The computational cost of active information sampling before decision-making under uncertainty

Pierre Petitet^{1,*}, Bahaaeddin Attaallah², Sanjay G. Manohar^{1,2} & Masud Husain^{1,2,*}

Affiliations

¹*Department of Experimental Psychology, University of Oxford, Oxford OX1 3PH, United Kingdom*

²*Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford OX3 9DU, United Kingdom*

* Corresponding authors: pierre.petitet@gmail.com (P.P.); masud.husain@ndcn.ox.ac.uk (M.H.)

Humans often seek information to minimise the pervasive effect of uncertainty on decisions. Current theories explain how much knowledge people should gather prior to a decision, based on the cost-benefit structure of the problem at hand. Here, we demonstrate that this framework omits a crucial agent-related factor: the cognitive effort expended while collecting information. Using an active sampling paradigm, we unveil a speed-efficiency trade-off whereby more informative samples take longer to find. Crucially, under sufficient time pressure, humans can break this trade-off, sampling both faster and more efficiently. Computational modelling demonstrates the existence of a cost of cognitive effort which, when incorporated into theoretical models, provides a better account of people’s behaviour, and also relates to self-reported fatigue accumulated during active sampling. Thus, the way people seek knowledge to guide their decisions is shaped not only by task-related costs and benefits, but also crucially by the quantifiable computational costs incurred.

Introduction

Obtaining information is a crucial part of normal decision-making under uncertainty. For example, deciding where to go out for dinner requires building predictions about how good the food, atmosphere and service will be in nearby restaurants. Incorrect predictions can ultimately result in poor choices, leading to deviations from the intended outcome (e.g., a bad dinner). Such disappointing experiences can be prevented, however, by spending time to gather information before making a decision in order to reduce uncertainty and improve predictions. To do this, humans need to solve two key problems. First,

1 they need to identify, in a complex high-dimensional world, what constitutes potentially useful infor-
2 mation, and devise a strategy to extract it efficiently. For instance, should they read online customer
3 reviews, ask a friend, or check a guidebook? Second, they must set a stopping policy that determines
4 when to stop gathering information and use the knowledge accumulated to guide their final decision.

5 Current accounts of active information gathering¹⁻⁶ are grounded in neuroeconomic theories of
6 decision-making⁷⁻¹¹. According to this perspective, people behave to maximise the expected utility
7 (EU) of their final decision, i.e., how rewarding the outcome of their choice is likely to be. A discrete
8 piece of information (sample) is worth acquiring in this framework only if the instrumental benefit it
9 provides outweighs the cost of obtaining it. This instrumental benefit is the increase in EU that results
10 from making the final decision with the information, compared to without it (ΔEU). Such an ap-
11 proach has proven useful in understanding why people seek more knowledge when sampling is cheap,
12 when errors are strongly penalised, or when information content of individual samples is low^{1,5,6,12-17}.
13 In other words, it provides a useful conceptual framework for understanding how stopping policies
14 respond to changes in the cost-benefit structure of the task environment.

15 Previous investigations of cognitive information sampling have largely focused on character-
16 ising such stopping policies. They have used simple paradigms in which resources, e.g., monetary
17 credits, could solely be used towards acquiring *more*, but crucially not *better*, samples^{1,5,13,18}. Such
18 an approach, however, neglects the fact that, in day-to-day life, not all potential samples are equally
19 informative. Instead, a well-chosen piece of information may have the potential to strongly resolve
20 uncertainty, while others may be irrelevant. It is therefore important to determine *which* data to acquire
21 in addition to *how much*, prior to making a decision. Crucially, the process of identifying the best
22 sample to collect may require cognitive effort as well as time, adding to a computational cost that is
23 currently not accounted for by theoretical models of active information sampling^{19,20}.

24 Two factors which theoretically might contribute to this computational cost are worth consider-
25 ing in more detail. First, identifying samples that are more informative potentially takes longer, and
26 thus carries an opportunity cost^{21,22}. In the restaurant example, asking a friend may provide more use-
27 ful information than consulting a guidebook, but we risk not obtaining that information before dinner
28 if it requires texting, then awaiting a response. Here, the opportunity cost is the potentially useful
29 information missed because of the decision deadline. Numerically, it is therefore a function of the

1 agent's sampling strategy (speed, efficiency), as well as the task's cost-benefit and temporal structure.
2 The second factor is a potential cognitive effort cost^{23,24} which arises from mechanisms that enhance
3 the amount of information gathered per unit of time²⁵, e.g., increasing the rate at which candidate sam-
4 ples are assessed, inhibiting unhelpful automatic or habitual responses, or allocating more attentional
5 resources to the search for knowledge. For example, if we decide to phone our friend, we might obtain
6 more information by turning the conversation towards the features of a restaurant that interest us most,
7 and paying close attention to the responses. The cognitive effort expended in honing in on the most
8 valuable information might lead to better decisions but also incur a cost in terms of neural resources
9 and/or energy consumption.

10 We hypothesised that if people discount the computational cost expended from information ben-
11 efits, then active sampling for information prior to decision-making should be more flexible than cur-
12 rently assumed. Testing this empirically requires experimental features that are missing from existing
13 paradigms^{1,13,18,26,27}: 1) some candidate samples should be more informative than others, and 2) it
14 should be non-trivial to find them. Here, we designed a more complex, active information sampling
15 task that incorporates these key features. In our paradigm, people search spatially for information, by
16 touching screen areas, before committing to an answer. Importantly, some sample locations give more
17 information than others, but finding them is not straightforward.

18 Using this paradigm, we demonstrate over five separate experiments that information sampling
19 prior to decision making is extremely flexible and adapts to contextual constraints. People are able to
20 increase the rate at which they accumulate knowledge either by acquiring data faster, or by slowing
21 down to select better observations at each step. Thus the results demonstrate a speed-efficiency trade-
22 off whereby more informative samples take longer to find. But crucially, when under sufficient time
23 pressure (when there is a substantial opportunity cost), humans are capable of breaking this trade-off,
24 sampling both faster and more efficiently. To account for such flexibility, we posit a cognitive control
25 mechanism that arbitrates how quickly and efficiently people sample information. In our model of cost-
26 benefit evaluation, both speed and efficiency are treated as costly, incurring a cognitive effort cost for
27 the individual. Further, this effort cost relates to subjective self-reported fatigue accumulated during
28 active sampling. Taken together, these findings show that incorporating a computational (cognitive)
29 cost into classical models of active information sampling provides a better account of human behaviour

1 prior to making decisions under uncertainty.

2 Results

3 **Active information sampling paradigm.** A schematic representation of the paradigm – framed as a
4 game called “Circle Quest” – is depicted in Fig. 1a. The goal was to maximise earnings in the form
5 of credits. Players searched for information about the location of a hidden circle (radius: 130 Px) by
6 tapping a touchscreen with their index finger. If they touched inside the hidden circle, a purple dot
7 appeared; or else, if they touched outside the hidden circle, a white dot appeared. These dots, denoting
8 the results of sampling at various locations, stayed on the screen throughout the information sampling
9 phase and until the end of the trial in order to limit memory load. Participants started each trial with
10 a certain number of credits (starting credit, R_0), which decreased by a fixed amount every time they
11 acquired a sample (sampling cost, η_s). During the time allocated to sampling (t_{max}), they were free to
12 acquire as many samples as judged necessary, at the rate they preferred, anywhere in the search space.
13 Once this time elapsed, individuals were asked to place a blue disk where they thought the hidden circle
14 was located. This response was subsequently rewarded according to the following scoring function:
15 $R_0 - s.\eta_s - e.\eta_e$; where s is the number of samples acquired, e is the placement error (in Px), and
16 η_e is the error cost (linear cost, in credits/Px). Thus, performing well in this task required minimising
17 uncertainty to avoid large error penalties, while also minimising sampling-related losses.

18 **Quantifying uncertainty with the Expected Error.** In this paradigm, “uncertainty” pertains to loca-
19 tion uncertainty of the hidden circle in the search space. It was quantified objectively as the Expected
20 Error (EE , in Px), i.e., the error an ideal observer is most likely to obtain by placing the blue circle
21 at the best possible location given the observations on the screen (Fig. 1b; details in *Supplementary*
22 *Methods*). This formalism is more suited than the Entropy of the posterior belief – another measure
23 of uncertainty quantifying the size of the solution space on a *log* scale – because it has direct instru-
24 mental implications in the game ($EV = R - \eta_e.EE$; where EV is the rational Expected Value).
25 Because the EE is a rather abstract dimension that may be difficult to intuit, we first asked, using a
26 simplified counterpart of the game, whether participants were able to manipulate this variable within a
27 decision-making problem.

Figure 1: Active information gathering task, with large and complex sampling space. a. In our paradigm, people search spatially for a hidden circle, by touching screen areas, before committing to an answer. The hidden circle is 5.80% the size of the arena. They win credits as a function of the number of samples used ($R_0 - s.\eta_s$), and the precision of their response ($e.\eta_e$). **b.** For any visual display, the posterior belief is the set of all locations with equal, non-zero, probability of being the centre of the hidden circle. The Expected Error (EE) is the probability-weighted average of all possible errors an agent could incur by placing the blue disk at the best possible location (centroid of the posterior belief, red circle). **c.** The large sampling space allows for many ways to gather information, i.e., reduce EE . The information extraction rate, α , quantifies sampling efficiency, i.e., how much EE decays from one sample to the next. It was estimated by fitting this simple exponential relationship to individual trials (see *Methods*). For illustration purposes, a slow-efficient (purple) and fast-inefficient (amber) trial are plotted. Extended Data Fig. 3 illustrates the evolution of the posterior belief over the first 6 samples of these example trials. Extended Data Fig. 4 shows the average reduction in EE over successive samples. This figure uses resources from Freepik.com.

1 In Exp. 1a, eight samples were presented on the screen (4 white, 4 purple), and individuals ($n =$
2 48, Mean age: 24.52 yr, s.d.: 3.75 yr; Extended Data Fig. 1) simply had to rate, on a visual analogue
3 scale, how well they thought they could localise the hidden circle (Fig. 2a; see *Methods*). In this task,
4 there was no active sampling; instead participants simply made judgements on a single “snapshot”
5 which provided them information. Uncertainty (EE) was manipulated experimentally by varying
6 the geometrical distribution of the samples. On a trial-by-trial basis, directly after their confidence
7 rating, a number of credits available appeared (40, 70, 100 or 125 credits), and individuals had to
8 indicate whether they would accept the offer to place the blue disk for this combination of reward and
9 uncertainty, or not. To incentivise authentic choices, participants were instructed that, at the end of
10 the experiment, they would have to play (i.e., place the blue circle and win credits according to their
11 performance) for a fixed number of trials, selected from the preferences they expressed.

12 Uncertainty ratings were calculated by sign-flipping and z-scoring confidence ratings for each in-
13 dividual. Objective uncertainty (EE) showed a significant linear relationship to these subjective ratings
14 (LMM - standardised $\beta = 0.80$, 95% CI = [0.76, 0.83], $t_{(8062)} = 49.89$, $p < 0.001$; Supplementary Ta-
15 ble 1; Fig. 2b), suggesting that participants were reporting an internal variable that resembled the EE .
16 This metric in fact provided a good predictive estimate of actual placement errors (Supplementary Fig.
17 1). Further, as demonstrated by logistic linear modelling of accept-reject decisions, individuals cor-
18 rectly weighed the EE against reward in order to decide whether or not to place the blue circle (Mixed
19 effects logistic regression - Main effect of EE : standardised $\beta = -2.53$, 95% CI = [-2.80, -2.27],
20 $t_{(8060)} = -8.57$, $p < 0.001$; Main effect of reward: standardised $\beta = 1.87$, 95% CI = [1.49, 2.26],
21 $t_{(8060)} = 7.04$, $p < 0.001$; Supplementary Table 1; Fig. 2b-c). This control task therefore confirmed
22 that human participants were able to report the EE as well as use it to inform decision-making within
23 our paradigm.

Figure 2: A passive counterpart of the game shows participants are sensitive to reward and uncertainty. **a.** In Exp. 1a, a passive version of the task was used to characterise how uncertainty (expected error) is 1) perceived (Confidence rating), and 2) combined with reward into a subjective value (accept-reject decisions). **b.** Participants ($n = 48$) reported an uncertainty estimate that closely resembled the Expected Error ($EE \times$ Uncertainty Score: standardised $\beta = 0.80$, 95% CI = [0.76, 0.83], $t_{(8062)} = 49.89$, $p < 0.001$). **c.** They were also more likely to accept placing the blue circle when the EE was low (Main effect of EE on the decision to accept the offer: standardised $\beta = -2.53$, 95% CI = [-2.80, -2.27], $t_{(8060)} = -8.57$, $p < 0.001$) and the number of credits available high (Main effect of reward on the decision to accept the offer: standardised $\beta = 1.87$, 95% CI = [1.49, 2.26], $t_{(8060)} = 7.04$, $p < 0.001$). In both panels, dots show group means, and error bars show s.e.m. Full statistical details are reported in Supplementary Table 1. This figure uses resources from Freepik.com.

1 **A speed-efficiency trade-off constrains sampling behaviour.** In a separate, counter-balanced ses-
2 sion, the same participants also underwent the active version of the task which was the main focus of
3 our study (Exp. 1b; see *Methods*). In this condition, they were asked to actively sample the search
4 space before deciding where to place the blue circle in order to win credits (Fig. 1a). Both the starting
5 credit (R_0) and sampling cost (η_s) were manipulated in a two-by-two block design (η_e & t_{max} were
6 fixed; Extended Data Fig. 2).

7 The game was designed such that, at any given stage, some candidate samples carried useful
8 information (i.e. they greatly reduced the EE), while others did not (Extended Data Fig. 3). For ex-
9 ample, sampling in-between dots of the same colour, or more than two radiuses away from a purple dot
10 did not provide any new information (i.e. $\Delta EE = 0$). As a result, the EE could decay more or less
11 sharply over successive observations, depending on the quality of the agent’s choices of sampling loca-
12 tions. We quantified this dimension of behaviour (sampling efficiency) with the information extraction
13 rate, α , estimated separately for each trial (see *Methods*; Fig. 1c). This parameter quantifies the rate of
14 exponential decay of uncertainty over successive trials (Extended Data Fig. 4). Throughout the paper,
15 greater α means sharper EE reduction from one sample to the next, i.e., more efficient information
16 gathering.

17 Unlike in many existing paradigms^{1,5,13,18,26,27}, α was not imposed or fixed in our task. Instead,
18 participants had agency over the efficiency of their own exploration strategy. At any stage of the search,
19 they could choose to deliberate longer to acquire a sample at a more informative location, which in turn
20 would result in a sharper reduction of uncertainty. Regardless of the experimental variable manipulated
21 (R_0 or η_s), slower acquisition rate – greater inter-sampling interval (ISI) – was indeed systematically
22 associated with more efficient sampling in Exp. 1b (LMM - Main effect of ln-transformed ISI on
23 α : $t_{(4606)} = 10.16$, $p < 0.001$; Supplementary Table 2; Fig. 3a). Further analysis of the strategies
24 underlying exploration behaviour suggested that time-dependent efficiency benefits involve both the
25 inhibition of unhelpful automatic sampling actions, and a closer approximation of the optimal search
26 algorithm (details in *Supplementary Results*).

Figure 3: Economic and time constraints influence the extent, speed and efficiency of active information sampling. **a.** Across all four experimental conditions of Exp. 1 ($n = 48$ participants), slower sampling rate (longer inter-sampling intervals) led to sharper reductions of uncertainty from one sample to the next (larger α ; Main effect of ln-transformed ISI on α : $t_{(4606)} = 10.16$, $p < 0.001$). The ISI and α were calculated separately for each trial. Conditional plots were generated by sliding, for each subject and condition separately, 25% quantile bins of ISI and computing the mean α for each bin. The error bar shows the group s.e.m. **b-d.** Participants sampled slower, more efficiently but less extensively when acquiring data was more expensive (high vs. low η_s). By contrast, a larger starting credit led to a small increase in sampling speed and amount of data acquired by had no significant effect on efficiency (high vs. low R_0). Points and error bars show group mean $\pm$ s.e.m. Asterisks show significant effects (vertical bars: main effect of η_s ; horizontal bars: main effect of R_0 ; $p < 0.05$). Full statistical details are reported in Supplementary Tables 2-6

1 **A model of active sampling penalising speed and efficiency.** The existence of the trade-off illus-
2 trated in Fig. 3a allowed us to ask the central question of this paper: How do humans choose the
3 speed and efficiency at which to gather information prior to making a decision? Further, when do they
4 prioritise speed over efficiency (or vice versa)? And can they sometimes prioritise both?

