong, M. & Li, S. Learned reward association improves visual working memory. *J. Exp. Psychol. Hum. Percept. Perform.* **40**, 841–856 <https://doi.org/10.1037/a0035131> (2014).
33. Brissenden, J. A., Adkins, T. J., Hsu, Y. T. & Lee, T. G. Reward influences the allocation but not the availability of resources in visual working memory. *J. Exp. Psychol. Gen.* **152**, 1825–1839 <https://doi.org/10.1037/xge0001370> (2023).
34. Van Den Berg, R., Zou, Q., Li, Y. & Ma, W. J. No effect of monetary reward in a visual working memory task. *PLoS ONE* **18**, e0280257 <https://doi.org/10.1371/journal.pone.0280257> (2023).
35. Atkinson, A. L. et al. Are there multiple ways to direct attention in working memory? *Ann. N. Y. Acad. Sci.* **1424**, 115–126 <https://doi.org/10.1111/nyas.13634> (2018).
36. Gazzaley, A. & Nobre, A. C. Top-down modulation: bridging selective attention and working memory. *Trends Cogn. Sci.* **16**, 129–135 <https://doi.org/10.1016/j.tics.2011.11.014> (2012).
37. Awh, E., Vogel, E. K. & Oh, S.-H. Interactions between attention and working memory. *Neuroscience* **139**, 201–208 <https://doi.org/10.1016/j.neuroscience.2005.08.023> (2006).
38. Bröker, F., Holt, L. L., Roads, B. D., Dayan, P. & Love, B. C. Demystifying unsupervised learning: how it helps and hurts. *Trends Cogn. Sci.* **28**, 974–986 <https://doi.org/10.1016/j.tics.2024.09.005> (2024).
39. Cortese, A. Metacognitive resources for adaptive learning. *Neurosci. Res.* **178**, 10–19 <https://doi.org/10.1016/j.neures.2021.09.003> (2022).
40. Fleming, S. M. Metacognition and confidence: a review and synthesis. *Annu. Rev. Psychol.* **75**, 241–268 <https://doi.org/10.1146/annurev-psych-022423-032425> (2024).
41. Daniel, R. & Pollmann, S. Striatal activations signal prediction errors on confidence in the absence of external feedback. *NeuroImage* **59**, 3457–3467 <https://doi.org/10.1016/j.neuroimage.2011.11.058> (2012).
42. Guggenmos, M., Wilbertz, G., Hebart, M. N. & Sterzer, P. Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *Elife* **5**, e13388 <https://doi.org/10.7554/eLife.13388> (2016).
43. Schultz, W., Dayan, P. & Read Montague, P. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 <https://doi.org/10.1126/science.275.5306.1593> (1997).
44. O'Doherty, J. P., Buchanan, T. W., Seymour, B. & Dolan, R. J. Predictive neural coding of reward preference involves dissociable responses in human ventral midbrain and ventral striatum. *Neuron* **49**, 157–166 <https://doi.org/10.1016/j.neuron.2005.11.014> (2006).
45. Hebart, M. N., Schriever, Y., Donner, T. H. & Haynes, J.-D. The relationship between perceptual decision variables and confidence in the human brain. *Cereb. Cortex* **26**, 118–130 <https://doi.org/10.1093/cercor/bhu181> (2016).
46. Ding, L. & Gold, J. I. Caudate encodes multiple computations for perceptual decisions. *J. Neurosci.* **30**, 15747–15759 <https://doi.org/10.1523/JNEUROSCI.2894-10.2010> (2010).
47. Balsdon, T., Andrea Pisau, M. & Philastides, M. G. Distinct basal ganglia contributions to learning from implicit and explicit value signals in perceptual decision-making. *Nat. Commun.* **15**, 5317 <https://doi.org/10.1038/s41467-024-49538-w> (2024).
48. Chen, H., Verbeke, P., Mattioni, S., Calderon, C. B. & Verguts, T. Neural and computational evidence for a predictive learning account of the testing effect. *Proc. Natl Acad. Sci. USA* **122**, e2506530122 <https://doi.org/10.1073/pnas.2506530122> (2025).
49. Holroyd, C. B. et al. Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nat. Neurosci.* **7**, 497–498 <https://doi.org/10.1038/nn1238> (2004).
50. Satterthwaite, T. D. et al. Being right is its own reward: load and performance related ventral striatum activation to correct responses during a working memory task in youth. *NeuroImage* **61**, 723–729 <https://doi.org/10.1016/j.neuroimage.2012.03.060> (2012).
51. Barto, A. G. Intrinsic motivation and reinforcement learning. in *Intrinsically Motivated Learning in Natural and Artificial Systems* (eds Baldassarre, G. & Mirolli, M.) 17–47 (Springer, 2013).
52. Chentanez, N., Barto, A. & Singh, S. Intrinsically motivated reinforcement learning. In *Advances in Neural Information Processing Systems 17 (NIPS 2004)* (eds Saul, L. et al.) https://proceedings.neurips.cc/paper_files/paper/2004/file/4be5a36cbaca8ab9d2066debfe4e65c1-Paper.pdf (2004).
53. Singh, S., Lewis, R. L., Barto, A. G. & Sorg, J. Intrinsically motivated reinforcement learning: an evolutionary perspective. *IEEE Trans. Auton. Ment. Dev.* **2**, 70–82 <https://doi.org/10.1109/TAMD.2010.2051031> (2010).
54. Prinzmetal, W., Amiri, H., Allen, K. & Edwards, T. Phenomenology of attention: I. Color, location, orientation, and spatial frequency. *J. Exp. Psychol. Hum. Percept. Perform.* **24**, 261–282 <https://doi.org/10.1037/0096-1523.24.1.261> (1998).
55. Tomić, I., Adamcová, D., Fehér, M. & Bays, P. M. Dissecting the components of error in analogue report tasks. *Behav. Res. Methods* **56**, 8196–8213 <https://doi.org/10.3758/s13428-024-02453-w> (2024).
56. Schneegans, S., Taylor, R. & Bays, P. M. Stochastic sampling provides a unifying account of visual working memory limits. *Proc. Natl Acad. Sci. USA* **117**, 20959–20968 <https://doi.org/10.1073/pnas.2004306117> (2020).
57. Van Marcke, H., Denmat, P. L., Verguts, T. & Desender, K. Manipulating prior beliefs causally induces under- and overconfidence. *Psychol. Sci.* **35**, 358–375 <https://doi.org/10.1177/09567976241231572> (2024).

58. Kiani, R., Corthell, L. & Shadlen, M. N. Choice certainty is informed by both evidence and decision time. *Neuron* **84**, 1329–1342 <https://doi.org/10.1016/j.neuron.2014.12.015> (2014).
59. Chetverikov, A. & Jehee, J. F. M. Motion direction is represented as a bimodal probability distribution in the human visual cortex. *Nat. Commun.* **14**, 7634 <https://doi.org/10.1038/s41467-023-43251-w> (2023).
60. Kwak, Y. & Curtis, C. E. Unveiling the abstract format of mnemonic representations. *Neuron* **110**, 1822–1828.e5 <https://doi.org/10.1016/j.neuron.2022.03.016> (2022).
61. Wolf, D. H. et al. Striatal intrinsic reinforcement signals during recognition memory: relationship to response bias and dysregulation in schizophrenia. *Front. Behav. Neurosci.* **5**, 81 <https://doi.org/10.3389/fnbeh.2011.00081> (2011).
62. Han, S., Huettel, S. A., Raposo, A., Alison Adcock, R. & Dobbins, I. G. Functional significance of striatal responses during episodic decisions: recovery or goal attainment? *J. Neurosci.* **30**, 4767–4775 <https://doi.org/10.1523/JNEUROSCI.3077-09.2010> (2010).
63. Shadmehr, R., Reppert, T. R., Summerside, E. M., Yoon, T. & Ahmed, A. A. Movement vigor as a reflection of subjective economic utility. *Trends Neurosci.* **42**, 323–336 <https://doi.org/10.1016/j.tins.2019.02.003> (2019).
64. Summerside, E. M., Shadmehr, R. & Ahmed, A. A. Vigor of reaching movements: reward discounts the cost of effort. *J. Neurophysiol.* **119**, 2347–2357 <https://doi.org/10.1152/jn.00872.2017> (2018).
65. Manohar, S. G., Finzi, R. D., Drew, D. & Husain, M. Distinct motivational effects of contingent and noncontingent rewards. *Psychol. Sci.* **28**, 1016–1026 <https://doi.org/10.1177/0956797617693326> (2017).
66. Berridge, K. C. & Robinson, T. E. What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res. Rev.* **28**, 309–369 [https://doi.org/10.1016/S0165-0173\(98\)00019-8](https://doi.org/10.1016/S0165-0173(98)00019-8) (1998).
67. Yoo, A. H. & Collins, A. G. E. How working memory and reinforcement learning are intertwined: a cognitive, neural, and computational perspective. *J. Cogn. Neurosci.* **34**, 551–568 https://doi.org/10.1162/jocn_a_01808 (2022).
68. Serences, J. T. Value-based modulations in human visual cortex. *Neuron* **60**, 1169–1181 <https://doi.org/10.1016/j.neuron.2008.10.051> (2008).
69. Gregory Ashby, F. & Todd Maddox, W. Human category learning. *Annu. Rev. Psychol.* **56**, 149–178 <https://doi.org/10.1146/annurev.psych.56.091103.070217> (2005).
70. Hattie, J. & Timperley, H. The power of feedback. *Rev. Educ. Res.* **77**, 81–112 <https://doi.org/10.3102/003465430298487> (2007).
71. Haddara, N. & Rahnev, D. The impact of feedback on perceptual decision-making and metacognition: reduction in bias but no change in sensitivity. *Psychol. Sci.* **33**, 259–275 <https://doi.org/10.1177/09567976211032887> (2022).
72. Rouault, M. & Fleming, S. M. Formation of global self-beliefs in the human brain. *Proc. Natl Acad. Sci. USA* **117**, 27268–27276 <https://doi.org/10.1073/pnas.2003094117> (2020).
73. Denmat, P. L., Desender, K. & Verguts, T. Learning to be confident: how agents learn confidence based on prediction errors. *Cognition* **266**, 106332 <https://doi.org/10.1016/j.cognition.2025.106332> (2026).
74. Denmat, P. L., Verguts, T. & Desender, K. A low-dimensional approximation of optimal confidence. *PLoS Comput. Biol.* **20**, e1012273 <https://doi.org/10.1371/journal.pcbi.1012273> (2024).
75. Bays, P. M. A signature of neural coding at human perceptual limits. *J. Vis.* **16**, 4 <https://doi.org/10.1167/16.11.4> (2016).