5 From an optimal control^{28–30} perspective, solving this problem requires assigning a subjective
6 cost to both sampling speed and efficiency. Fig. 4 illustrates this framework. The Expected Utility
7 (EU) is calculated as the reward available (remaining credits: $R_0 - \eta_s \cdot s$) minus the expected error
8 penalty ($\eta_e \cdot EE$) minus a cumulative cognitive effort penalty ($\eta_c \cdot s$). We posit that the rate at which
9 this cognitive effort penalty grows over successive samples (η_c) may be a function of sampling speed
10 ($1/ISI$) and/or efficiency (α). Within this framework, a rational decision maker chooses the course
11 of action (combination of sampling speed, efficiency, and extent) that maximises return (i.e., the EU
12 at the time of the decision). In Exp. 1b, an increase in starting credit (R_0) should therefore not affect
13 behaviour much because it merely influences the magnitude of the EU but not the course of action
14 required to maximise it. By contrast, an increase in sampling cost (η_s) should act as an incentive to
15 sample less, slower and more efficiently to compensate for the additional loss incurred at each step of
16 the search.

Figure 4: **Choosing how quickly, efficiently and extensively to sample.** **a.** Standard accounts of active information sampling posit that individuals behave so as to maximise the expected value of stopping, which is here calculated as $EV = R_0 - \eta_s \cdot s - \eta_e \cdot EE$ (i.e. EV = Starting credit - Sampling Penalty - Expected Error Penalty; see also Supplementary Fig. 3). This formalism is useful to decide how many samples to acquire but cannot predict how quickly or efficiently to sample. In order to do so, we propose to introduce a cognitive effort cost, η_c , that agents have to pay for each sample acquired, and that may be a function of both sampling speed ($1/ISI$) and/or efficiency (α). **b.** The Expected Utility (EU) is then defined as the previous EV plus the expected information-related benefit that would result from sampling minus the economical (η_s) and cognitive effort cost (η_c) expanded finding the sample. The acquirability equals to one for samples obtainable during the time allocated to the search, and zero otherwise. **c.** Various candidate cognitive effort cost functions were evaluated in this paper (linear, quadratic, exponential, logistic). When $\eta_c = 0$ (Models 1 & 2), our framework reduces to previous accounts. Other models (Models 3-14) differ in the way sampling speed ($1/ISI$) and efficiency (α) are penalised.

1 Raw behavioural data in Exp. 1b were qualitatively consistent with the predictions of this frame-
 2 work (Fig. 3b-d). When acquiring samples was more expensive, participants indeed explored less
 3 extensively – choosing to make fewer samples (LMM - Main effect of η_s on the number of samples ac-
 4 quired: $t_{(4604)} = -18.64, p < 0.001$; Supplementary Table 3; Fig. 3d). Remarkably, they also slowed
 5 down their exploration (LMM - Main effect of η on ln-transformed ISI: $t_{(4604)} = 13.83, p < 0.001$;
 6 Supplementary Table 4; Fig. 3b), and became more efficient as indexed by the information extraction
 7 rate (LMM - Main effect of η_s on ln-transformed α : $t_{(4604)} = 2.54, p = 0.011$; Supplementary Table
 8 5; Fig. 3c). Participants also showed a very small but significant sensitivity to the starting credit: they
 9 sampled slightly faster (LMM - Main effect of R_0 on ln-transformed ISI: $t_{(4604)} = -3.94, p < 0.001$;
 10 Supplementary Table 4; Fig. 3b) and more extensively (LMM - Main effect of R_0 on the number of
 11 samples acquired: $t_{(4604)} = 7.08, p < 0.001$; Supplementary Table 3; Fig. 3d) when more reward
 12 was available. This sensitivity effect should not exist if economical variables are represented veridi-
 13 cally (rational decision-maker), and may instead reflect the use of a non-linear utility function. For
 14 completeness, the effect of all experimental variables on the placement error and score are reported in
 15 Supplementary Tables 7 & 8.

16 In order to formally assess whether positing a cognitive effort cost in terms of sampling speed and
 17 efficiency was indeed required to explain human behaviour quantitatively in our task, we fitted variants
 18 of the modelling framework we proposed to individual datasets. Model fitting required the introduction
 19 of free parameters to the framework, estimated at the individual level. All models included two free
 20 weights, w_e and w_s , which quantified the influence of information-related benefits and sampling cost
 21 to the *EU* (Fig. 4b; see *Methods*). The reference model (Model 1) included no cognitive effort cost
 22 at all, and was in effect similar to standard views which do not account for this factor. Because it
 23 is known that humans often exhibit non-linear utility functions when deciding under uncertainty¹¹ –
 24 as suggested here by the small sensitivity to starting credit (Fig. 3b-d) – we first evaluated a model
 25 with a power term on the sampling cost (Model 2). In this model, sampling speed and efficiency were
 26 not penalised but the contribution of sampling costs to the Expected Utility could increase or decrease
 27 non-linearly as a function of the number of samples acquired (details in *Methods*). It was used as a
 28 second reference model to check whether a non-linear utility function could account for the data in the
 29 absence of cognitive effort cost.

1 The subsequent models (Models 3-14) evaluated the four classes of cognitive effort cost depicted
2 in Fig. 4c: linear, quadratic, exponential and sigmoid. Thus, they differed in the way sampling speed
3 ($1/ISI$) and efficiency (α) were penalised. Note that exponential and logistic cost functions required
4 an extra free parameter relative to linear and quadratic ones (see *Methods*; Extended Data Fig. 5).
5 To avoid over-fitting by overly complex models, a veridical specification of economic variables was
6 used when a cognitive effort cost was included (i.e., no power term on sampling cost in Models 3-
7 14). For each class of cost function, three versions of the model were built, based on whether 1)
8 speed, 2) efficiency, or 3) both speed and efficiency were penalised. The free weights w_{speed} and w_{α}
9 quantified how severely sampling speed and efficiency were penalised within the cognitive effort cost
10 function. Free parameters were estimated separately for each individual dataset using a maximum
11 (multivariate) likelihood procedure, implemented in MATLAB with Bayesian Adaptive Direct Search
12 (BADS³¹). Model comparison was carried out with Bayesian Information Criterion (BIC), chosen for
13 its stringent penalisation of model complexity. Smaller BIC indicates better trade-off between model
14 fit and complexity.

15 In the absence of a cognitive effort cost, positing a non-linear utility function provided a better
16 account of the data than a purely rational specification (Model 2 vs. 1: $\Delta BIC = -1,036$; Fig. 5a).
17 However, this was a rather small explanatory power benefit compared to those observed with models
18 including a cognitive effort cost (Models 3-14). Overall, the winning model (Model 8) penalised for
19 both speed and efficiency using a quadratic cost function (Model 8 vs. 1: $\Delta BIC = -22,779$; Fig.
20 5a). The winning cognitive effort cost function was:

$$21 \quad \eta_c(ISI, \alpha) = w_0 + w_{speed} \cdot \frac{1}{ISI^2} + w_{\alpha} \cdot \alpha^2 \quad (\text{Quadratic}) \quad (1)$$

22 Thus, as expected, positing the existence of a cognitive effort cost – due to both a cost of speed
23 and of efficiency – improved our ability to explain participants' choices of sampling speed, efficiency
24 and extent in Exp. 1b (see *Supplementary Information* for other models that did not outperform Model
25 8; Supplementary Fig. 2).

26 A potential limitation of the winning model is that, unlike in simple motor tasks³², the intrinsic
27 relationship between speed and efficiency is poorly understood here. As a result, the addition of the

Figure 5: **A model comparison robustly shows across 5 experiments that a model with a quadratic cognitive effort cost best explains active information sampling behaviour.** Model 1 and 2 included no cognitive effort cost (Model 1: rational specification; Model 2: non-linear utility function; in grey) and were used as a reference. The other 12 models included a cognitive effort cost specified in terms of either speed (Models 3, 6, 9, 12), efficiency (Models 4, 7, 10, 13) or both (Models 5, 8, 11, 14). They differed with respect to the shape of the cognitive effort cost function used (green: linear; orange: quadratic; blue: exponential; pink: sigmoid). **a.** In Exp. 1b ($n = 48$ participants), individuals' choices of sampling speed, efficiency and extent were best explained by a model penalising both speed and efficiency with a quadratic cost function (Model 8, lowest Bayesian Information Criterion, BIC). **b.** Model comparison from Exp. 2-5 (BIC summed across all four experiments, $n = 79$ participants) replicated the winning model from Exp. 1b (Model 8). This was also the case when considering each experiment separately (Extended Data Fig. 7).

1 degrees of freedom we proposed might merely reflect the need for the model to incorporate a specifica-
2 tion of this relationship, rather than a true cognitive effort cost. The last two sections provide evidence
3 in favour of the existence of a true cognitive effort cost that influences active information sampling.

4 **Efficiency penalties reflect a cognitively effortful process.** To seek independent confirmation that
5 the cost introduced in our framework indeed captures an effortful process, participants were asked
6 to rate, on a visual analogue scale, how tired they felt throughout the experimental session in Exp.
7 1b (see *Methods*). On average, fatigue built up linearly over successive task blocks (LMM - Main
8 effect of Block on Fatigue Ratings: $t_{(1150)} = 4.16$, $p < 0.001$; Supplementary Table 9; Fig. 6a).
9 Importantly, the amount by which it increased showed a significant association with the contribution
10 of efficiency to individuals' cognitive effort cost (w_α , as extracted through model fitting; see *Methods*).
11 Those participants in whom increasing efficiency was subjectively more costly (high w_α) also suffered
12 a greater increase in subjective fatigue as a result of performing the task (standardised $\beta = 0.49$,
13 95% CI = [0.17, 0.80], $t_{(43)} = 3.13$, $p = 0.0032$; Robust multiple linear regression controlling for
14 w_0 , w_{speed} and Experiment duration; Supplementary Table 10; Fig. 6b). By contrast, those for whom
15 increasing sampling efficiency was less penalised got less tired over the course of the experiment.
16 Performing the analysis on z-scored ratings did not change the conclusion (Supplementary Tables 9 &
17 10). Thus, the efficiency cost introduced in our modelling framework captures an effortful process that
18 was meaningfully related to how subjective fatigue developed throughout the task.

19 **Breaking the speed-efficiency trade-off under time pressure.** A second way to adjudicate whether
20 the speed-efficiency trade-off we uncovered involves a cognitive effort cost is to ask whether it is
21 “breakable” under certain experimental conditions. If the trade-off solely reflects the fact that sampling
22 efficiency is naturally a function of speed, then participants should be constrained to manipulating
23 these two dimensions in an anti-correlated fashion. In other words, sampling faster comes at the cost
24 of efficiency, and vice versa. If, alternatively, the trade-off emerges from a cognitive effort cost, then,
25 under sufficient experimental pressure, participants should be able to increase speed without sacrificing
26 efficiency, or vice versa, or even increase both at the same time. Note that these two hypotheses are
27 not mutually exclusive, and the key prediction is that of a breakable trade-off if a cognitive effort cost
28 contributes to information sampling behaviour.

29 In order to test whether the speed-efficiency trade-off could be broken in some conditions, the

Figure 6: **The cost of efficiency is associated with fatigue development.** **a.** In Exp. 1b, participants ($n = 48$) rated how tired they felt throughout the experimental session. The average rating during the exposure task (blocks -3 -1) was used as baseline fatigue, which was subtracted from subsequent ratings for normalisation purposes. On average, fatigue increased linearly over successive blocks (Main effect of Block on Fatigue Ratings: $t_{(1150)} = 4.16$, $p < 0.001$). **c.** The cost of efficiency (w_α) was associated with the increase in fatigue observed as a result of the experimental session (standardised $\beta = 0.49$, 95% CI = [0.17, 0.80], $t_{(43)} = 3.13$, $p = 0.0032$; Robust multiple linear regression controlling for w_0 , w_{speed} and Experiment duration). The scatter plot shows the mean normalised fatigue rating in the last experimental block (y-axis) as a function of the cost of efficiency (w_α), controlling for task duration, w_0 , and w_{speed} ; task-duration (w_α and w_{speed} were ln-transformed for normalisation purposes). Colour shows the contribution of individual data points to the robust multiple linear regression (multivariate outliers are down-weighted). Full statistical details are reported in Supplementary Tables 9 & 10.

1 four experimental variables (R_0 , η_s , η_e , t_{max}) were manipulated in isolation in four separate experi-
 2 ments (Exps. 2-5; $n = 79$ participants divided into 4 groups, Mean age: 25.13 yr, s.d.: 4.05 yr; see
 3 *Methods*; Extended Data Figs. 1 & 2). Full statistical analysis of sampling behaviour for these exper-
 4 iments is provided in *Supplementary Results* (Extended Data Fig. 6). None of the economic variables
 5 (R_0 , η_s , η_e ; Exps. 2-4) led to violations of the speed efficiency trade-off (Fig. 7). That is, even when
 6 they adjusted speed and efficiency in response to changes in η_s (Figs. 3a & 7), individuals remained
 7 confined within a fixed anti-correlation between the two. This was evidence by the disappearing of
 8 the effect of η_s on α when controlling for sampling speed in Exp. 3 (LMM - Main effect of η_s on ln-
 9 transformed α : $t_{(1918)} = 3.97$, $p < 0.001$, Supplementary Table 5; Main effect of η_s on ln-transformed
 10 α while controlling for ln-transformed ISI: $t_{(1917)} = -0.60$, $p = 0.55$, Supplementary Table 6; also
 11 observed in Exp. 1b).

12 By contrast, time pressure (small t_{max}) acted as a drive for participants to break the speed-
 13 efficiency trade-off. In Exp. 5 ($n = 19$ participants), the time available for people to sample before
 14 making their final decision (t_{max}) was either 18 (as in previous experiments), or 14, 10 or 6 seconds
 15 (Extended Data Fig. 2). Compared to experiments in which economic variables were individually
 16 manipulated (R_0 , η_s , η_e in Exps. 2-4; $n = 60$ participants), players under overall heightened time
 17 pressure sampled on average not only faster (LMM - Contrast Exp. 5 vs. Exps. 2-4 for ln-transformed
 18 ISI: $F_{(1,7580)} = 47.65$, $p < 0.001$) but also more efficiently (LMM - Contrast Exp. 5 vs. Exps. 2-4 for
 19 ln-transformed α : $F_{(1,7580)} = 5.26$, $p = 0.022$; Fig. 7a). In other words, they broke the speed-efficiency
 20 trade-off observed in previous experiments.

21 Because this result relies on a between-group comparison, we checked that it could not be ac-
 22 counted for by inter-individual differences in IQ (Raven's progressive matrices³³), or self-reported
 23 motivation indices of apathy (AMI³⁴) and impulsivity (BIS-11³⁴). Including these external measures
 24 within the previous analysis made no difference to the result (LMM - Contrast Exp. 5 vs. Exps. 2-4 for
 25 ln-transformed ISI, while controlling for questionnaires: $F_{(1,7577)} = 48.30$, $p < 0.001$; Contrast Exp.
 26 5 vs. Exps. 2-4 for ln-transformed α : $F_{(1,7577)} = 4.82$, $p = 0.028$). Further and crucially, even within
 27 Exp. 5, there was evidence of a breaking of the speed-efficiency trade-off at the within-individual level.
 28 Shorter t_{max} led to faster sampling (LMM - Main effect of t_{max} on ln-transformed ISI: $t_{(1822)} = 11.71$,
 29 $p < 0.001$; Supplementary Table 4) with no significant loss in terms of efficiency (LMM - Main effect

1 of t_{max} on ln-transformed α : $t_{(1822)} = -0.02$, $p = 0.98$; Supplementary Table 5). This time, when
2 controlling for sampling speed in the analysis, α showed a significant increase with shorter t_{max} (LMM
3 - Main effect of t_{max} on ln-transformed α while controlling for ln-transformed ISI: $t_{(1821)} = -3.86$,
4 $p < 0.001$; Supplementary Table 6).

5 Taken together, the results demonstrate that participants could break the speed-efficiency trade-
6 off, sampling both faster and more efficiently. This qualitative feature can be seen as a direct mani-
7 festation of the existence of cognitive effort cost (Fig. 4). Formal model comparison for Exps. 2-5
8 confirmed this conclusion, replicating the results of Exp. 1b. Once again, human performance was
9 best explained when positing the existence of a quadratic cognitive effort cost penalising both speed
10 and efficiency (Model 8 vs. 1: $\Delta BIC = -37,800$; Fig. 5b; Extended Data Fig. 7). Taken together,
11 our results suggest that the way humans seek knowledge to guide their decisions is shaped not only by
12 the economic costs and benefits of information, but also crucially by the cognitive effort expenditure
13 involved in seeking this information.

Figure 7: A breakable speed-efficiency trade-off governs active information sampling. **a.** Overall, in Exps. 2-5 ($n = 79$ participants), longer inter-sampling intervals led to larger information gains at each step, as in Exp. 1b. Participants of Exp. 5 ($n = 19$, green line), however, were able to sample both faster (Contrast Exp. 5 vs. Exps. 2-4 for ln-transformed ISI: $F_{(1,7580)} = 47.65$, $p < 0.001$) and more efficiently (Contrast Exp. 5 vs. Exps. 2-4 for ln-transformed α : $F_{(1,7580)} = 5.26$, $p = 0.022$; see also Extended Data Fig. 6). In the framework we propose (Fig. 4), maintaining the same level of efficiency while sampling faster, or extracting more information given a constant inter-sampling interval, are considered to be cognitively effortful (e.g. consumes more energetic resource). These costs are captured by the model parameters w_{speed} and w_{α} . **b.** To visualise the effect of economic (Exps. 2-4) and time (Exp. 5) pressure on the speed-efficiency trade-off, the same plot was generated for each experimental condition (highest and lowest two levels grouped together). Individuals lost in speed what they gained in efficiency when sampling costs increased (Exp. 3), but were able to break this trade-off under time pressure (Exp. 5). In both panels conditional plots were generated by sliding, for each subject and condition separately, 25% quantile bins of ISI and computing the mean α for each bin. The error bar shows the s.e.m. Full statistical details are reported in Supplementary Tables 4-6.