76. Bröder, A., Navarro-Báez, S. & Undorf, M. Reducing cheap talk? How monetary incentives affect the accuracy of metamemory judgments. *Mem. Cogn.* **53**, 1725–1737 <https://doi.org/10.3758/s13421-024-01679-5> (2025).
77. Lebreton, M. et al. Two sides of the same coin: monetary incentives concurrently improve and bias confidence judgments. *Sci. Adv.* **4**, eaq0668 <https://doi.org/10.1126/sciadv.aq0668> (2018).
78. van den Berg, R., Yoo, A. H. & Ma, W. J. Fechner's law in metacognition: a quantitative model of visual working memory confidence. *Psychol. Rev.* **124**, 197–214 <https://doi.org/10.1037/rev000060> (2017).
79. Li, H.-H., Sprague, T. C., Yoo, A. H., Ma, W. J. & Curtis, C. E. Joint representation of working memory and uncertainty in human cortex. *Neuron* **109**, 3699–3712.e6 <https://doi.org/10.1016/j.neuron.2021.08.022> (2021).
80. Ma, W. J., Beck, J. M., Latham, P. E. & Pouget, A. Bayesian inference with probabilistic population codes. *Nat. Neurosci.* **9**, 1432–1438 <https://doi.org/10.1038/nn1790> (2006).
81. Pouget, A., Dayan, P. & Zemel, R. S. Inference and computation with population codes. *Annu. Rev. Neurosci.* **26**, 381–410 <https://doi.org/10.1146/annurev.neuro.26.041002.131112> (2003).
82. Schneegans, S. & Bays, P. M. Drift in neural population activity causes working memory to deteriorate over time. *J. Neurosci.* **38**, 4859–4869 <https://doi.org/10.1523/JNEUROSCI.3440-17.2018> (2018).
83. Schaffner, J., Bao, S. D., Tobler, P. N., Hare, T. A. & Polania, R. Sensory perception relies on fitness-maximizing codes. *Nat. Hum. Behav.* **7**, 1135–1151 <https://doi.org/10.1038/s41562-023-01584-y> (2023).
84. Shuler, M. G. & Bear, M. F. Reward timing in the primary visual cortex. *Science* **311**, 1606–1609 <https://doi.org/10.1126/science.1123513> (2006).
85. Ten, A., Oudeyer, P.-Y., Sakaki, M. & Murayama, K. The curious U: integrating theories linking knowledge and information-seeking behavior. *Open Mind* **9**, 1763–1785 <https://doi.org/10.1162/OPMI.a.41> (2025).
86. Kidd, C., Piantadosi, S. T. & Aslin, R. N. The Goldilocks effect: human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS ONE* **7**, e36399 <https://doi.org/10.1371/journal.pone.0036399> (2012).
87. Kidd, C., Piantadosi, S. T. & Aslin, R. N. The Goldilocks effect in infant auditory attention. *Child Dev.* **85**, 1795–1804 <https://doi.org/10.1111/cdev.12263> (2014).
88. Badre, D. Cognitive control. *Annu. Rev. Psychol.* **76**, 167–195 <https://doi.org/10.1146/annurev-psych-022024-103901> (2024).
89. Inzlicht, M., Shenhav, A. & Olivola, C. Y. The effort paradox: effort is both costly and valued. *Trends Cogn. Sci.* **22**, 337–349 <https://doi.org/10.1016/j.tics.2018.01.007> (2018).
90. Westbrook, A. & Braver, T. S. Cognitive effort: a neuroeconomic approach. *Cogn. Affect. Behav. Neurosci.* **15**, 395–415 <https://doi.org/10.3758/s13415-015-0334-y> (2015).
91. Shenhav, A. et al. Toward a rational and mechanistic account of mental effort. *Annu. Rev. Neurosci.* **40**, 99–124 <https://doi.org/10.1146/annurev-neuro-072116-031526> (2017).
92. Kool, W., McGuire, J. T., Rosen, Z. B. & Botvinick, M. M. Decision making and the avoidance of cognitive demand. *J. Exp. Psychol. Gen.* **139**, 665–682 <https://doi.org/10.1037/a0020198> (2010).
93. Corlazzoli, G., Desender, K. & Gevers, W. Feeling and deciding: subjective experiences rather than objective factors drive the decision to invest cognitive control. *Cognition* **240**, 105587 <https://doi.org/10.1016/j.cognition.2023.105587> (2023).
94. Kool, W. & Botvinick, M. The intrinsic cost of cognitive control. *Behav. Brain Sci.* **36**, 697–698 <https://doi.org/10.1017/S0140525X1300109X> (2013).