1 Discussion

2 Modern theories of active sampling assume that the brain optimises the amount of information needed
3 prior to decision-making on the basis of the costs and benefits of resolving uncertainty. Accordingly,
4 informative samples are worth acquiring only when the subjective benefits they provide outweigh the
5 cost of obtaining them. This simple axiom has proven extremely powerful in explaining how much
6 information people seek before decision-making^{1,5,6,12–17}, and provides a theoretical bridge to other
7 types of motivated behaviours such as foraging^{4,35,36}. However, a natural key question is: What do
8 people treat as costs and benefits during active information sampling?

9 Here we focused on one cost whose contribution has been elusive so far – the cognitive effort
10 cost^{19,20}. We hypothesised that if the brain discounts the cognitive effort expended from information
11 benefits, then active sampling behaviour should be more flexible than currently assumed. More specif-
12 ically, if finding more informative samples comes at an extra cognitive cost, then people should search
13 fast and efficiently only when pressured by task constraints. Using a complex active sampling exper-
14 imental paradigm, we showed that humans indeed flexibly adjust the amount of cognitive resources
15 they allocate to active sampling, in response to economic and time pressures. They sample more ef-
16 ficiently but slower when the economic losses incurred increase (Exp. 1b & 3), and are able to break
17 this trade-off under time pressure, sampling both faster and more efficiently (Exp. 5; Fig. 7).

18 To account for this flexibility, we posited the existence of a cognitive control mechanism^{22,37} that
19 determines how quickly and efficiently to sample information. Importantly, in this cost-benefit evalu-
20 ation, both speed and efficiency are treated as costly. Incorporating this notion into a standard infor-
21 mation sampling model explained human behaviour better than simpler models. Further confirmation
22 of the psychological validity of our framework was provided by the significant relationship between
23 subjective fatigue reports and the model-derived cost of efficiency (Fig. 6). The findings have some res-
24 onance with existing behavioural theories that invoke cognitive effort to explain why consumers do not
25 compare products that differ along multiple dimensions³⁸, or why they do not read contracts³⁹. Within
26 a very distinct situation – one in which people have to actively sample their environment in order to
27 reduce uncertainty about a spatial dimension – we demonstrated that similar principles may apply.
28 Crucially, however, the model provided here is quantitative and accounts for changes in behaviour, that
29 is, it generates quantifiable predictions of sampling behaviour that are testable experimentally.

1 Consistent with current theories of control allocation in other domains^{30,40}, our results converged
2 on the view that the brain actually penalises fast and efficient information sampling, thereby promot-
3 ing what is effectively sub-optimal exploration in the absence of strong environmental pressure. It
4 was evident that under greater time pressure, individuals could sample faster and more efficiently, but
5 chose not to do so when there was a longer deadline (Fig. 7). There are at least three possible reasons
6 why such a mechanism might exist: metabolic cost, shared computing resources, and reliability test-
7 ing. First and foremost, fast and efficient sampling might increase metabolic consumption because
8 of additional neural resources requirements – for example, higher neural firing rate⁴¹, optimisation
9 of brain network organisation^{42,43}, more efficient information processing^{25,44}. An indirect proxy for
10 such a possible increase in (metabolically costly) processing demands was provided by the analysis
11 of choices of sampling location, which supported the flexible switching between search algorithms of
12 various computational complexity^{45,46} (see *Supplementary Results*).

13 The second explanation is that, because other cognitive processes also depend on the same shared
14 limited neural resources, prioritisation of one process might negatively impact on the others, resulting
15 in an opportunity cost in terms of unprocessed task-unrelated stimuli²². In this context, sub-optimal
16 sampling performance might “leave room” (capacity) for parallel processing of other potentially useful
17 information. Finally, although it impinges upon search efficiency, there might be indirect benefits to
18 acquiring apparently uninformative, confirmatory samples⁴⁷. Indeed, they provide timely checks of the
19 reliability of the information source that can only be performed where the agent has perfect knowledge
20 over what the outcome should be (i.e. null entropy). Challenging what is already known might confer
21 an ecological advantage in non-deterministic environments where the trustworthiness of past and future
22 observations may drift over time.

23 The computational framework used throughout this paper assumes that individuals choose how to
24 explore (i.e., how quickly, efficiently, and extensively to sample) on the basis of internal simulations of
25 what the future might hold for all alternative courses of action – or at least a simplified, representative
26 sub-sample of them⁴⁸. This principle is well-established within the sensorimotor control⁴⁹ and goal-
27 directed⁵⁰ literature. Deciding within this strategy space therefore requires metacognitive abilities^{51,52}
28 that may only be acquirable through experience (e.g., an expectation of the actual uncertainty reduction
29 for a given “internally-selected” efficiency, α ; an expectation of the actual error for a response under

1 a certain uncertainty level). It is possible that this meta-knowledge may, to a certain extent, be learnt
2 online as participants try out various strategies (see *Supplementary Information* for a visualisation of
3 behaviour as a function of task exposure). This idea is compatible with a recent account of cognitive
4 control in which a reinforcement meta-learner learns how to select behavioural states (e.g., physi-
5 cal/cognitive effort) through error-based updating coupled with value-based action selection⁵³. Our
6 paradigm could therefore be used in the future to investigate the computational architecture underlying
7 *de novo* learning of a cognitive effort cost function.

8 Previous information sampling tasks – the card sampling task¹⁸, tile flipping task¹, beads/urn
9 task²⁶, dart/fishing task⁵ – could not isolate cognitive effort cost from mere economic sampling cost
10 because participants had little to no control over the information extraction rate (here called α). If
11 all candidate samples are expected to reduce uncertainty by the same amount, regardless of the time
12 or cognitive effort spent choosing them, then active sampling becomes solely about whether or not to
13 sample (i.e., when to stop). Cognitive effort needs to confer an instrumental advantage for it to be
14 traded against other resources⁴⁰. It is worth emphasising, however, that in the simpler instance where
15 the information extraction rate is constant (i.e., fixed α), our active sampling model actually reduces
16 to previous accounts^{5,17} (Fig. 4). But if – as in the active sampling task introduced here and most real
17 life scenarios – agents do have control over the informativeness of individual samples, then optimal
18 allocation of computational resources becomes a central problem, which our model offers a solution
19 to.

20 Our study has three main limitations. First, the task relies heavily on motor performance, both to
21 sample the search space, and to produce a response at the end of every trial (Fig. 1a). Inter-individual
22 differences in motor abilities may affect performance in ways that are not – or wrongly – captured by
23 the model (Fig. 4). Although this source of variance is likely to be negligible in the population of young
24 healthy adults tested here, it may become a more significant confound in certain age groups or clinical
25 populations. Second, the definition of uncertainty used throughout this work (Expected Error) relies
26 on the assumption that people were able to use of all of the information that was available in the visual
27 display. It can be argued, however, that the level of uncertainty depends not only on the information
28 presented (i.e., the visual display), but also on the agent’s ability to use this information (i.e., which
29 solution they select within the solution space). Thus, for the same visual display, an individual who

1 is well able to identify the best solution within the solution space should rightly expect a lower error
2 than an individual who typically picks worse solutions. For simplification purposes, and because EE
3 provided a good estimate of actual placement errors (Supplementary Fig. 1), this layer of complexity
4 was omitted in the present work. Lastly, the model developed here assumes a constant sampling
5 speed and efficiency within a given trial, and therefore cannot account for within-trial effects (e.g.,
6 speeding-up or slowing-down as individuals get closer to the decision deadline). Considering within-
7 trial effects may allow future studies to better characterise dynamic decision processes unfolding during
8 the exploration. More specifically, it may enable the segregation of two distinct decisions: 1) whether
9 or not to sample, and 2) where to sample.

10 In conclusion, the results presented here identified cognitive effort as a central cost the brain
11 trades off against reward during active information sampling prior to decision-making under uncer-
12 tainty. The active sampling experimental paradigm introduced here, together with our modelling ap-
13 proach, offers a compact way to quantify the contribution of this cost.

14 **Methods**

15 **Participants.** All participants provided written informed consent, and the study was approved by the
16 University of Oxford ethics committee. Sample size was determined a priori, based on counterbalanc-
17 ing constraints, and previous studies with similar objectives^{5,54,55}. Recruitment took place via the Ox-
18 ford Psychology Research participant recruitment scheme website (<https://opr.sona-systems.com/>), and via online study advertisement. Participating in an experiment was a criterion of exclusion
19 for the other experiments, such that there was no overlap in participants between Exps. 1-5. Volunteers
20 were predominantly university students and young professionals. None of them reported undertaking
21 psycho-active drugs, or having personal history of neurological or psychiatric condition.

23 Experiment 1 included 48 participants (24 females), Exps. 2-4 included 20 participants each (10
24 females per Exp.), and Exp. 5 included 19 participants (10 females). Thus, in total, 127 different partic-
25 ipants were tested in this work. Demographics and participant characteristics are reported in Extended
26 Data Fig. 1. In Exp. 1, prospective volunteers (148 individuals volunteered) were selected based on
27 their motivation level, gender and age in order to ensure a large spread of motivational levels. The on-
28 line screening questionnaire included the Apathy Motivation Index (AMI), a validated self-report index

1 of apathy and motivation³⁴, and demographics information (gender, age, years of education, handed-
2 ness, medication, personal history of neurological or psychiatric condition). No participant dropped
3 out after inclusion, and all participants tested were included in the analysis.

4 **Apparatus.** All tasks were presented on a 17-inch touchscreen PC using MATLAB version 2018a
5 (MathWorks; <https://uk.mathworks.com>) and Psychtoolbox^{56,57} version 3. Participants sat in
6 a quiet room, within reaching distance of the screen (~ 50 cm). An experimenter was present in the
7 room at all time during behavioural testing.

8 **Active information sampling task (“Circle Quest”).** The goal of the task was to win as many credits
9 as possible. Participants were instructed that the visual display contained a purple circle, hidden behind
10 a circular grey mask (search space). The size of the hidden circle (radius: 130 Px; area: 5.80% of the
11 search space) was displayed on both sides for reference. The number of credits that could be won by
12 perfectly localising the hidden circle was displayed on the side at all time, and decreased every time
13 participants touched the screen during the sampling phase. This quantity represented an endowment
14 from which sampling and error costs were subtracted. When manipulated, the error cost (η_e) was
15 visually depicted on the side as a progress bar with the words “Scoring severity” at the top (Exp. 4).
16 This progress bar could contain 1, 2, 3 or 4 squares that corresponded to an error cost of 0.5, 1.5, 1.9 or
17 2.6 cds/Px respectively. It could change from one trial to the next but remained constant within a given
18 trial. Finally, when manipulated, a timer representing the remaining time allocated to the search was
19 displayed on both sides of the screen (Exp. 5). It took the shape of a timer that got filled clockwise.

20 Information about the location of the hidden circle could be gathered by tapping the screen. If
21 participants touched inside the hidden circle, a purple dot (radius: 4 Px) appeared. If they touched
22 outside the hidden circle, a white dot appeared. Trials comprised of two phases. During the active
23 sampling phase, participants could collect as many samples as they wanted, at any location inside the
24 search space. They were instructed that they could stop at any moment but had to wait until the end of
25 the allocated time before making their final response. A single purple dot was displayed on the screen
26 at the beginning of this phase, to act as an initial hint and shorten the length of the search. The location
27 of this first purple dot was drawn randomly from inside the hidden circle, and was always at least 260
28 Px away (diameter of the hidden circle) from the edge of the search space to ensure it always carried
29 the same amount of information ($EE_0 = 86.67$ Px). The number of credits available decreased every

1 time participants acquired a sample. At the end of the sampling phase, a blue disk of the same size as
2 the hidden circle (i.e. radius: 130 Px) appeared in the middle of the screen. Participants then had 6
3 seconds to drag it on top of where they thought the hidden circle was. Finally, the score was presented.
4 It was calculated using the formula displayed in Fig. 1a. Participants won all the remaining credits
5 $(R_0 - s.\eta_s)$ minus a penalty reflecting how far they placed the blue disk from the actual hidden circle
6 $(e.\eta_e)$. In other words, there were two ways of losing credits in this task: 1) by sampling information,
7 and 2) by mis-localising the hidden circle. The duration of each trial was experimentally set such that
8 participants could not complete more trials by sampling less/faster or placing the blue circle quickly.

9 **Passive decision task.** To characterise the extent to which participant were sensitive to the number of
10 credits and expected error on offer, we designed a version of the task in which these two variables were
11 experimentally manipulated on a trial-by-trial basis (Fig. 2a). The design of this task was inspired by
12 other decision-making task developed in the laboratory^{54,58–60}. This task was counter-balanced with
13 the active sampling version in Exp. 1.

14 First, participants saw dot samples on the screen. All trials contained four positive (purple)
15 and four negative (white) samples. The uncertainty (EE) was manipulated by varying the spatial
16 configuration of these samples. For example, the EE was greater when the positive samples were
17 clustered together as opposed to spread out. The same set of stimuli was used for all individuals but
18 the order of presentation was randomised. The EE ranged from 10 to 70 Px. Participants were asked
19 to first report on a scale from 0 to 100 how confident they felt about their ability to localise the hidden
20 circle, given the samples displayed on the screen. After the confidence rating, the number of credits
21 available appeared (either 40, 70, 100 or 125). Participants were then asked to accept or reject the
22 offer to place the blue circle. Note that participants simply expressed preferences in this version of the
23 task, such that there was no “correct” or “incorrect” response. They received no feedback during their
24 series of choices. To incentivise the context, they were instructed that a sub-sample of their “accepted”
25 trials would be selected at the end of the session for them to position the blue circles and win credits
26 accordingly.

27 In total, all participants underwent the same set of 168 trials (42 per reward level), divided into 6
28 experimental blocks (28 trials per block). The order of presentation was randomised.

1 **External measures.** Interleaved with the tasks, participants completed established self-report ques-
2 tionnaire measures of apathy (Apathy Motivation Index³⁴, AMI), impulsivity (Barratt Impulsiveness
3 Scale-11⁶¹, BIS-11), depression (Beck Depression Inventory-II⁶², BDI-II) and uncertainty preference
4 (Intolerance of Uncertainty Scale⁶³, IUS). They also performed a time-limited (10 minutes) test of
5 non-verbal fluid intelligence (Ravens Progressive Matrices test³³).

6 **Fatigue Ratings.** In Exp. 1b, participants were asked to rate, between 0 and 100, their level of sub-
7 jective fatigue at various points throughout the experimental session. When doing so, a vertical visual
8 analogue scale (VAS) appeared in the middle of the screen with the words “Not at all” and “Extremely”
9 at its lowest and highest extremes, and the question “How TIRED do you feel?” was displayed at the
10 top.

11 Individuals were instructed to touch the screen to place a cursor at the level that best reflected
12 their current fatigue state. The corresponding rating (a number between 0 and 100) was displayed on
13 the side of the cursor to facilitate the response. For any given rating, the starting position of the cursor
14 was set to the value of the preceding rating. Participants pressed the space bar to validate their response
15 and resume the task.

16 Individuals were asked to rate their fatigue at the beginning, middle and end of each experimental
17 block (i.e. 9 times during the exposure task, then 24 times during the active sampling task). Both raw
18 and z-scored ratings were analysed.

19 **Experimental design.** The average duration of each experiment is provided in Extended Data Fig. 1.
20 Exp. 1 included two separate sessions (Exp. 1a: Passive decision task; Exp. 1b: Active information
21 sampling task), taking place at the same time on consecutive days. The order of the two sessions was
22 fully counterbalanced among gender category. Then, participants underwent 3 blocks of 28 “exposure”
23 trials (84 trials in total) requiring them to place the blue disk for different combinations of reward and
24 uncertainty (average duration: 10.0 minutes, s.d.: 1.7 minutes; performance reported in Supplementary
25 Fig. 1). The aim of this short task was to expose participants to the scoring function so that they could
26 perform the passive choice experiment even on the first session. After the exposure task, participants
27 underwent either 8 blocks of 12 active sampling trials (96 total; active information sampling task) or 6
28 blocks of 28 choice trials (168 total; passive decision task). Both the starting credit ($R_0 = [100, 125]$)

1 cds) and sampling cost ($\eta_s = [1, 5]$ cds/sample) were experimentally manipulated in a two-by-two
 2 design in the active information sampling task in Exp. 1b (24 repetitions of each of the 4 combinations;
 3 Extended Data Fig. 2). In addition, a visual analogue scale appeared on the screen every 14 trials
 4 during the exposure, and every 6 trials during the active information sampling task for participants to
 5 rate how fatigued they felt. The average value during the exposure was used as the baseline measure
 6 of fatigue.

7 In Exps. 2-5, all four experimental variables (R_0 , η_s , η_e , t_{max}) were manipulated in separate
 8 experimental groups. They each included 4 blocks of 24 trials (96 trials in total). Every block fea-
 9 tured six repetitions of each of the four levels of the experimental variable manipulated (Exp. 2:
 10 $R_0 = [50, 75, 100, 125]$ cds, $\eta_s = 3$ cds/sample, $\eta_e = 1.2$ cds/Px, $t_{max} = 18$ secs; Exp. 3:
 11 $R_0 = 100$ cds, $\eta_s = [1, 3, 5, 7]$ cds/sample, $\eta_e = 1.2$ cds/Px, $t_{max} = 18$ secs; Exp. 4: $R_0 =$
 12 100 cds, $\eta_s = 3$ cds/sample, $\eta_e = [0.5, 1.5, 1.9, 2.6]$ cds/Px, $t_{max} = 18$ secs; Exp. 5: $R_0 =$
 13 100 cds, $\eta_s = 3$ cds/sample, $\eta_e = 1.2$ cds/Px, $t_{max} = [6, 10, 14, 18]$ secs; Extended Data Fig. 2).
 14 The four trial types were inter-mixed and the location of the hidden circle was chosen randomly for
 15 each trial.

16 In all experiments, task instructions were delivered at the beginning of the session using an
 17 automated script. This script presented the rules of the game, and included a practice trial with step-
 18 by-step instructions popping up on the screen. The experimenter proceeded to the rest of the testing
 19 if participants demonstrated sufficient understanding of the task (no participant was excluded for not
 20 understanding the task). All participants also completed a debriefing questionnaire at the end of the
 21 experimental session to ensure they understood the task properly.

22 **Quantifying information sampling efficiency.** As illustrated in Fig. 1c, information sampling effi-
 23 ciency was quantified as the rate at which the EE decreased over successive samples. On average, this
 24 relationship was exponential (Extended Data Fig. 4). The information extraction rate, α , was estimated
 25 by fitting the following relationship to individual datasets separately:

$$\begin{aligned} \hat{EE}_{(n,1)} &= \frac{\sum_{i=1}^{96} EE_{(i,1)}}{96} \\ \hat{EE}_{(n,s)} &= (\hat{EE}_{(n,1)} - \hat{EE}_{\infty}) \cdot (1 - \alpha_n)^{s-1} + \hat{EE}_{\infty} \\ 0 &< \alpha < 1 \quad \& \quad \hat{EE}_{\infty} > 0 \end{aligned} \tag{2}$$

1 where n is the trial number, and s is the serial position of the sample within the trial (the 0^{th} sample
2 is the one displayed on the screen at the beginning of the trial, before participants touched the screen;
3 Fig. 1a). This simple model was fitted using a least mean squared error procedure, implemented
4 in MATLAB (The MathWorks inc., version 2018a) with the function *fmincon*. For each participant,
5 parameter estimation yielded one information extraction rate α per trial, and one positive asymptote
6 $\hat{E}E_{\infty}$ for all trials. The former captured how much individual observations reduced uncertainty within
7 a given trial (higher α means more efficient sampling, i.e., sharper reduction of EE). The latter
8 captures the asymptotic limit of information gathering for a given participant.

9 **Computational model of sampling speed, efficiency and extent.** From a neuroeconomic perspec-
10 tive, behaving *optimally* means adopting the course of action – parameterised here as the combination
11 of (ISI, α, s) – that maximises the expected value, EV , of the final response (Fig. 4a). This variable
12 provides a predictive estimate of the reward expected to arise from placing the blue disk. Numerically,
13 it is calculated as:

$$EV_0 = R_0 - \eta_e \cdot EE_0$$

$$EV_s(ISI, \alpha, t_{max}) = EV_{s-1} + p(s|ISI, t_{max}) \cdot [\eta_e \cdot (1 - \alpha) \cdot (EE_{s-1} - \hat{E}E_{\infty}) - \eta_s] \quad (3)$$

14 Previous EV + Probability of acquiring the sample given the current time

. [Expected information benefit – Sampling cost]

15 where the probability of acquiring a sample is 1 during the time allocated to the search, and null
16 otherwise:

$$17 \quad \begin{cases} p(s|ISI, t_{max}) = 1 & \text{if } s \cdot ISI \leq t_{max} \\ p(s|ISI, t_{max}) = 0 & \text{if } s \cdot ISI > t_{max} \end{cases} \quad (4)$$

18 Note that the interaction between economic variables and time pressure is built in the architecture
19 of the model. The time allocated to the search (t_{max}) influences the number of acquirable samples for
20 a given sampling rate, while economic variables (R_0, η_s, η_e) influence on the magnitude of the EV at
21 each step. Thus, the two interact in a multiplicative manner. For example, manipulating time pressure
22 when errors are strongly penalised (i.e., every sample induces a large ΔEV) has a greater influence on
23 utility than when errors are weakly penalised (i.e., every sample induces a small ΔEV). Although this
24 type of interaction was not tested here for practical reasons, it could be done in future studies.

1 The rational EV model was modified as follows to build the models evaluated in this work. Note
 2 that the *EV* is replaced by an Expected Utility (*EU*) because its specification departs from the rational
 3 definition:

$$EU_0 = R_0 - \eta_e \cdot EE_0$$

$$EU_s(ISI, \alpha, t_{max}) = EU_{s-1} + p(s|ISI, t_{max}) \cdot [w_e \cdot \eta_e \cdot (1 - \alpha) \cdot (EE_{s-1} - E\hat{E}_\infty) - w_s \cdot \eta_s^{1+\gamma \cdot s} - \eta_c(ISI, \alpha)]$$

Previous EU + Probability of acquiring the sample given the current time

. [Expected information benefit – Sampling cost – Cognitive effort cost]

4 (5)

5 where w_e and w_s are subjective weights quantifying the contribution of the expected information ben-
 6 efits and sampling cost to the *EU*, γ is a power term introducing a non-linear transformation of the
 7 sampling cost over successive samples, and $H_c(ISI, \alpha)$ is a cognitive effort cost function. The in-
 8 formation gathering asymptotic limit, $E\hat{E}_\infty$, was estimated for each individual by fitting equation 2
 9 beforehand. Thus, the modelling of active information sampling behaviour (ISI, α, s) took into con-
 10 sideration inter-individual differences in asymptotic information gathering performance.

11 Various cognitive effort cost functions were evaluated in the paper:

$$\eta_c(ISI, \alpha) = 0 \quad (\text{No cost})$$

$$\eta_c(ISI, \alpha) = w_0 + w_{speed} \cdot \frac{1}{ISI} + w_\alpha \cdot \alpha \quad (\text{Linear})$$

$$12 \quad \eta_c(ISI, \alpha) = w_0 + w_{speed} \cdot \frac{1}{ISI^2} + w_\alpha \cdot \alpha^2 \quad (\text{Quadratic}) \quad (6)$$

$$\eta_c(ISI, \alpha) = A \cdot e^{w_0 + w_{speed} \cdot \frac{1}{ISI} + w_\alpha \cdot \alpha} \quad (\text{Exponential})$$

$$\eta_c(ISI, \alpha) = \frac{A}{1 + e^{w_0 - w_{speed} \cdot \frac{1}{ISI} - w_\alpha \cdot \alpha}} \quad (\text{Sigmoid})$$

13 where w_{speed} and w_α are subjective weights – estimated through the fitting procedure – quantifying the
 14 contribution of sampling speed and efficiency to the cognitive effort cost. All models reported in the
 15 paper also included a free intercept w_0 (versions without an intercept are reported in *Supplementary*
 16 *Information*; Supplementary Fig. 2). Logistic and exponential cost functions required the addition of
 17 an extra free scaling parameter, A .

18 For each of the linear, quadratic, exponential and sigmoid classes, three versions of the model
 19 were built: 1) only speed penalised ($w_\alpha = 0$), 2) only efficiency penalised ($w_{speed} = 0$), and 3) both

1 speed and efficiency penalised. The sampling cost power term, γ , was only included in the absence
 2 of a cognitive effort cost ($\eta_c(ISI, \alpha) = 0$; Model 2) to act as a second reference model, and ask
 3 whether a non-linear utility function could account for the data without the cognitive effort costs. It was
 4 never used in combination to a linear, quadratic, exponential or logistic cognitive effort cost function
 5 to prevent over-fitting by overly complex models. An exhaustive description of the free parameters
 6 included in each model is provided in Extended Data Fig. 5.

7 The likelihood function was obtained by applying a *softmax* function over the 3-dimensional
 8 space of EU (the EU is a function of ISI, α, s) for a given task environment (combination of $R_0, \eta_s, \eta_e, t_{max}$),
 9 as follows:

$$10 \quad p_s(stop|ISI, \alpha, t_{max}) = \frac{\exp(EU_s(ISI, \alpha, t_{max}))}{\sum_i \sum_a \sum_t^{t_{max}} \exp(EU_s(i, a, t))} \quad (7)$$

11 Model fitting involved finding, for each individual separately, the parameters that minimised the neg-
 12 ative log-likelihood of observing the multivariate distribution of number of samples acquired (s , one
 13 value per trial), inter-sampling interval (ISI , averaged across samples within a trial to obtain one value
 14 per trial) and sampling efficiency (α , one value per trial). Note that sampling speed and efficiency were
 15 assumed to be stationary within a trial (i.e., one combination of $[s, ISI, \alpha]$ per trial). Although it is an
 16 approximation of actual human behaviour, this assumption was necessary to simplify the framework.
 17 Parameter optimisation was carried in MATLAB (The MathWorks inc., version 2018a) using Bayesian
 18 Adaptive Direct Search (BADSS³¹), chosen for its high performance with non-smooth cost functions.
 19 Model fits were compared using Bayesian Information Criterion (BIC), which quantifies goodness of
 20 fit while penalising models for the number of free parameters to avoid over-fitting.

21 **Statistical analysis.** Data handling and statistical analysis were performed in MATLAB (The Math-
 22 Works inc., version 2018a). Mixed-effect models with maximal random structure were used for statis-
 23 tical analysis of behavioural measures. When relevant, “trial index” was included as a random effect to
 24 control for potential learning effects throughout the task (see Supplementary Fig. 4 for a visualisation
 25 of these effects). Extended details (link function, data transformation, model specification, parameter
 26 estimates, statistics) are provided in the tables in *Supplementary Information*. Robust multiple linear
 27 regression was used to investigate the relationship between fatigue development (Block 8 *minus* Base-
 28 line) and the parameters of the winning cognitive effort cost function in Exp. 1b (w_0, w_{speed}, w_α) while
 29 controlling for the experiment duration. This method uses an iteratively re-weighted least squares pro-

cedure to limit the contribution of potential outliers to the linear fit. All statistical tests were two-sided. No *a priori* hypotheses/predictions were altered after the data were analysed or during the course of writing/revising the paper.

Data Availability

Anonymised participant data have been deposited on Open Science Framework, and can be found at <https://osf.io/25wkh/>.

Code Availability

Code for replicating the main results reported the manuscript has been deposited on Open Science Framework, and can be found at <https://osf.io/25wkh/>. Additional analysis codes are available on request from the corresponding author.

References

1. Clark, L., Robbins, T. W., Ersche, K. D. & Sahakian, B. J. Reflection impulsivity in current and former substance users. *Biological Psychiatry* **60**, 515 – 522 (2006). URL <http://www.sciencedirect.com/science/article/pii/S0006322305013971>.
2. Battaglia, P. W. & Schrater, P. R. Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *Journal of Neuroscience* **27**, 6984–6994 (2007). URL <https://www.jneurosci.org/content/27/26/6984>. <https://www.jneurosci.org/content/27/26/6984.full.pdf>.
3. Gottlieb, J., Oudeyer, P. Y., Lopes, M. & Baranes, A. Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences* **17**, 585–593 (2013). URL <http://dx.doi.org/10.1016/j.tics.2013.09.001>.
4. Averbach, B. B. Theory of choice in bandit, information sampling and foraging tasks. *PLoS computational biology* **11**, e1004164 (2015).
5. Juni, M. Z., Gureckis, T. M. & Maloney, L. T. Information sampling behavior with explicit sampling costs: Correction to Juni, Gureckis, and Maloney (2015). *Decision* **3**, 168–168 (2016).

- 1 6. Hauser, T. U., Moutoussis, M., Dayan, P. & Dolan, R. J. Increased decision thresholds trigger
2 extended information gathering across the compulsivity spectrum. *Translational Psychiatry* **7**
3 (2017). URL <http://dx.doi.org/10.1038/s41398-017-0040-3>.
- 4 7. Savage, L. J. *The foundations of statistics* (Courier Corporation, 1972).
- 5 8. Tversky, A. & Kahneman, D. Advances in prospect theory: Cumulative representation of uncer-
6 tainty. *Journal of Risk and Uncertainty* **5**, 297–323 (1992). URL [https://doi.org/10.](https://doi.org/10.1007/BF00122574)
7 1007/BF00122574.
- 8 9. Von Neumann, J., Morgenstern, O. & Kuhn, H. W. *Theory of games and economic behavior*
9 (*commemorative edition*) (Princeton university press, 2007).
- 10 10. Johnson, J. G. & Busemeyer, J. R. Decision making under risk and uncertainty. *Wiley Interdisci-*
11 *plinary Reviews: Cognitive Science* **1**, 736–749 (2010).
- 12 11. Kahneman, D. & Tversky, A. Prospect Theory: An Analysis of Decision Under Risk. In *Handbook*
13 *of the Fundamentals of Financial Decision Making*, vol. Volume 4 of *World Scientific Handbook*
14 *in Financial Economics Series*, 99–127 (WORLD SCIENTIFIC, 2012). URL [https://doi.](https://doi.org/10.1142/9789814417358_{_}0006)
15 [org/10.1142/9789814417358_{_}0006](https://doi.org/10.1142/9789814417358_{_}0006).
- 16 12. Chamberlain, S. R. *et al.* A neuropsychological comparison of obsessive-compulsive disorder and
17 trichotillomania. *Neuropsychologia* **45**, 654–662 (2007).
- 18 13. Furl, N. & Averbeck, B. B. Parietal Cortex and Insula Relate to Evidence Seeking Relevant to
19 Reward-Related Decisions. *Journal of Neuroscience* **31**, 17572–17582 (2011).
- 20 14. Djamshidian, A. *et al.* Decision making, impulsivity, and addictions: Do Parkinson’s disease
21 patients jump to conclusions? *Movement Disorders* **27**, 1137–1145 (2012).
- 22 15. Hauser, T. U. *et al.* Increased decision thresholds enhance information gathering performance in
23 juvenile Obsessive-Compulsive Disorder (OCD). *PLoS Computational Biology* **13**, 1–19 (2017).
- 24 16. Hauser, T. U., Moutoussis, M., Purg, N., Dayan, P. & Dolan, R. J. Beta-Blocker Propranolol Mod-
25 ulates Decision Urgency During Sequential Information Gathering. *The Journal of Neuroscience*
26 **38**, 7170–7178 (2018).
- 27 17. Jones, P. R. *et al.* Efficient visual information sampling develops late in childhood. *Journal of*
28 *Experimental Psychology: General* **148**, 1138–1152 (2019).

18. Hertwig, R., Barron, G., Weber, E. U. & Erev, I. Decisions from experience and the effect of rare events in risky choice. *Psychological Science* **15**, 534–539 (2004). URL <https://doi.org/10.1111/j.0956-7976.2004.00715.x>. PMID: 15270998, <https://doi.org/10.1111/j.0956-7976.2004.00715.x>.
19. Gottlieb, J. & Oudeyer, P. Y. Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience* **19**, 758–770 (2018). URL <http://dx.doi.org/10.1038/s41583-018-0078-0>.
20. Bossaerts, P., Yadav, N. & Murawski, C. Uncertainty and computational complexity. *Philosophical Transactions of the Royal Society B* **374**, 20180138 (2019). URL <https://doi.org/10.1098/rstb.2018.0138>.
21. Kurzban, R., Duckworth, A., Kable, J. W. & Myers, J. An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences* **36**, 661679 (2013).
22. Boureau, Y.-L., Sokol-Hessner, P. & Daw, N. D. Deciding how to decide: Self-control and meta-decision making. *Trends in Cognitive Sciences* **19**, 700 – 710 (2015). URL <http://www.sciencedirect.com/science/article/pii/S1364661315002041>.
23. Shenhav, A., Botvinick, M. & Cohen, J. The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron* **79**, 217 – 240 (2013). URL <http://www.sciencedirect.com/science/article/pii/S0896627313006077>.
24. Westbrook, A. & Braver, T. S. Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience* **15**, 395–415 (2015). URL <https://doi.org/10.3758/s13415-015-0334-y>.
25. Horan, M., Daddaoua, N. & Gottlieb, J. Parietal neurons encode information sampling based on decision uncertainty. *Nature Neuroscience* **22**, 1327–1335 (2019). URL <https://doi.org/10.1038/s41593-019-0440-1>.
26. Phillips, L. D. & Edwards, W. Conservatism in a simple probability inference task. *Journal of experimental psychology* **72**, 346 (1966).
27. Tversky, A. & Edwards, W. Information versus reward in binary choices. *Journal of Experimental Psychology* **71**, 680 (1966).