95. Botvinick, M. M., Huffstetler, S. & McGuire, J. T. Effort discounting in human nucleus accumbens. *Cogn. Affect. Behav. Neurosci.* **9**, 16–27 <https://doi.org/10.3758/CABN.9.1.16> (2009).
96. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 <https://doi.org/10.1163/156856897X00357> (1997).
97. Pelli, D. G. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10**, 437–442 (1997).
98. Scase, M. O., Braddick, O. J. & Raymond, J. E. What is noise for the motion system? *Vis. Res.* **36**, 2579–2586 [https://doi.org/10.1016/0042-6989\(95\)00325-8](https://doi.org/10.1016/0042-6989(95)00325-8) (1996).
99. JASP v.0.18.3 <https://jasp-stats.org/> (JASP Team, 2024).
100. Liang, F., Paulo, R., Molina, G., Clyde, M. A. & Berger, J. O. Mixtures of g priors for Bayesian variable selection. *J. Am. Stat. Assoc.* **103**, 410–423 <https://doi.org/10.1198/016214507000001337> (2008).
101. Wagenmakers, E.-J. et al. Bayesian inference for psychology. Part II: example applications with JASP. *Psychon. Bull. Rev.* **25**, 58–76 <https://doi.org/10.3758/s13423-017-1323-7> (2018).
102. Lee, M. D. & Wagenmakers, E.-J. *Bayesian Cognitive Modeling: A Practical Course* (Cambridge Univ. Press, 2013).
103. Tomić, I. & Bays, P. M. Perceptual similarity judgments do not predict the distribution of errors in working memory. *J. Exp. Psychol. Learn. Mem. Cogn.* **50**, 535–549 <https://doi.org/10.1037/xlm0001172> (2024).
104. Tomić, I. & Bays, P. M. A dynamic neural resource model bridges sensory and working memory. *Elife* **12**, RP91034 <https://doi.org/10.7554/eLife.91034.3> (2024).
105. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 <https://doi.org/10.1038/nrn3136> (2012).
106. Tomić, I. & Bays, P. M. Internal but not external noise frees working memory resources. *PLoS Comput. Biol.* **14**, e1006488 <https://doi.org/10.1371/journal.pcbi.1006488> (2018).
107. Glimcher, P. W. & Rustichini, A. Neuroeconomics: the consilience of brain and decision. *Science* **306**, 447–452 <https://doi.org/10.1126/science.1102566> (2004).
108. Starmer, C. Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *J. Econ. Lit.* **38**, 332–382 <https://doi.org/10.1257/jel.38.2.332> (2000).