- 1 28. Dixon, M. L. & Christoff, K. The decision to engage cognitive control is driven by expected
2 reward-value: Neural and behavioral evidence. *PLOS ONE* **7**, 1–12 (2012). URL <https://doi.org/10.1371/journal.pone.0051637>.
3
- 4 29. Shenhav, A., Botvinick, M. & Cohen, J. The expected value of control: An integrative the-
5 ory of anterior cingulate cortex function. *Neuron* **79**, 217 – 240 (2013). URL <http://www.sciencedirect.com/science/article/pii/S0896627313006077>.
6
- 7 30. Manohar, S. G. *et al.* Reward pays the cost of noise reduction in motor and cognitive control. *Cur-*
8 *rent Biology* **25**, 1707–1716 (2015). URL [http://dx.doi.org/10.1016/j.cub.2015.](http://dx.doi.org/10.1016/j.cub.2015.05.038)
9 05.038.
- 10 31. Acerbi, L. & Ma, W. J. Practical bayesian optimization for model fitting with bayesian adaptive di-
11 rect search. In Guyon, I. *et al.* (eds.) *Advances in Neural Information Processing Systems*, vol. 30,
12 1836–1846 (Curran Associates, Inc., 2017). URL [https://proceedings.neurips.cc/](https://proceedings.neurips.cc/paper/2017/file/df0aab058ce179e4f7ab135ed4e641a9-Paper.pdf)
13 [paper/2017/file/df0aab058ce179e4f7ab135ed4e641a9-Paper.pdf](https://proceedings.neurips.cc/paper/2017/file/df0aab058ce179e4f7ab135ed4e641a9-Paper.pdf).
- 14 32. Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U. & Nieuwenhuis, S. The neural basis of the
15 speedaccuracy tradeoff. *Trends in Neurosciences* **33**, 10 – 16 (2010). URL <http://www.sciencedirect.com/science/article/pii/S0166223609001556>.
16
- 17 33. Carpenter, P. A., Just, M. A. & Shell, P. What one intelligence test measures: A theoretical
18 account of the processing in the Raven progressive matrices test. *Psychological Review* **97**, 404–
19 431 (1990).
- 20 34. Ang, Y. S., Lockwood, P., Apps, M. A., Muhammed, K. & Husain, M. Distinct subtypes of apathy
21 revealed by the apathy motivation index. *PLoS ONE* **12**, 1–15 (2017).
- 22 35. Charnov, E. L. *et al.* *Optimal foraging, the marginal value theorem* (Academic Press Inc., 1976).
- 23 36. Kolling, N., Behrens, T. E. J., Mars, R. B. & Rushworth, M. F. S. Neural mechanisms of forag-
24 ing. *Science* **336**, 95–98 (2012). URL [https://science.sciencemag.org/content/](https://science.sciencemag.org/content/336/6077/95)
25 336/6077/95. [https://science.sciencemag.org/content/336/6077/95.](https://science.sciencemag.org/content/336/6077/95.full.pdf)
26 [full.pdf](https://science.sciencemag.org/content/336/6077/95.full.pdf).
- 27 37. Shenhav, A. *et al.* Toward a rational and mechanistic account of mental effort. *An-*
28 *ual Review of Neuroscience* **40**, 99–124 (2017). URL <https://doi.org/10.1146/>

annurev-neuro-072116-031526. PMID: 28375769, <https://doi.org/10.1146/>

annurev-neuro-072116-031526.

38. Payne, J. W., Payne, J. W., Bettman, J. R. & Johnson, E. J. *The adaptive decision maker* (Cambridge university press, 1993).

39. Stark, D. P. & Choplin, J. M. A license to deceive: Enforcing contractual myths despite consumer psychological realities academic article. *New York University Journal of Law & Business* **5**, 617–744 (2009). URL <https://heinonline.org/HOL/P?h=hein.journals/nyujolbu5&i=621>.

40. Otto, A. R. & Daw, N. D. The opportunity cost of time modulates cognitive effort. *Neuropsychologia* **123**, 92 – 105 (2019). URL <http://www.sciencedirect.com/science/article/pii/S0028393218302021>. Cognitive Effort.

41. Hasenstaub, A., Otte, S., Callaway, E. & Sejnowski, T. J. Metabolic cost as a unifying principle governing neuronal biophysics. *Proceedings of the National Academy of Sciences* **107**, 12329–12334 (2010). URL <https://www.pnas.org/content/107/27/12329>. <https://www.pnas.org/content/107/27/12329.full.pdf>.

42. Kitzbichler, M. G., Henson, R. N. A., Smith, M. L., Nathan, P. J. & Bullmore, E. T. Cognitive effort drives workspace configuration of human brain functional networks. *Journal of Neuroscience* **31**, 8259–8270 (2011). URL <https://www.jneurosci.org/content/31/22/8259>. <https://www.jneurosci.org/content/31/22/8259.full.pdf>.

43. Bullmore, E. & Sporns, O. The economy of brain network organization. *Nature Reviews Neuroscience* **13**, 336–349 (2012). URL <https://doi.org/10.1038/nrn3214>.

44. Kostal, L. & Kobayashi, R. Optimal decoding and information transmission in hodgkinhuxley neurons under metabolic cost constraints. *Biosystems* **136**, 3 – 10 (2015). URL <http://www.sciencedirect.com/science/article/pii/S030326471500091X>. Selected papers presented at the Eleventh International Workshop on Neural Coding, Versailles, France, 2014.

45. Hartmanis, J. & Stearns, R. E. On the computational complexity of algorithms. *Transactions of the American Mathematical Society* **117**, 285–306 (1965).

- 1 46. Arora, S. & Barak, B. *Computational complexity: a modern approach* (Cambridge University
2 Press, 2009).
- 3 47. Jones, M. & Sugden, R. Positive confirmation bias in the acquisition of information. *Theory and*
4 *Decision* **50**, 59–99 (2001). URL <https://doi.org/10.1023/A:1005296023424>.
- 5 48. Huys, Q. J. M. *et al.* Bonsai trees in your head: How the pavlovian system sculpts goal-directed
6 choices by pruning decision trees. *PLOS Computational Biology* **8**, 1–13 (2012). URL <https://doi.org/10.1371/journal.pcbi.1002410>.
- 7 49. Krding, K. P. & Wolpert, D. M. Bayesian decision theory in sensorimotor control. *Trends in Cogni-*
8 *tive Sciences* **10**, 319 – 326 (2006). URL [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S1364661306001276)
9 [article/pii/S1364661306001276](http://www.sciencedirect.com/science/article/pii/S1364661306001276). Special issue: Probabilistic models of cognition.
- 10 50. Husain, M. & Roiser, J. P. Neuroscience of apathy and anhedonia: A transdiagnostic approach.
11 *Nature Reviews Neuroscience* **19**, 470–484 (2018). URL [http://dx.doi.org/10.1038/](http://dx.doi.org/10.1038/s41583-018-0029-9)
12 [s41583-018-0029-9](http://dx.doi.org/10.1038/s41583-018-0029-9).
- 13 51. Fleming, S. M. & Lau, H. C. How to measure metacognition. *Frontiers in Human Neuroscience*
14 **8**, 443 (2014). URL [https://www.frontiersin.org/article/10.3389/fnhum.](https://www.frontiersin.org/article/10.3389/fnhum.2014.00443)
15 [2014.00443](https://www.frontiersin.org/article/10.3389/fnhum.2014.00443).
- 16 52. Vaghi, M. M. *et al.* Compulsivity reveals a novel dissociation between action and confidence.
17 *Neuron* **96**, 348 – 354.e4 (2017). URL [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S0896627317308413)
18 [article/pii/S0896627317308413](http://www.sciencedirect.com/science/article/pii/S0896627317308413).
- 19 53. Silvetti, M., Vassena, E., Abrahamse, E. & Verguts, T. Dorsal anterior cingulate-brainstem en-
20 semble as a reinforcement meta-learner. *PLOS Computational Biology* **14**, 1–32 (2018). URL
21 <https://doi.org/10.1371/journal.pcbi.1006370>.
- 22 54. Bonnelle, V., Manohar, S., Behrens, T. & Husain, M. Individual Differences in Premotor Brain
23 Systems Underlie Behavioral Apathy. *Cerebral Cortex* **26**, 807–819 (2016).
- 24 55. Lockwood, P. L. *et al.* Prosocial apathy for helping others when effort is required. *Nature Human*
25 *Behaviour* **1**, 1–10 (2017).
- 26 56. Brainard, D. H. The psychophysics toolbox. *Spatial vision* **10**, 433–436 (1997).
- 27

- 1 57. Kleiner, M., Brainard, D. & Pelli, D. What's new in psychtoolbox-3? (2007).
- 2 58. Chong, T. T. *et al.* Neurocomputational mechanisms underlying subjective valuation of effort
3 costs. *PLoS Biology* **15**, 1–28 (2017). URL [http://dx.doi.org/10.1371/journal.](http://dx.doi.org/10.1371/journal.pbio.1002598)
4 [pbio.1002598](http://dx.doi.org/10.1371/journal.pbio.1002598).
- 5 59. Le Heron, C. *et al.* Dysfunctional effort-based decision-making underlies apathy in genetic cere-
6 bral small vessel disease. *Brain* **141**, 3193–3210 (2018).
- 7 60. Le Heron, C. *et al.* Distinct effects of apathy and dopamine on effort-based decision-making in
8 parkinsons disease. *Brain* **141**, 1455–1469 (2018).
- 9 61. Patton, J. H., Stanford, M. S. & Barratt, E. S. Factor structure of Barratt impulsiveness scale.
10 *Journal of Clinical Psychology* **51**, 768–774 (1995).
- 11 62. Beck, A. T., Ward, C. H., Mendelson, M., Mock, J. & Erbaugh, J. An Inventory for Measuring
12 Depression. *Archives of General Psychiatry* **4**, 561–571 (1961). URL [https://doi.org/](https://doi.org/10.1001/archpsyc.1961.01710120031004)
13 [10.1001/archpsyc.1961.01710120031004](https://doi.org/10.1001/archpsyc.1961.01710120031004).
- 14 63. Buhr, K. & Dugas, M. The intolerance of uncertainty scale: psychometric properties of the en-
15 glish version. *Behaviour Research and Therapy* **40**, 931 – 945 (2002). URL [http://www.](http://www.sciencedirect.com/science/article/pii/S0005796701000924)
16 [sciencedirect.com/science/article/pii/S0005796701000924](http://www.sciencedirect.com/science/article/pii/S0005796701000924).

17 **Acknowledgements** We thank Luigi Acerbi for his help designing an optimal search algorithm (infoMax), and
18 María Serrano Bonilla for her help with figure design. We also thank all members of the Cognitive Neurology
19 Research Group for their assistance as experimental confederates. This work was supported by the Wellcome
20 Trust (206330/Z/17/Z). The funders had no role in study design, data collection and analysis, decision to publish
21 or preparation of the manuscript. P.P. and M.H. were funded by the Wellcome Trust (206330/Z/17/Z). B.A. was
22 funded by a Rhodes scholarship. S.G.M. was funded by an MRC Clinician Scientist Fellowship (MR/P00878/X).

23 **Author Contributions Statement** P.P., B.A. and M.H. designed the study. P.P. and B.A. collected the data.
24 P.P., B.A. and S.G.M. analysed the data. P.P. and M.H. wrote the paper.

25 **Competing Interests Statement** The authors declare no competing interests.

Extended Data Figures

	Exp. 1	Exp. 2	Exp. 3	Exp. 4	Exp. 5
Age	24.5 (3.8)	24.9 (3.8)	25.7 (5.2)	26.1 (4.0)	23.7 (2.6)
Sex	24F/24M	10F/10M	10F/10M	10F/10M	10F/9M
YOE	18.4 (3.2)	18.3 (4.0)	17.1 (3.2)	17.8 (2.1)	17.4 (1.3)
Raven's matrices	8.7 (2.4)	7.5 (2.6)	8.7 (2.0)	7.5 (1.8)	7.6 (2.9)
BDI total score	7.3 (5.5)	7.1 (5.2)	5.4 (5.0)	7.2 (6.5)	7.6 (7.8)
AMI total score	1.4 (0.5)	1.3 (0.4)	1.3 (0.4)	1.3 (0.3)	1.1 (0.4)
BIS total score	62.3 (9.8)	64.8 (10.7)	58.7 (9.1)	63.0 (13.0)	62.6 (13.6)
IUS total score	60.5 (17.3)	60.8 (21.9)	61.7 (20.0)	59.9 (13.0)	66.4 (24.4)
Active task duration (min)	47.3 (1.3)	51.5 (4.1)	50.0 (2.2)	50.4 (2.3)	42.2 (3.0)
Passive task duration (min)	41.7 (2.0)	-	-	-	-

Extended Data Figure 1: **Demographic and questionnaire measures.** The table shows group means with standard deviation in parenthesis. F: Female; M: Male; YOE: years of education (since kindergarten); BDI: Beck Depression Inventory-II; AMI: Apathy Motivation Index; BIS-11: Barratt Impulsiveness Scale-11; IUS: Intolerance of Uncertainty Scale.

	R_0	η_s	η_e	t_{max}	n trials
Exp. 1b	125	1			24
	100	1	1.2	18 secs	24
	125	5			24
	100	5			24
Exp. 2	50				24
	75	3	1.2	18 secs	24
	100				24
	125				24
Exp. 3		1			24
		3	1.2	18 secs	24
	100	5			24
		7			24
Exp. 4			0.5		24
			1.2	18 secs	24
	100	3	1.9		24
			2.6		24
Exp. 5				6 secs	24
				10 secs	24
	100	3	1.2	14 secs	24
				18 secs	24

Extended Data Figure 2: **Experimental design.** All experiments included four types of trial (24 repetitions per trial type; 96 trials in total). The cost-benefit and temporal structure of these trial types is depicted here. R_0 : initial credit; η_s : sampling cost; η_e : error cost; t_{max} : time allocated to the sampling phase.

Extended Data Figure 3: **A visual illustration of efficient versus inefficient sampling.** A map of the posterior belief (top row) and of the expected change in EE (bottom row) is plotted for the two example trials plotted in Fig. 1c (panel a: High α trial; panel b: Low α). The posterior belief map represents (in blue) all pixels that may be the centre of the hidden circle given the information on the screen (grey: search space; purple dots: observations inside the hidden circle; white dots: observations outside the hidden circle). The smaller this area, the lower the overall uncertainty (EE). For each sample, the associated map represents the change in EE expected to occur for every possible candidate sampling location. An efficient sampler systematically samples where the expected ΔEE is the lowest. **a.** During efficient sampling, the area covered by the posterior belief decays because the participant samples where the expected ΔEE is the lowest. **b.** By contrast, when sampling is inefficient, the area covered by the posterior belief decays less sharply because the participant does not sample where the expected ΔEE is the lowest.

Extended Data Figure 4: **On average, the expected error decreases exponentially over successive samples.** Grey lines represent, for each participant ($n = 127$ across all experiments), the average Expected Error as a function of sample index. The black line shows the average across all participants ($\pm$ s.d.).

	w_e	w_s	γ	w_0	w_{speed}	w_α	A	n free parameters
Model 1	✓	✓	x	x	x	x	x	2
Model 2	✓	✓	✓	x	x	x	x	3
Model 3	✓	✓	x	✓	✓	x	x	4
Model 4	✓	✓	x	✓	x	✓	x	4
Model 5	✓	✓	x	✓	✓	✓	x	5
Model 6	✓	✓	x	✓	✓	x	x	4
Model 7	✓	✓	x	✓	x	✓	x	4
Model 8*	✓	✓	x	✓	✓	✓	x	5
Model 9	✓	✓	x	✓	✓	x	✓	5
Model 10	✓	✓	x	✓	x	✓	✓	5
Model 11	✓	✓	x	✓	✓	✓	✓	6
Model 12	✓	✓	x	✓	✓	x	✓	5
Model 13	✓	✓	x	✓	x	✓	✓	5
Model 14	✓	✓	x	✓	✓	✓	✓	6
Model 15	✓	✓	x	x	✓	x	x	3
Model 16	✓	✓	x	x	x	✓	x	3
Model 17	✓	✓	x	x	✓	✓	x	4
Model 18	✓	✓	x	x	✓	x	x	3
Model 19	✓	✓	x	x	x	✓	x	3
Model 20	✓	✓	x	x	✓	✓	x	4
Model 21	✓	✓	x	x	✓	x	✓	4
Model 22	✓	✓	x	x	x	✓	✓	4
Model 23	✓	✓	x	x	✓	✓	✓	5
Model 24	✓	✓	x	x	✓	x	✓	4
Model 25	✓	✓	x	x	x	✓	✓	4
Model 26	✓	✓	x	x	✓	✓	✓	5

Extended Data Figure 5: **Free parameters included in our models of speed, efficiency and extent of active information sampling.** Models were defined using the equation reported in the *Methods*. Tick: parameter included in the model and estimated through the fitting procedure; Cross: parameter not included in the model (effectively set to zero). *: Winning model (Model comparison in Fig. 5, Extended Data Fig. 7 & Supplementary Fig. 2)

Extended Data Figure 6: **Economic and time constraints influence the extent, speed and efficiency of active information sampling.** **a.** Individuals sampled more with large initial reward reserves (Exp. 2), when acquiring data was cheap (Exp. 3), when final error was strongly penalised (Exp. 4), and when they had time to search (Exp. 5). **b-c.** The initial credit and severity of error discounting did not influence how quickly or efficiently participants resolved localisation uncertainty (Exps. 2 & 4). Players sampled slower but more efficiently when acquiring data was more expensive (Exp. 3). By contrast, they speeded up but maintained a high level of efficiency (indexed by information extraction rate) under tighter time pressure (Exp. 5). Points and error bars show group mean $\pm$ s.e.m. Lines show condition averages for each participants. Positive/negative main effect of the experimental variable manipulated is shown in red/blue. Full statistical details are reported in Supplementary Tables 3-5. This figure uses resources from [Freepik.com](https://www.freepik.com).

a. Exp. 2

b. Exp. 3

c. Exp. 4

d. Exp. 5

Extended Data Figure 7: **Model comparison from Experiments 2-5.** The same models as in Fig. 5 are compared here. Model 1 and 2 included no cognitive effort cost (Model 1: rational specification; Model 2: non-linear utility function; in grey) and were used as a reference. The other 12 models included a cognitive effort cost specified in terms of either speed (Models 3, 6, 9, 12), efficiency (Models 4, 7, 10, 13) or both (Models 5, 8, 11, 14). They differed with respect to the shape of the cost function (green: linear; orange: quadratic; blue: exponential; pink: sigmoid). Across all four experiments, participants' choices of sampling speed, efficiency and extent were best explained by a model penalising both speed and efficiency with a quadratic cost function (Model 8, lowest Bayesian Information Criterion, BIC).