Acknowledgements

We thank D. Aagten-Murphy and R. Taylor, who worked on earlier iterations of this project. We thank N. Abraham, P. Alex, A. Anand, P. McMeekin, A. Sabo, T. Wenban-Smith, A. Zhu and A. Triabhall for assisting with data collection. We used resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service.

Author contributions

I.T. contributed to conceptualization, methodology, software, data collection, investigation, formal analysis, modelling, visualizations

and writing—original draft and revisions. R.R. contributed to methodology, software and data collection. P.M.B. contributed to conceptualization, funding acquisition, supervision, methodology, formal analysis, modelling, visualizations and writing—editing and revisions.

Funding

This work was funded by the Wellcome Trust (grant 106926 to P.M.B.). Additional support for I.T. was provided by the Croatian Science Foundation (grant UIP-2025-02-6353). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41562-026-02573-7>.

Correspondence and requests for materials should be addressed to Ivan Tomić or Paul M. Bays.

Peer review information *Nature Human Behaviour* thanks Andrey Chetverikov, Gabrielle Weidemann and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

- | | |
|-------------------------------------|--|
| n/a | Confirmed |
| ☐ | ☒ The exact sample size (<i>n</i>) for each experimental group/condition, given as a discrete number and unit of measurement |
| ☐ | ☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly |
| ☐ | ☒ The statistical test(s) used AND whether they are one- or two-sided
<i>Only common tests should be described solely by name; describe more complex techniques in the Methods section.</i> |
| ☒ | ☐ A description of all covariates tested |
| ☐ | ☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons |
| ☐ | ☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals) |
| ☒ | ☐ For null hypothesis testing, the test statistic (e.g. <i>F</i> , <i>t</i> , <i>r</i>) with confidence intervals, effect sizes, degrees of freedom and <i>P</i> value noted
<i>Give P values as exact values whenever suitable.</i> |
| ☐ | ☒ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings |
| ☒ | ☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes |
| ☐ | ☒ Estimates of effect sizes (e.g. Cohen's <i>d</i> , Pearson's <i>r</i>), indicating how they were calculated |

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection	All experimental tasks were implemented using custom code developed by the authors. MATLAB R2023b (The MathWorks Inc.) was used for desktop-based experiments. Psychophysics Toolbox Version 3 (Psychtoolbox-3, v3.0.18.13) was used in combination with MATLAB for stimulus presentation and response collection. JavaScript was used for browser-based experiments. No commercial or open-source code beyond the listed versions was used in the implementation of the experiments.
Data analysis	All analyses were conducted using a combination of commercial software and custom scripts developed by the authors. MATLAB R2024b (The MathWorks Inc.) was used for data preprocessing, model fitting, and visualisation. JASP version 0.18.3 (JASP Team) was used for conducting statistical analyses. All custom analysis code will be made publicly available upon publication.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

Data and analysis code are publicly available at <https://doi.org/10.17605/OSF.IO/EHKQA>.

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	Two hundred fifteen naive observers (118 females, 91 males, 6 preferred not to say) took part in the study.
Reporting on race, ethnicity, or other socially relevant groupings	n/a
Population characteristics	Two hundred fifteen naive observers (mean and SD age, M = 27.7, SD = 5.0) took part in the study.
Recruitment	For the online experiments, observers were recruited using Prolific (https://www.prolific.co). For the laboratory experiments, observers were recruited through the Cambridge Psychology research sign-up system.
Ethics oversight	Procedures were approved by the University of Cambridge Psychology Research Ethics Committee. All observers gave written informed consent in accordance with the Declaration of Helsinki.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	For the online experiments, we used a Bayesian stopping rule to determine the sample size. For the laboratory experiment (Experiment 3b), we aimed to collect a number of participants similar to that in Experiment 3a.
Data exclusions	In Experiment 3a, data for one observer were excluded from the analysis of optimal resource allocation, as explained in the main text. We excluded these data in Figure 4D, and the comparison of observed and optimal allocations is based on the remaining observers. Including this observer's data and performing a non-parametric test did not change our conclusions. In Experiment 3b, four observers who had difficulty maintaining fixation or could not be successfully calibrated on the eye tracker were replaced.
Replication	Experiment 3a, conducted online, was replicated in a laboratory setting as Experiment 3b.
Randomization	Experimental parameters, including the order of motion directions and the colour-reward associations (i.e., which object was associated with high or low reward), were randomised across observers.
Blinding	Experimenters were not blinded during data collection, as a within-subjects design was used in which all observers completed all experimental conditions.

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Briefly describe the study type including whether data are quantitative, qualitative, or mixed-methods (e.g. qualitative cross-sectional, quantitative experimental, mixed-methods case study).
-------------------	---