Supplementary Information

Supplementary Methods

Quantifying uncertainty with the Expected Error (EE). An ideal observer maintains and sequentially updates a posterior distribution over all hidden circle locations (belief), based on previous knowledge (prior distribution) and new incoming observations (likelihood distribution). Simply put, active inference can be seen as the process of ruling out potential hidden circles from the space of all possible solutions. According to Bayes' rule, the probability of a pixel λ to be the centre of the hidden circle given an observation o (o^+ if the dot is purple, o^- if the dot is white) at the location σ , is the product of the prior probability and the likelihood, normalised by the probability of making this particular observation at this location. Because it is applied sequentially, the posterior probability becomes the prior probability from one sample to the next. Here, s refers to the stage of the search, i.e., the number of samples currently on the screen. Mathematically, sequential Bayesian inference is written as follows:

$$\begin{aligned} p_s(\lambda|o, \sigma) &= \frac{p_s(\lambda) \cdot p_s(o, \sigma|\lambda)}{p_s(o, \sigma)} \\ p_{s+1}(\lambda) &= p_s(\lambda|o, \sigma) \end{aligned} \tag{1}$$

where $p_s(\lambda|o, \sigma)$ is the probability that λ is the centre of the hidden circle, given the observation o at the location σ . Because of the deterministic way information is delivered, the likelihood function is defined as:

$$\left\{ \begin{array}{l} p_s(o^+, \sigma|\lambda) = 1 \text{ if } |\lambda - \sigma| \leq r \\ p_s(o^+, \sigma|\lambda) = 0 \text{ if } |\lambda - \sigma| > r \\ p_s(o^-, \sigma|\lambda) = 1 - p_s(o^+, \sigma|\lambda) \end{array} \right. \tag{2}$$

where r is the radius of the hidden circle. In other words, the likelihood of observing a purple dot (o^+) at a location σ is 1 for pixels within a one radius distance of σ , and zero otherwise. The opposite is true for the likelihood of observing a white dot (o^-). Note that this definition assumes no uncertainty regarding the radius of the hidden circle r . This can be safely assumed because the size of the hidden circle was displayed on the side of the search space at all time (Fig. 7a). Finally, the probability of the observation o at the sampling location σ is the sum over all possible hidden circle centres λ of the

probability of the observation given λ , weighted by the probability of λ to be the hidden circle centre:

$$p_s(o, \sigma) = \sum_{\lambda} p_s(o, \sigma | \lambda) \cdot p_s(\lambda) \quad (3)$$

Sequentially applying Bayes' rule allows an ideal observer to maintain a posterior belief over all locations on the screen (Supplementary Fig. 5a, top row). For every possible circle placement, an expected error can be computed as:

$$EE_s(\lambda) = \sum_i p_s(\lambda_i) \cdot |\lambda - \lambda_i| \quad (4)$$

The EE represents the error that would occur on average after placing the blue disk at the location λ . At any stage of the search, it is minimal at the optimal placement location λ^* , and increase as a function of the distance to λ^* . In the paper, EE refers to the expected error at the optimal placement location ($EE(\lambda^*)$). In the *Supplementary Material*, however, we use the notation EE^* to avoid confusion.

An alternative way to quantify uncertainty is to empirically compute the average placement error observed for any given visual display (as opposed to theoretically calculating it as proposed here). The advantage of this empirical approach is that it does not require making an assumption regarding participants' ability to place the blue circle. However, it is not practical for two reasons. First, it would require an extremely large training dataset to obtain a distribution of placement errors for any given visual display. Second, this method can only be applied to stimuli for which a circle has actually been placed. By contrast, a key advantage of the EE formalism is that it can be applied to any visual display (even those for which a circle has not been placed), thereby providing a method to quantify uncertainty beyond the observed data. Furthermore, we show in the main text that the EE closely resembles confidence ratings (Fig. 2b) and provides a good predictive estimate of actual placement errors (Supplementary Fig. 1).

Optimal search algorithm (Optimal policy index). The objective of the search is to minimise EE^* as much as possible from one sample to the next. The optimal sampling policy is therefore to acquire data where the expected ΔEE^* is minimal. For any candidate sampling location σ , the latter is calculated as:

$$\mathbb{E}[\Delta EE^*](\sigma) = p_s(o^+, \sigma) \cdot (EE_{s+1}^{*+} - EE_s^*) + p_s(o^-, \sigma) \cdot (EE_{s+1}^{*-} - EE_s^*) \quad (5)$$

where EE_{s+1}^{*+} and EE_{s+1}^{*-} are the expected error after a positive or negative observation at the location σ , respectively. Example maps are shown in the middle row of Supplementary Fig. 5a. By definition, the optimal next sampling location is where $\mathbb{E}[\Delta EE^*](\sigma)$ is the lowest.

We used $\mathbb{E}[EE^*](\sigma)$ to compute the probability of an optimal agent to sample at any location on the screen (Supplementary Fig. 5a, bottom row). The latter is simply calculated as a softmax over the whole map, normalised between zero and one. The optimal policy index was the value of this probability at the location where participants actually sampled.

Heuristic search algorithm (Habitual policy index). It was noticed by the researchers running the experimental sessions that participants tended to follow stereotypical search patterns (Supplementary Fig. 6). Typically, they explored a direction radially until a negative sample was encountered, at which point they moved away and searched in another direction. This type of behaviour – non-predicted by the optimal search algorithm – was quantified using two data features: 1) the change in exploration direction occurring from one sample to the next (φ); 2) the distance between successive samples (ρ). Consistent with our empirical observation, the bi-variate distribution of (φ, ρ) was non-uniform, and was influenced by the last observation (Supplementary Fig. 5b). Participants were indeed more likely to explore locally ($\rho < r$) and unidirectionally ($\varphi \sim 0$) when the last sample was located inside the hidden circle, but revert direction ($\varphi \sim \pi$) and move away ($\rho \sim 2r$) when the last sample was negative.

The habitual policy index was calculated, for each sample, as the probability for the participant to be following a rigid exploration strategy (consisting in a stereotypical reaction to the previous observation only), given the observed combination of (φ, ρ) . According to Bayes' rule, this probability is:

$$p_s(\text{Habitual}|\varphi, \rho) = \frac{p(\varphi, \rho|\text{Habitual}) \cdot p_{s-1}(\text{Habitual})}{p_{s-1}(\varphi, \rho)} \quad (6)$$

where:

- $p(\varphi, \rho|\text{Habitual})$ is the value of the participant-specific bi-variate probability distribution function for the combination (φ, ρ) , given the previous observation (e.g. left distribution after purple observation, right distribution after white observation in Fig. 5b);
- $p_{s-1}(\text{Habitual})$ is uniform across all possible combinations of (φ, ρ) observable at the time the

sample was placed;

- $p_{s-1}(\varphi, \rho)$ is the probability of observing this specific combination of (φ, ρ) if the agent was selecting a sample randomly within the search space.

Note that this index could only be calculated for samples in serial position 2 or above because the calculation of φ required at least three dots on the visual display (Supplementary Fig. 5b).

Interested versus disinterested search algorithms. The optimal search algorithm used in this paper provides a solution that resembles an *infoMax* sampler, commonly used in the eye movement literature¹⁻³. The key difference between the two, however, is that our *EE*-based algorithm is concerned not only with reducing the entropy of the solution space (like the *infoMax* search) but also with minimising its anisotropy, which is functionally advantageous given the scoring function of the task. Thus, to use a neuroeconomic terminology⁴, the *EE*-based algorithm can be seen as “interested” (i.e., it is informed by the value of information in the context of the task), while the *infoMax* one is “disinterested” (i.e., it is only concerned with localising the hidden circle). We did not set out to distinguish the contribution of these two search strategies because they yielded very similar predictions in our experiments.

Nevertheless, the active sampling paradigm introduced here could be used in future work to investigate how “interested” the search pattern is. To do so, the scoring function would need to be anisotropic, (e.g., vertical errors are more penalised than horizontal ones). Under these circumstances, reducing the expected error along the most penalised directions should become a priority for an interested agent. By contrast, if the search is disinterested, it should not be affected by this manipulation.

Supplementary Results

Optimal vs. habitual search policies. To better understand the factors underlying time-dependent efficiency benefits, we derived two search policy indices based on participants’ choices of sampling locations. These capture two alternative ways in which players could decide where to touch on the screen. The optimal policy index (OPI) posits that individuals are more likely to sample in regions expected to yield the largest reduction in *EE*. The location of these regions changed dynamically in a non-trivial way as a function of incoming information, such that a new “map” of where best to sample

effectively had to be re-computed at every step of the search. Thus, we expected this strategy to be time consuming, and cognitively demanding. By contrast, a computationally cheaper though less efficient strategy was to follow heuristic exploration rules of the type “*Sample locally and unidirectionally if the last observation was inside the hidden circle, or else, reverse direction and explore away*”. This class of rigid and stereotypical behaviours was captured by the habitual policy index (HPI). Details on how these two indices were computed are provided in the *Supplementary Methods* and Supplementary Fig. 5. Note that these two policies were not mutually exclusive, and that, in some cases, they predicted overlapping sampling locations. Thus, these indices should be interpreted as complementary metrics capturing how optimal choices of sampling locations were (OPI), and whether these choices could also be explained by a simple rigid rule (HPI).

Using a linear mixed-effect model (LMM) with maximal randomness, we found that slower sampling was associated with more optimal choices of sampling locations (Exp. 1-5: all $p < 0.01$), and reduced reliance on habitual policies (Exp. 1-4: $p < 0.04$, Exp. 5: $p = 0.14$; Supplementary Table 12; Supplementary Fig. 7). In other words, when taking more time to think before touching the screen, players tended to make more optimal choices of sampling location, and be less influenced by rigid, habitual behaviour.

These policy indices were also used to better understand the shifts in strategy underlying efficiency adjustments in responses to increases in sampling costs and time pressure. When doing so, they were averaged across the first five samples of each trial (97.69% of trials included at least 5 samples) to avoid confounding the analysis by the number of samples acquired, and to focus on the part of the exploration responsible for most of the EE reduction (Extended Data Fig. 4) and where the use of stereotypical exploration patterns was most evident (Supplementary Fig. 6). In Exp. 1b, the increase in α with higher sampling costs appeared to involve both a closer approximation of the optimal search behaviour (Main effect of η_s on OPI: $t_{(4546)} = 2.52$, $p = 0.012$; Supplementary Table 13; Supplementary Fig. 8a), and the inhibition of habitual sampling actions (Main effect of η_s on HPI: $t_{(4546)} = -3.36$, $p < 0.001$; Supplementary Table 14; Supplementary Fig. 8b). A similar pattern was observed in Exp. 3 when η_s was manipulated in isolation (Supplementary Tables 13 & 14; Supplementary Fig. 8).

In Exp. 5, the overall efficiency benefit appeared to be related to a strong down-weighting of heuristic exploration behaviour all together (Exp. 5 vs. Exp. 2-4: $F_{(1,7356)} = 8.38$, $p = 0.0038$), while

the use of the optimal search policy was unaffected (Exp. 5 vs. Exp. 2-4: $F_{(1,7356)} = 0.94, p = 0.33$; Supplementary Figs. 7 & 8). Similarly, within Exp. 5, shorter decision deadlines were associated with a reduction in habitual sampling (Main effect of t_{max} on HPI: $t_{(1794)} = 2.14, p = 0.033$; Supplementary Table 14; Supplementary Fig. 8b) but had no significant influence on the use of the optimal policy (Main effect of t_{max} on OPI: $t_{(1794)} = -0.42, p = 0.67$; Supplementary Table 13; Supplementary Fig. 8a).

Taken together, the analysis of optimal and habitual policy indices indicate that the mechanisms by which individuals increased sampling efficiency involved a closer approximation of optimal behaviour and/or the inhibition of unhelpful automatic exploration patterns. This is also evident in the raw density maps that show qualitatively less evidence of a cross pattern in Exp. 5 compared to the other experiments (Supplementary Fig. 6).

Additional models of active information sampling. In addition to Models 3-14, a variant in of each model in which no free intercept, w_0 , was included was also tested (Models 15-26; Extended Data Fig. 5). The summed BIC is provided in Supplementary Fig. 2. None of these additional models outperformed Model 8.

Lastly, the following two cognitive effort cost functions were also evaluated to ask whether a simpler version of Model 8 could explain the data better:

$$H_c(ISI, \alpha) = w. \left(\frac{1}{ISI^2} + \alpha^2 \right) \quad (\text{Model 27})$$

$$H_c(ISI, \alpha) = w. \left(\frac{\alpha}{ISI} \right)^2 \quad (\text{Model 28})$$
(7)

None of these cognitive effort cost specifications outperformed Model 8 (Comparison relative to Model 8: $\Delta BIC > 5,500$ in all experiments).

Analysis of sampling extent, speed, and efficiency in Exps. 2-5. The cost-benefit structure of the game was determined objectively by three economic factors, expressed in monetary credits: 1) Starting credit (R_0), 2) Cost of individual sampling actions (sampling cost, η_s), and 3) Severity of penalty for error distance (error cost, η_e). Further, time pressure was manipulated by postponing or advancing

when the final response had to be made (t_{max}), thereby manipulating the opportunity cost (i.e., the instantaneous cost of not sampling). Extended Data Fig. 6 shows the influence of each task parameter on active information sampling in young healthy adults, assessed in four separate experiments (Exps. 2-5: total $n = 79$ participants; see *Methods*; Extended Data Fig. 2).

Consistent with the literature, individuals sampled more extensively when they started with larger reward reserves (Main effect of R_0 on s in Exp. 2: $t_{(1918)} = 4.30$, $p < 0.001$), when acquiring data was cheap (Main effect of η_s on s in Exp. 3: $t_{(1918)} = -9.38$, $p < 0.001$), when final error was strongly penalised (Main effect of η_e on s in Exp. 4: $t_{(1918)} = 5.08$, $p < 0.001$), and when they were not pressured by time (Main effect of t_{max} on s in Exp. 5: $t_{(1822)} = 11.14$, $p < 0.001$; Supplementary Table 3; Extended Data Fig. 6a). This, in turn, influenced the placement error and score obtained (Supplementary Tables 7-8).

The main rationale for developing our paradigm, however, was to characterise active information gathering behaviour beyond mere number of samples acquired. The data show that the way people searched their environment also responded to economic and time pressures (sampling rate and efficiency; Extended Data Fig. 6b-c). When acquiring samples costed more (Exp. 3), information sampling was more efficient as indexed by the information extraction rate (Main effect of η_s on α in Exp. 3: $t_{(1918)} = 3.97$, $p < 0.001$; Supplementary Table 5) but slower as measured by inter-sampling interval (Main effect of η_s on ISI in Exp. 3: $t_{(1918)} = 7.41$, $p < 0.001$; Supplementary Table 4). Thus, at higher costs, participants required fewer samples to obtain the same level of information, but took longer to decide where to sample at each step.

By contrast, under tighter time pressure when there are greater opportunity costs, individuals were able to explore faster (Main effect of t_{max} on ISI in Exp. 5: $t_{(1822)} = 11.71$, $p < 0.001$; Supplementary Table 4) whilst maintaining the highest level of per-sample information gain observed (Main effect of t_{max} on α in Exp. 5: $t_{(1822)} = -0.02$, $p = 0.98$; Exp. 5 vs. Exp 2-4: $F_{(1,7580)} = 4.86$, $p = 0.028$ – note how information extraction rate is maintained at the same level even as inter-sampling interval decreased with shorter available search times; see also Fig. 7b). Finally, the manipulation of reward magnitude (Exp. 2) and severity of error discounting (Exp. 4) did not influence the way participants sampled (all $p > 0.21$; Supplementary Tables 4 & 5; Extended Data Fig. 6b-c), demonstrating that the flexible adjustments reported above were not a mere by-product of the number of samples

obtained.