Research sample	State the research sample (e.g. Harvard university undergraduates, villagers in rural India) and provide relevant demographic information (e.g. age, sex) and indicate whether the sample is representative. Provide a rationale for the study sample chosen. For studies involving existing datasets, please describe the dataset and source.
Sampling strategy	Describe the sampling procedure (e.g. random, snowball, stratified, convenience). Describe the statistical methods that were used to predetermine sample size OR if no sample-size calculation was performed, describe how sample sizes were chosen and provide a rationale for why these sample sizes are sufficient. For qualitative data, please indicate whether data saturation was considered, and what criteria were used to decide that no further sampling was needed.
Data collection	Provide details about the data collection procedure, including the instruments or devices used to record the data (e.g. pen and paper, computer, eye tracker, video or audio equipment) whether anyone was present besides the participant(s) and the researcher, and whether the researcher was blind to experimental condition and/or the study hypothesis during data collection.
Timing	Indicate the start and stop dates of data collection. If there is a gap between collection periods, state the dates for each sample cohort.
Data exclusions	If no data were excluded from the analyses, state so OR if data were excluded, provide the exact number of exclusions and the rationale behind them, indicating whether exclusion criteria were pre-established.
Non-participation	State how many participants dropped out/declined participation and the reason(s) given OR provide response rate OR state that no participants dropped out/declined participation.
Randomization	If participants were not allocated into experimental groups, state so OR describe how participants were allocated to groups, and if allocation was not random, describe how covariates were controlled.

Ecological, evolutionary & environmental sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Briefly describe the study. For quantitative data include treatment factors and interactions, design structure (e.g. factorial, nested, hierarchical), nature and number of experimental units and replicates.
Research sample	Describe the research sample (e.g. a group of tagged <i>Passer domesticus</i> , all <i>Stenocereus thurberi</i> within Organ Pipe Cactus National Monument), and provide a rationale for the sample choice. When relevant, describe the organism taxa, source, sex, age range and any manipulations. State what population the sample is meant to represent when applicable. For studies involving existing datasets, describe the data and its source.
Sampling strategy	Note the sampling procedure. Describe the statistical methods that were used to predetermine sample size OR if no sample-size calculation was performed, describe how sample sizes were chosen and provide a rationale for why these sample sizes are sufficient.
Data collection	Describe the data collection procedure, including who recorded the data and how.
Timing and spatial scale	Indicate the start and stop dates of data collection, noting the frequency and periodicity of sampling and providing a rationale for these choices. If there is a gap between collection periods, state the dates for each sample cohort. Specify the spatial scale from which the data are taken
Data exclusions	If no data were excluded from the analyses, state so OR if data were excluded, describe the exclusions and the rationale behind them, indicating whether exclusion criteria were pre-established.
Reproducibility	Describe the measures taken to verify the reproducibility of experimental findings. For each experiment, note whether any attempts to repeat the experiment failed OR state that all attempts to repeat the experiment were successful.
Randomization	Describe how samples/organisms/participants were allocated into groups. If allocation was not random, describe how covariates were controlled. If this is not relevant to your study, explain why.
Blinding	Describe the extent of blinding used during data acquisition and analysis. If blinding was not possible, describe why OR explain why blinding was not relevant to your study.

Did the study involve field work? ☐ Yes ☐ No

Field work, collection and transport

Field conditions	Describe the study conditions for field work, providing relevant parameters (e.g. temperature, rainfall).
Location	State the location of the sampling or experiment, providing relevant parameters (e.g. latitude and longitude, elevation, water depth).
Access & import/export	Describe the efforts you have made to access habitats and to collect and import/export your samples in a responsible manner and in

Access & import/export *compliance with local, national and international laws, noting any permits that were obtained (give the name of the issuing authority, the date of issue, and any identifying information).*

Disturbance *Describe any disturbance caused by the study and how it was minimized.*

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

- n/a Involved in the study
- ☒ ☐ Antibodies
- ☒ ☐ Eukaryotic cell lines
- ☒ ☐ Palaeontology and archaeology
- ☒ ☐ Animals and other organisms
- ☒ ☐ Clinical data
- ☒ ☐ Dual use research of concern
- ☒ ☐ Plants

Methods

- n/a Involved in the study
- ☒ ☐ ChIP-seq
- ☒ ☐ Flow cytometry
- ☒ ☐ MRI-based neuroimaging

Antibodies

Antibodies used *Describe all antibodies used in the study; as applicable, provide supplier name, catalog number, clone name, and lot number.*

Validation *Describe the validation of each primary antibody for the species and application, noting any validation statements on the manufacturer's website, relevant citations, antibody profiles in online databases, or data provided in the manuscript.*

Eukaryotic cell lines

Policy information about [cell lines and Sex and Gender in Research](#)

Cell line source(s) *State the source of each cell line used and the sex of all primary cell lines and cells derived from human participants or vertebrate models.*

Authentication *Describe the authentication procedures for each cell line used OR declare that none of the cell lines used were authenticated.*

Mycoplasma contamination *Confirm that all cell lines tested negative for mycoplasma contamination OR describe the results of the testing for mycoplasma contamination OR declare that the cell lines were not tested for mycoplasma contamination.*

Commonly misidentified lines (See [ICLAC](#) register) *Name any commonly misidentified cell lines used in the study and provide a rationale for their use.*

Palaeontology and Archaeology

Specimen provenance *Provide provenance information for specimens and describe permits that were obtained for the work (including the name of the issuing authority, the date of issue, and any identifying information). Permits should encompass collection and, where applicable, export.*

Specimen deposition *Indicate where the specimens have been deposited to permit free access by other researchers.*

Dating methods *If new dates are provided, describe how they were obtained (e.g. collection, storage, sample pretreatment and measurement), where they were obtained (i.e. lab name), the calibration program and the protocol for quality assurance OR state that no new dates are provided.*

☐ Tick this box to confirm that the raw and calibrated dates are available in the paper or in Supplementary Information.

Ethics oversight *Identify the organization(s) that approved or provided guidance on the study protocol, OR state that no ethical approval or guidance was required and explain why not.*

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Animals and other research organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research, and [Sex and Gender in Research](#)

Laboratory animals	<i>For laboratory animals, report species, strain and age OR state that the study did not involve laboratory animals.</i>
Wild animals	<i>Provide details on animals observed in or captured in the field; report species and age where possible. Describe how animals were caught and transported and what happened to captive animals after the study (if killed, explain why and describe method; if released, say where and when) OR state that the study did not involve wild animals.</i>
Reporting on sex	<i>Indicate if findings apply to only one sex; describe whether sex was considered in study design, methods used for assigning sex. Provide data disaggregated for sex where this information has been collected in the source data as appropriate; provide overall numbers in this Reporting Summary. Please state if this information has not been collected. Report sex-based analyses where performed, justify reasons for lack of sex-based analysis.</i>
Field-collected samples	<i>For laboratory work with field-collected samples, describe all relevant parameters such as housing, maintenance, temperature, photoperiod and end-of-experiment protocol OR state that the study did not involve samples collected from the field.</i>
Ethics oversight	<i>Identify the organization(s) that approved or provided guidance on the study protocol, OR state that no ethical approval or guidance was required and explain why not.</i>

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Clinical data

Policy information about [clinical studies](#)

All manuscripts should comply with the ICMJE [guidelines for publication of clinical research](#) and a completed [CONSORT checklist](#) must be included with all submissions.

Clinical trial registration	<i>Provide the trial registration number from ClinicalTrials.gov or an equivalent agency.</i>
Study protocol	<i>Note where the full trial protocol can be accessed OR if not available, explain why.</i>
Data collection	<i>Describe the settings and locales of data collection, noting the time periods of recruitment and data collection.</i>
Outcomes	<i>Describe how you pre-defined primary and secondary outcome measures and how you assessed these measures.</i>

Dual use research of concern

Policy information about [dual use research of concern](#)

Hazards

Could the accidental, deliberate or reckless misuse of agents or technologies generated in the work, or the application of information presented in the manuscript, pose a threat to:

No	Yes	
☐	☐	Public health
☐	☐	National security
☐	☐	Crops and/or livestock
☐	☐	Ecosystems
☐	☐	Any other significant area

Experiments of concern

Does the work involve any of these experiments of concern:

No	Yes
☐	☐ Demonstrate how to render a vaccine ineffective
☐	☐ Confer resistance to therapeutically useful antibiotics or antiviral agents
☐	☐ Enhance the virulence of a pathogen or render a nonpathogen virulent
☐	☐ Increase transmissibility of a pathogen
☐	☐ Alter the host range of a pathogen
☐	☐ Enable evasion of diagnostic/detection modalities
☐	☐ Enable the weaponization of a biological agent or toxin
☐	☐ Any other potentially harmful combination of experiments and agents

Plants

Seed stocks	n/a
Novel plant genotypes	n/a
Authentication	n/a

ChIP-seq

Data deposition

- ☐ Confirm that both raw and final processed data have been deposited in a public database such as [GEO](#).
- ☐ Confirm that you have deposited or provided access to graph files (e.g. BED files) for the called peaks.

Data access links <i>May remain private before publication.</i>	<i>For "Initial submission" or "Revised version" documents, provide reviewer access links. For your "Final submission" document, provide a link to the deposited data.</i>
Files in database submission	<i>Provide a list of all files available in the database submission.</i>
Genome browser session (e.g. UCSC)	<i>Provide a link to an anonymized genome browser session for "Initial submission" and "Revised version" documents only, to enable peer review. Write "no longer applicable" for "Final submission" documents.</i>