Individuals sub-optimally give-up reward for information. To determine the degree to which participants under- or over- sampled in the active paradigm, we calculated the rational number of samples to acquire for a given task environment (R_0, η_s, η_e) and sampling efficiency ($\alpha, \hat{E}E_\infty$), using an approach comparable to previous work^{5,6}. This analysis is summarised in Supplementary Fig. 3.

The optimal stopping policy is to terminate the search when the expected value (EV) is maximal (Eq. 3). We calculated the average EV timecourse for every participant given the parameter estimates of their relationship between sample and EE ($\alpha, \hat{E}E_\infty$), and assuming that they sampled fast enough to acquire as many samples as needed during the time allocated to the search (i.e., no opportunity cost; Supplementary Fig. 3a). The probability of an optimal agent to acquire a certain number of samples was calculated as the *softmax* over the EV timecourse. This approach allowed us to predict the full probability distribution of number of samples acquired, as opposed to single point predictions (Supplementary Fig. 3b).

Over- and under-sampling behaviour was quantified as the signed difference in the number of samples acquired relative to the optimal solution, normalised by the optimal standard deviation (z-score; $(s - s^*)/SD^*$). This measure showed clear over-sampling behaviour, particularly in task environment requiring participants to acquire only a few samples (Supplementary Fig. 3c). Note that, by definition, this over-sampling was functionally maladaptive in our task: participants who over-sampled the most obtained the lowest scores (Supplementary Fig. 3d). A questionnaire-based way of reaching the same conclusion is developed in the next section and summarised in Supplementary Table 15 & Supplementary Fig. 9.

Relationships between questionnaire measures and active sampling behaviour. A secondary objective of Exp. 1b was to investigate the relationship between information sampling performance as assessed by our active paradigm, and psychopathological traits as assessed by the questionnaire measures. The total scores of the AMI, BIS-11, IUS, BDI (ln-transformed), and Raven’s matrices were entered in a robust multiple linear regression controlling for age and gender to assess how well they predicted: 1) the average number of samples acquired, 2) the average inter-sampling rate, 3) the average information extraction rate (α), 3) the average uncertainty level (EE) at the end of the search, and

5) the average score. Full statistical details are reported in Supplementary Table 15.

On average, participants who scored higher on the intolerance of uncertainty scale sampled more in the active task ($\beta = 0.38, p = 0.027$), but at the same speed ($\beta = -0.26, p = 0.14$) and efficiency ($\beta = -0.09, p = 0.63$). As a result, they did not decrease uncertainty further (likely because they reached their asymptotic limit of information gathering; $\beta = -0.11, p = 0.53$), and therefore obtained lower scores ($\beta = -0.55, p < 0.001$; Supplementary Fig. 9). A follow-up LMM including the effect of the experimental variables manipulated (R_0 and η_s) showed that intolerance of uncertainty was associated with lower scores across all conditions, but did not significantly interact with the sensitivity to economic constraints (IUS $\times R_0$: $\beta = 0.01, t_{(186)} = 0.19, p = 0.85$; IUS $\times \eta_s$: $\beta = -0.10, t_{(186)} = -1.24, p = 0.22$).

In the present sample, no significant relationship was observed between parameter estimates of the winning model, and any of the questionnaire measures.

References

1. Butko, N. J. & Movellan, J. R. Infomax control of eye movements. *IEEE Transactions on Autonomous Mental Development* **2**, 91–107 (2010).
2. Yang, S. C.-H., Lengyel, M. & Wolpert, D. M. Active sensing in the categorization of visual patterns. *eLife* **5**, e12215 (2016). URL <https://doi.org/10.7554/eLife.12215>.
3. Song, M., Wang, X., Zhang, H. & Li, J. Proactive information sampling in value-based decision-making: Deciding when and where to saccade. *Frontiers in Human Neuroscience* **13**, 35 (2019). URL <https://www.frontiersin.org/article/10.3389/fnhum.2019.00035>.
4. Markant, D. B. & Gureckis, T. M. Does the utility of information influence sampling behavior? *Proceedings of the 34th Annual Conference of the Cognitive Science Society* 719–724 (2012).
5. Juni, M. Z., Gureckis, T. M. & Maloney, L. T. Information sampling behavior with explicit sampling costs: Correction to Juni, Gureckis, and Maloney (2015). *Decision* **3**, 168–168 (2016).
6. Jones, P. R. *et al.* Efficient visual information sampling develops late in childhood. *Journal of Experimental Psychology: General* **148**, 1138–1152 (2019).

Supplementary figures

Supplementary Figure 1: **The Expected Error at the best possible placement provides the best predictive estimate of actual errors.** An Expected Error (EE) can be calculated for any possible circle placement, given the information displayed on the screen. Here, the difference between EE and actual placement error is shown for the worst (left) and best (middle) possible placements. On the right, the EE value used is the average of all possible placements given the information on the screen (all circles of the solution space). On the x-axis, a positive (negative) value means that the uncertainty measure overestimates (underestimates) actual placement errors. These distributions are shown for the exposure task (start of Exps. 1a & 1b) and Exps. 2-5. Across all experiments, the EE at the best possible placement location (middle column, i.e., EE^*) provided the best predictive estimate of actual errors, as illustrated by a distribution of prediction errors centred on zero. It is therefore the measure of uncertainty we used throughout the paper.

Supplementary Figure 3: **Individuals sub-optimally give-up reward for information.** **a.** The expected value (EV) is calculated by replacing the error term by the expected error (EE) in the scoring formula ($EV = R_0 - s \cdot \eta_s - EE \cdot \eta_e$; Fig. 1). Its timecourse over successive samples is shown for every experimental condition and every participant. **b.** Predictions of the optimal model were generated by computing the softmax over the EV timecourse for each condition and participant (solid distribution). For comparison, individual and group distributions are shown as thin and thick lines respectively. Overall, participants tended to acquire more samples than required by the optimal solution, especially in task environments requiring to tolerate uncertainty (high η_s , low η_e). **c.** Deviations from the optimal solution were quantified as a z-score: $(s - s^*)/SD^*$ where s^* and SD^* are the optimal mean and standard deviation of the number of samples to acquire. Error bar: bootstrapped 95% confidence interval. **d.** By definition, departures from optimal stopping behaviour had negative consequences on reward. Lines show the individual EV timecourses (same as a.), centred on their peaks (optimal number of samples to acquire). Ellipses contain 90% confidence interval of the means for each participant.

Supplementary Figure 4: **Active information sampling behaviour shows evidence of non-stationarity.** The three parameters of interest (number of samples acquired, inter-sampling interval, information extraction rate) are plotted as a function of task progression (trial bins containing 25% of the data each). Points indicate group means and error bars show the s.e.m.

Supplementary Figure 5: **Sampling policy indices.** **a. Optimal policy index.** Throughout the search, an ideal observer sequentially updates a posterior belief over all possible centres of the hidden circle (blue pixels). Search efficiency is maximised by sampling where the expected ΔEE^* is lowest (yellow regions). The probability of an ideal observer to sample at any given location (bottom row) was obtained by applying a softmin function to the map of expected ΔEE^* (middle row). **b. Habitual policy index.** The bi-variate distribution of change in sampling direction (φ , wrapped between 0 and π) and distance to previous sample (ρ) differed as a function of the outcome of the last sample (left distribution: purple dot) or white (right distribution: white dot). This is illustrated here for all participants of Exp. 2-4 pooled together. Participants tended to explore locally (small φ and ρ) when the last sample was inside the hidden circle, and explore away (large φ and ρ) otherwise. This feature of the data was used to derive the habitual policy index.

Supplementary Figure 6: **Raw sampling density maps.** Raw density maps pooled across all participants show clear presence of heuristic exploration patterns (cross) in Exps. 1-4 (manipulation of economic variables) but not in Exp. 5 (manipulation of time pressure). Coordinates are relative to the observation presented at the beginning of the trial (centre of the map). Red and white circles correspond to one and two r distance to the zeroth sample, respectively. Yellow indicates higher density.

Supplementary Figure 7: **Better search algorithms take longer to run. a.** In all experiments, people chose where to sample more optimally when taking their time (Optimal Policy Index). Conversely, the use of habitual exploration rules increased with faster sampling (Habitual Policy Index), except in Exp. 5 where individuals showed very little heuristic behaviour overall. Inter-sampling time is z-scored per-subject and per-condition to remove unrelated individual and experimental effects. The habitual and optimal policy indices are averaged across the first 5 samples for visualisation purposes. These conditional plots were generated by sliding, for each subject separately, 25% quantile bins of z-scored ISI and computing the mean policy index for each bin. The error bar shows the s.e.m. See also Fig. 5-6.

Supplementary Figure 8: **Effect of experimental variables on the search policies guiding active sampling.** The Optimal Policy Index (**a.**) and Habitual Policy Index (**b.**) – averaged across the first five samples acquired – is plotted as a function the task variable manipulated in each experiment. Points indicate group means and error bars show the bootstrapped 95% confidence interval of the mean.

Supplementary Figure 9: **Intolerance of uncertainty had a negative effect on score (Exp. 5).** The plot shows partial correlation, controlling for all covariates of Supplementary Table 15 (AMI total score, BIS-11 total score, ln-transformed BDI score, Raven’s matrices score, age, gender). Black intensity shows the contribution of individual data points to the robust multiple linear regression (multivariate outliers are down-weighted).

Supplementary tables

	Uncertainty Score	Choices
(Intercept)	$\beta = +8.27 \times 10^{-17}$ $SE = 0.0149$ $t_{8062} = +0.00$ $p = 1.00$	$\beta = +0.356$ $SE = 0.286$ $t_{8060} = +1.24$ $p = 0.21$
<i>EE</i>	$\beta = +0.799$ $SE = 0.016$ $t_{8062} = +49.89$ $p < 0.0001$	$\beta = -2.53$ $SE = 0.135$ $t_{8060} = -18.77$ $p < 0.0001$
<i>R</i>		$\beta = +1.87$ $SE = 0.196$ $t_{8060} = +9.58$ $p < 0.0001$
<i>R:EE</i>		$\beta = +0.101$ $SE = 0.102$ $t_{8060} = +0.99$ $p = 0.32$
<i>adj - R</i> ²	0.69	1.00
<i>N</i> _{obs}	8064	8064
BIC	14046.98	48710.74

Supplementary Table 1: **Fixed effects of the analysis of uncertainty score and accept-reject decisions.** Models were specified as follows. Uncertainty Score $\sim 1 + EE + (1 + EE \mid \text{participant}) + (1 + EE \mid \text{Session Order}) + (1 \mid \text{trial})$; Choice $\sim 1 + R*EE + (1 + R*EE \mid \text{participant}) + (1 + R*EE \mid \text{Session Order}) + (1 \mid \text{trial})$. A linear model was used for uncertainty scores, while a logistic one was used for choices (coded as 1 = accept, 0 = reject). To obtain standardised coefficient estimates, *R* and *EE* were z-scored.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +0.214$ $SE = 0.00675$ $t_{4606} = +31.73$ $p < 0.0001$	$\beta = +0.205$ $SE = 0.00689$ $t_{1918} = +29.69$ $p < 0.0001$	$\beta = +0.214$ $SE = 0.00793$ $t_{1918} = +26.98$ $p < 0.0001$	$\beta = +0.209$ $SE = 0.00806$ $t_{1918} = +25.90$ $p < 0.0001$	$\beta = +0.263$ $SE = 0.00902$ $t_{1822} = +29.15$ $p < 0.0001$
<i>lnISI</i>	$\beta = +0.0762$ $SE = 0.0075$ $t_{4606} = +10.16$ $p < 0.0001$	$\beta = +0.0988$ $SE = 0.00931$ $t_{1918} = +10.62$ $p < 0.0001$	$\beta = +0.0639$ $SE = 0.0127$ $t_{1918} = +5.02$ $p < 0.0001$	$\beta = +0.0963$ $SE = 0.0127$ $t_{1918} = +7.60$ $p < 0.0001$	$\beta = +0.0695$ $SE = 0.0102$ $t_{1822} = +6.81$ $p < 0.0001$
<i>adj - R</i> ²	0.22	0.28	0.21	0.26	0.11
<i>N</i> _{obs}	4608	1920	1920	1920	1824
BIC	-10743.20	-4277.37	-4285.17	-4555.20	-3698.71

Supplementary Table 2: **Fixed effects of the relationship between inter-sampling interval and information extraction rate.** For each experiment, the linear mixed-effect model was specified as follows: $\alpha \sim 1 + ISI + (1 \mid \text{trial}) + (1 + ISI \mid \text{Exp. cond.}) + (1 + ISI \mid \text{participant})$. ISI : Inter-sampling interval.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +2.31$ $SE = 0.0287$ $t_{4604} = +80.53$ $p < 0.0001$	$\beta = +2.33$ $SE = 0.0741$ $t_{1918} = +31.45$ $p < 0.0001$	$\beta = +2.32$ $SE = 0.0484$ $t_{1918} = +47.98$ $p < 0.0001$	$\beta = +2.41$ $SE = 0.0492$ $t_{1918} = +48.92$ $p < 0.0001$	$\beta = +2.38$ $SE = 0.0503$ $t_{1822} = +47.26$ $p < 0.0001$
R_0	$\beta = +0.0392$ $SE = 0.00554$ $t_{4604} = +7.08$ $p < 0.0001$	$\beta = +0.0339$ $SE = 0.00788$ $t_{1918} = +4.30$ $p < 0.0001$			
η_s	$\beta = -0.191$ $SE = 0.0102$ $t_{4604} = -18.64$ $p < 0.0001$		$\beta = -0.209$ $SE = 0.0223$ $t_{1918} = -9.38$ $p < 0.0001$		
$\eta_s:R_0$	$\beta = +0.00784$ $SE = 0.00685$ $t_{4604} = +1.14$ $p = 0.25$				
η_e				$\beta = +0.0871$ $SE = 0.0171$ $t_{1918} = +5.08$ $p < 0.0001$	
t_{max}					$\beta = +0.14$ $SE = 0.0125$ $t_{1822} = +11.14$ $p < 0.0001$
$adj - R^2$	0.67	0.69	0.68	0.54	0.57
N_{obs}	4608	1920	1920	1920	1824
BIC	288.78	239.07	234.35	174.31	203.75

Supplementary Table 3: **Fixed effect of the analysis of task variables on the number of samples acquired.** Linear mixed-effect models all used a Poisson distribution, and were specified as follows. Exp. 1b: $\text{sample} \sim 1 + \eta_s * R_0 + (1 + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $\text{sample} \sim 1 + R_0 + (1 + R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $\text{sample} \sim 1 + \eta_s + (1 + \eta_s \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $\text{sample} \sim 1 + \eta_e + (1 + \eta_e \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $\text{sample} \sim 1 + t_{max} + (1 + t_{max} \mid \text{participant}) + (1 \mid \text{trial})$.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +0.25$ $SE = 0.025$ $t_{4604} = +9.97$ $p < 0.0001$	$\beta = +0.282$ $SE = 0.0528$ $t_{1918} = +5.34$ $p < 0.0001$	$\beta = +0.311$ $SE = 0.0477$ $t_{1918} = +6.53$ $p < 0.0001$	$\beta = +0.185$ $SE = 0.0418$ $t_{1918} = +4.42$ $p < 0.0001$	$\beta = -0.123$ $SE = 0.0485$ $t_{1822} = -2.54$ $p = 0.011$
R_0	$\beta = -0.0184$ $SE = 0.00468$ $t_{4604} = -3.94$ $p < 0.0001$	$\beta = -0.00383$ $SE = 0.00442$ $t_{1918} = -0.87$ $p = 0.39$			
η_s	$\beta = +0.122$ $SE = 0.00883$ $t_{4604} = +13.83$ $p < 0.0001$		$\beta = +0.127$ $SE = 0.0171$ $t_{1918} = +7.41$ $p < 0.0001$		
$\eta_s:R_0$	$\beta = -0.00145$ $SE = 0.00614$ $t_{4604} = -0.24$ $p = 0.81$				
η_e				$\beta = -0.00924$ $SE = 0.00745$ $t_{1918} = -1.24$ $p = 0.21$	
t_{max}					$\beta = +0.18$ $SE = 0.0154$ $t_{1822} = +11.71$ $p < 0.0001$
$adj - R^2$	0.59	0.62	0.65	0.53	0.63
N_{obs}	4608	1920	1920	1920	1824
BIC	-1333.54	-670.99	-369.33	-876.64	-30.15