Methodology

Replicates	<i>Describe the experimental replicates, specifying number, type and replicate agreement.</i>
Sequencing depth	<i>Describe the sequencing depth for each experiment, providing the total number of reads, uniquely mapped reads, length of reads and whether they were paired- or single-end.</i>
Antibodies	<i>Describe the antibodies used for the ChIP-seq experiments; as applicable, provide supplier name, catalog number, clone name, and lot number.</i>
Peak calling parameters	<i>Specify the command line program and parameters used for read mapping and peak calling, including the ChIP, control and index files used.</i>
Data quality	<i>Describe the methods used to ensure data quality in full detail, including how many peaks are at FDR 5% and above 5-fold enrichment.</i>
Software	<i>Describe the software used to collect and analyze the ChIP-seq data. For custom code that has been deposited into a community repository, provide accession details.</i>

Flow Cytometry

Plots

Confirm that:

- ☐ The axis labels state the marker and fluorochrome used (e.g. CD4-FITC).
- ☐ The axis scales are clearly visible. Include numbers along axes only for bottom left plot of group (a 'group' is an analysis of identical markers).
- ☐ All plots are contour plots with outliers or pseudocolor plots.
- ☐ A numerical value for number of cells or percentage (with statistics) is provided.

Methodology

Sample preparation

Describe the sample preparation, detailing the biological source of the cells and any tissue processing steps used.

Instrument

Identify the instrument used for data collection, specifying make and model number.

Software

Describe the software used to collect and analyze the flow cytometry data. For custom code that has been deposited into a community repository, provide accession details.

Cell population abundance

Describe the abundance of the relevant cell populations within post-sort fractions, providing details on the purity of the samples and how it was determined.

Gating strategy

Describe the gating strategy used for all relevant experiments, specifying the preliminary FSC/SSC gates of the starting cell population, indicating where boundaries between "positive" and "negative" staining cell populations are defined.

- ☐ Tick this box to confirm that a figure exemplifying the gating strategy is provided in the Supplementary Information.

Magnetic resonance imaging

Experimental design

Design type

Indicate task or resting state; event-related or block design.

Design specifications

Specify the number of blocks, trials or experimental units per session and/or subject, and specify the length of each trial or block (if trials are blocked) and interval between trials.

Behavioral performance measures

State number and/or type of variables recorded (e.g. correct button press, response time) and what statistics were used to establish that the subjects were performing the task as expected (e.g. mean, range, and/or standard deviation across subjects).

Acquisition

Imaging type(s)

Specify: functional, structural, diffusion, perfusion.

Field strength

Specify in Tesla

Sequence & imaging parameters

Specify the pulse sequence type (gradient echo, spin echo, etc.), imaging type (EPI, spiral, etc.), field of view, matrix size, slice thickness, orientation and TE/TR/flip angle.

Area of acquisition

State whether a whole brain scan was used OR define the area of acquisition, describing how the region was determined.

Diffusion MRI

☐ Used

☐ Not used

Preprocessing

Preprocessing software

Provide detail on software version and revision number and on specific parameters (model/functions, brain extraction, segmentation, smoothing kernel size, etc.).

Normalization

If data were normalized/standardized, describe the approach(es): specify linear or non-linear and define image types used for transformation OR indicate that data were not normalized and explain rationale for lack of normalization.

Normalization template

Describe the template used for normalization/transformation, specifying subject space or group standardized space (e.g. original Talairach, MNI305, ICBM152) OR indicate that the data were not normalized.

Noise and artifact removal

Describe your procedure(s) for artifact and structured noise removal, specifying motion parameters, tissue signals and physiological signals (heart rate, respiration).

Volume censoring

Define your software and/or method and criteria for volume censoring, and state the extent of such censoring.

Statistical modeling & inference

Model type and settings

Specify type (mass univariate, multivariate, RSA, predictive, etc.) and describe essential details of the model at the first and second levels (e.g. fixed, random or mixed effects; drift or auto-correlation).

Effect(s) tested

Define precise effect in terms of the task or stimulus conditions instead of psychological concepts and indicate whether ANOVA or factorial designs were used.

Specify type of analysis: ☐ Whole brain ☐ ROI-based ☐ Both

Statistic type for inference

Specify voxel-wise or cluster-wise and report all relevant parameters for cluster-wise methods.

(See [Eklund et al. 2016](#))

Correction

Describe the type of correction and how it is obtained for multiple comparisons (e.g. FWE, FDR, permutation or Monte Carlo).

Models & analysis

n/a | Involved in the study

☐☐ Functional and/or effective connectivity☐☐ Graph analysis☐☐ Multivariate modeling or predictive analysis

Functional and/or effective connectivity

Report the measures of dependence used and the model details (e.g. Pearson correlation, partial correlation, mutual information).

Graph analysis

Report the dependent variable and connectivity measure, specifying weighted graph or binarized graph, subject- or group-level, and the global and/or node summaries used (e.g. clustering coefficient, efficiency, etc.).

Multivariate modeling and predictive analysis

Specify independent variables, features extraction and dimension reduction, model, training and evaluation metrics.