Supplementary Table 4: **Fixed effect of the analysis of task variables on the inter-sampling interval (ln-transformed)**. Linear mixed-effect models were specified as follows. Exp. 1b: $\ln ISI \sim 1 + \eta_s * R_0 + (1 + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $\ln ISI \sim 1 + R_0 + (1 + R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $\ln ISI \sim 1 + \eta_s + (1 + \eta_s \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $\ln ISI \sim 1 + \eta_e + (1 + \eta_e \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $\ln ISI \sim 1 + t_{max} + (1 + t_{max} \mid \text{participant}) + (1 \mid \text{trial})$.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = -1.51$ $SE = 0.0248$ $t_{4604} = -60.97$ $p < 0.0001$	$\beta = -1.52$ $SE = 0.0418$ $t_{1918} = -36.45$ $p < 0.0001$	$\beta = -1.5$ $SE = 0.0349$ $t_{1918} = -42.90$ $p < 0.0001$	$\beta = -1.54$ $SE = 0.039$ $t_{1918} = -39.57$ $p < 0.0001$	$\beta = -1.43$ $SE = 0.0216$ $t_{1822} = -66.02$ $p < 0.0001$
R_0	$\beta = +0.00228$ $SE = 0.00632$ $t_{4604} = +0.36$ $p = 0.72$	$\beta = -0.00104$ $SE = 0.00714$ $t_{1918} = -0.15$ $p = 0.88$			
η_s	$\beta = +0.0173$ $SE = 0.00681$ $t_{4604} = +2.54$ $p = 0.011$		$\beta = +0.0286$ $SE = 0.0072$ $t_{1918} = +3.97$ $p < 0.0001$		
$\eta_s:R_0$	$\beta = -0.00966$ $SE = 0.00776$ $t_{4604} = -1.24$ $p = 0.21$				
η_e				$\beta = +0.0065$ $SE = 0.00707$ $t_{1918} = +0.92$ $p = 0.36$	
t_{max}					$\beta = -0.000163$ $SE = 0.00799$ $t_{1822} = -0.02$ $p = 0.98$
$adj - R^2$	0.19	0.26	0.21	0.24	0.07
N_{obs}	4608	1920	1920	1920	1824
BIC	3978.99	1105.14	958.39	985.56	994.65

Supplementary Table 5: **Fixed effect of the analysis of task variables on the information extraction rate (ln-transformed).** Linear mixed-effect models were specified as follows. Exp. 1b: $\ln \alpha \sim 1 + \eta_s * R_0 + (1 + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $\ln \alpha \sim 1 + R_0 + (1 + R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $\ln \alpha \sim 1 + \eta_s + (1 + \eta_s \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $\ln \alpha \sim 1 + \eta_e + (1 + \eta_e \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $\ln \alpha \sim 1 + t_{max} + (1 + t_{max} \mid \text{participant}) + (1 \mid \text{trial})$.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = -1.6$ $SE = 0.033$ $t_{4603} = -48.52$ $p < 0.0001$	$\beta = -1.64$ $SE = 0.0344$ $t_{1917} = -47.53$ $p < 0.0001$	$\beta = -1.59$ $SE = 0.0395$ $t_{1917} = -40.34$ $p < 0.0001$	$\beta = -1.61$ $SE = 0.0371$ $t_{1917} = -43.38$ $p < 0.0001$	$\beta = -1.4$ $SE = 0.0205$ $t_{1821} = -68.25$ $p < 0.0001$
$\ln ISI$	$\beta = +0.339$ $SE = 0.0487$ $t_{4603} = +6.96$ $p < 0.0001$	$\beta = +0.413$ $SE = 0.041$ $t_{1917} = +10.07$ $p < 0.0001$	$\beta = +0.261$ $SE = 0.0625$ $t_{1917} = +4.18$ $p < 0.0001$	$\beta = +0.375$ $SE = 0.0509$ $t_{1917} = +7.36$ $p < 0.0001$	$\beta = +0.241$ $SE = 0.0367$ $t_{1821} = +6.56$ $p < 0.0001$
R_0	$\beta = +0.0075$ $SE = 0.00611$ $t_{4603} = +1.23$ $p = 0.22$	$\beta = +0.000307$ $SE = 0.00716$ $t_{1917} = +0.04$ $p = 0.97$			
η_s	$\beta = -0.022$ $SE = 0.00703$ $t_{4603} = -3.13$ $p = 0.0018$		$\beta = -0.006$ $SE = 0.00994$ $t_{1917} = -0.60$ $p = 0.55$		
$\eta_s:R_0$	$\beta = -0.00747$ $SE = 0.00624$ $t_{4603} = -1.20$ $p = 0.23$				
η_e				$\beta = +0.0112$ $SE = 0.00718$ $t_{1917} = +1.56$ $p = 0.12$	
t_{max}					$\beta = -0.0445$ $SE = 0.0115$ $t_{1821} = -3.86$ $p = 0.00012$
$adj - R^2$	0.23	0.31	0.25	0.29	0.11
N_{obs}	4608	1920	1920	1920	1824
BIC	3779.78	1008.56	904.87	897.76	940.12

Supplementary Table 6: **Fixed effect of the analysis of task variables on information extraction rate (ln-transformed), controlling for inter-sampling interval (ln-transformed).** Linear mixed-effect models were specified as follows. Exp. 1b: $\ln \alpha \sim 1 + \ln ISI + \eta_s * R_0 + (1 + \ln ISI + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $\ln \alpha \sim 1 + R_0 + \ln ISI + (1 + R_0 + \ln ISI \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $\ln \alpha \sim 1 + \eta_s + \ln ISI + (1 + \eta_s + \ln ISI \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $\ln \alpha \sim 1 + \eta_e + \ln ISI + (1 + \eta_e + \ln ISI \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $\ln \alpha \sim 1 + t_{max} + \ln ISI + (1 + t_{max} + \ln ISI \mid \text{participant}) + (1 \mid \text{trial})$.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +2.51$ $SE = 0.0318$ $t_{4604} = +78.86$ $p < 0.0001$	$\beta = +2.48$ $SE = 0.0809$ $t_{1918} = +30.62$ $p < 0.0001$	$\beta = +2.5$ $SE = 0.0554$ $t_{1918} = +45.20$ $p < 0.0001$	$\beta = +2.33$ $SE = 0.059$ $t_{1918} = +39.50$ $p < 0.0001$	$\beta = +2.56$ $SE = 0.0461$ $t_{1822} = +55.39$ $p < 0.0001$
R_0	$\beta = -0.0234$ $SE = 0.0122$ $t_{4604} = -1.92$ $p = 0.05$	$\beta = -0.0229$ $SE = 0.0157$ $t_{1918} = -1.46$ $p = 0.15$			
η_s	$\beta = +0.194$ $SE = 0.0156$ $t_{4604} = +12.39$ $p < 0.0001$		$\beta = +0.212$ $SE = 0.0266$ $t_{1918} = +7.97$ $p < 0.0001$		
$\eta_s:R_0$	$\beta = -0.014$ $SE = 0.0123$ $t_{4604} = -1.13$ $p = 0.26$				
η_e				$\beta = -0.0688$ $SE = 0.0197$ $t_{1918} = -3.50$ $p = 0.00048$	
t_{max}					$\beta = -0.181$ $SE = 0.0189$ $t_{1822} = -9.59$ $p < 0.0001$
$adj - R^2$	0.17	0.24	0.21	0.12	0.13
N_{obs}	4608	1920	1920	1920	1824
BIC	9601.57	3924.38	3999.90	4130.15	3783.42

Supplementary Table 7: **Fixed effect of the analysis of task variables on the final placement error (ln-transformed)**. Linear mixed-effect models were specified as follows. Exp. 1b: $\ln Error \sim 1 + \eta_s * R_0 + (1 + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $\ln Error \sim 1 + R_0 + (1 + R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $\ln Error \sim 1 + \eta_s + (1 + \eta_s \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $\ln Error \sim 1 + \eta_e + (1 + \eta_e \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $\ln Error \sim 1 + t_{max} + (1 + t_{max} \mid \text{participant}) + (1 \mid \text{trial})$.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +67.1$ $SE = 0.51$ $t_{4604} = +131.53$ $p < 0.0001$	$\beta = +37.5$ $SE = 1.46$ $t_{1918} = +25.61$ $p < 0.0001$	$\beta = +44.2$ $SE = 1.04$ $t_{1918} = +42.55$ $p < 0.0001$	$\beta = +47.2$ $SE = 0.85$ $t_{1918} = +55.49$ $p < 0.0001$	$\beta = +48$ $SE = 1.19$ $t_{1822} = +40.23$ $p < 0.0001$
R_0	$\beta = +11.7$ $SE = 0.234$ $t_{4604} = +49.92$ $p < 0.0001$	$\beta = +27.5$ $SE = 0.317$ $t_{1918} = +86.76$ $p < 0.0001$			
η_s	$\beta = -18.6$ $SE = 0.448$ $t_{4604} = -41.46$ $p < 0.0001$		$\beta = -19.6$ $SE = 0.46$ $t_{1918} = -42.65$ $p < 0.0001$		
$\eta_s:R_0$	$\beta = -0.435$ $SE = 0.218$ $t_{4604} = -1.99$ $p = 0.046$				
η_e				$\beta = -11$ $SE = 0.664$ $t_{1918} = -16.63$ $p < 0.0001$	
t_{max}					$\beta = -0.841$ $SE = 0.465$ $t_{1822} = -1.81$ $p = 0.07$
$adj - R^2$	0.71	0.81	0.57	0.32	0.12
N_{obs}	4608	1920	1920	1920	1824
BIC	37852.29	15625.15	16563.14	16451.53	14870.46

Supplementary Table 8: **Fixed effect of the analysis of task variables on the score.** Linear mixed-effect models were specified as follows. Exp. 1b: $\text{Score} \sim 1 + \eta_s * R_0 + (1 + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $\text{Score} \sim 1 + R_0 + (1 + R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $\text{Score} \sim 1 + \eta_s + (1 + \eta_s \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $\text{Score} \sim 1 + \eta_e + (1 + \eta_e \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $\text{Score} \sim 1 + t_{max} + (1 + t_{max} \mid \text{participant}) + (1 \mid \text{trial})$.

	Raw ratings	Z-scored ratings
(Intercept)	$\beta = +28.9$ $SE = 3.07$ $t_{1150} = +9.42$ $p < 0.0001$	$\beta = -0.537$ $SE = 0.0976$ $t_{1150} = -5.50$ $p < 0.0001$
block	$\beta = +0.873$ $SE = 0.21$ $t_{1150} = +4.16$ $p < 0.0001$	$\beta = +0.153$ $SE = 0.0259$ $t_{1150} = +5.90$ $p < 0.0001$
$adj - R^2$	0.98	0.34
N_{obs}	1152	1152

Supplementary Table 9: **Fixed effects of the analysis of fatigue development in Exp. 1.** For each column, the linear mixed-effect model was specified as follows: $\text{Fatigue} \sim 1 + \text{block} + (1 + \text{block} \mid \text{participant})$. The first model used raw fatigue ratings, while the second used z-scored fatigue ratings as the dependent variable.

	Raw ratings	Z-scored ratings
(Intercept)	$\beta = -585$ $SE = 223$ $t_{43} = -2.63$ $p = 0.012$	$\beta = -40.8$ $SE = 32.4$ $t_{43} = -1.26$ $p = 0.21$
Exp. Duration	$\beta = +145$ $SE = 58.1$ $t_{43} = +2.50$ $p = 0.016$	$\beta = +10$ $SE = 8.45$ $t_{43} = +1.19$ $p = 0.24$
w_0	$\beta = -16.4$ $SE = 6.98$ $t_{43} = -2.35$ $p = 0.023$	$\beta = -1.53$ $SE = 1.02$ $t_{43} = -1.51$ $p = 0.14$
w_α	$\beta = +8.86$ $SE = 2.83$ $t_{43} = +3.13$ $p = 0.0032$	$\beta = +0.87$ $SE = 0.412$ $t_{43} = +2.11$ $p = 0.041$
w_{speed}	$\beta = -5.99$ $SE = 3.27$ $t_{43} = -1.83$ $p = 0.07$	$\beta = -0.712$ $SE = 0.476$ $t_{43} = -1.50$ $p = 0.14$
$adj - R^2$	0.31	0.14
N_{obs}	48	48

Supplementary Table 10: **Robust multiple linear regressions of the relationship between task-induced fatigue and model-derived cognitive effort cost.** For each column, the robust multiple linear regression was specified as follows: $\Delta\text{Fatigue} \sim 1 + \ln \text{Exp.Duration} + w_0 + \ln w_\alpha + \ln w_{speed}$. ln-transforms were applied to obtain a more normal distribution of the independent variables. The first model used raw fatigue ratings, while the second used z-scored fatigue ratings as the dependent variable. The Δ symbol refers to the difference between the average rating within the last block of active search and the 3 blocks of exposure (used as a measure of baseline fatigue).

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +4.17$ $SE = 1.2$ $t_{43916} = +3.48$ $p = 0.00049$	$\beta = +4.47$ $SE = 1.15$ $t_{18790} = +3.89$ $p = 0.0001$	$\beta = +3.69$ $SE = 1.17$ $t_{18610} = +3.15$ $p = 0.0016$	$\beta = +3.89$ $SE = 1.33$ $t_{20027} = +2.93$ $p = 0.0033$	$\beta = +3.36$ $SE = 1.07$ $t_{18509} = +3.14$ $p = 0.0017$
$EE_{(s-1)}$	$\beta = -0.355$ $SE = 0.0261$ $t_{43916} = -13.63$ $p < 0.0001$	$\beta = -0.381$ $SE = 0.0306$ $t_{18790} = -12.46$ $p < 0.0001$	$\beta = -0.359$ $SE = 0.0327$ $t_{18610} = -10.98$ $p < 0.0001$	$\beta = -0.352$ $SE = 0.0335$ $t_{20027} = -10.51$ $p < 0.0001$	$\beta = -0.277$ $SE = 0.0256$ $t_{18509} = -10.83$ $p < 0.0001$
$\ln(t_{(s)} - t_{(s-1)})$	$\beta = -0.735$ $SE = 0.323$ $t_{43916} = -2.27$ $p = 0.023$	$\beta = -1.03$ $SE = 0.348$ $t_{18790} = -2.96$ $p = 0.003$	$\beta = -0.647$ $SE = 0.303$ $t_{18610} = -2.14$ $p = 0.033$	$\beta = -1.33$ $SE = 0.462$ $t_{20027} = -2.88$ $p = 0.0039$	$\beta = -0.593$ $SE = 0.169$ $t_{18509} = -3.51$ $p = 0.00046$
$adj - R^2$	0.41	0.38	0.41	0.38	0.46
N_{obs}	43919	18793	18613	20030	18512
BIC	310197.46	134366.80	132690.67	143178.89	124277.57

Supplementary Table 11: **Fixed effects of the relationship between pre-sampling time and information gain.** For each experiment, the linear mixed-effect model was specified as follows: $\Delta EE \sim 1 + EE_{(s-1)} + \ln(t_{(s)} - t_{(s-1)}) + (1 + EE_{(s-1)} + \ln(t_{(s)} - t_{(s-1)}) \mid \text{sample}) + (1 + EE_{(s-1)} + \ln(t_{(s)} - t_{(s-1)}) \mid \text{trial}) + (1 + EE_{(s-1)} + \ln(t_{(s)} - t_{(s-1)}) \mid \text{Exp. cond.}) + (1 + EE_{(s-1)} + \ln(t_{(s)} - t_{(s-1)}) \mid \text{participant})$. $\Delta EE = EE_{(s)} - EE_{(s-1)}$

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +0.116$ $SE = 0.0772$ $t_{43914} = +1.51$ $p = 0.13$	$\beta = +0.185$ $SE = 0.0573$ $t_{18788} = +3.23$ $p = 0.0013$	$\beta = +0.1$ $SE = 0.0883$ $t_{18610} = +1.14$ $p = 0.26$	$\beta = +0.0811$ $SE = 0.0464$ $t_{20027} = +1.75$ $p = 0.08$	$\beta = -0.226$ $SE = 0.118$ $t_{18509} = -1.91$ $p = 0.06$
HPI	$\beta = -1.87$ $SE = 0.487$ $t_{43914} = -3.84$ $p = 0.00012$	$\beta = -1.76$ $SE = 0.518$ $t_{18788} = -3.40$ $p = 0.00067$	$\beta = -1.99$ $SE = 0.96$ $t_{18610} = -2.07$ $p = 0.038$	$\beta = -1.7$ $SE = 0.376$ $t_{20027} = -4.52$ $p < 0.0001$	$\beta = -1.86$ $SE = 1.25$ $t_{18509} = -1.49$ $p = 0.14$
OPI	$\beta = +0.0439$ $SE = 0.0154$ $t_{43914} = +2.86$ $p = 0.0043$	$\beta = +0.05$ $SE = 0.016$ $t_{18788} = +3.13$ $p = 0.0018$	$\beta = +0.0662$ $SE = 0.0125$ $t_{18610} = +5.30$ $p < 0.0001$	$\beta = +0.079$ $SE = 0.023$ $t_{20027} = +3.44$ $p = 0.00059$	$\beta = +0.0727$ $SE = 0.0257$ $t_{18509} = +2.83$ $p = 0.0047$
$adj - R^2$	0.28	0.34	0.35	0.22	0.35
N_{obs}	43917	18791	18613	20030	18512
BIC	44070.58	16754.25	19666.95	17023.11	19602.99

Supplementary Table 12: **Fixed effects of the relationship between pre-sampling time and sampling policies.** For each experiment, the linear mixed-effect model was specified as follows: $\ln(t_{(s)} - t_{(s-1)}) \sim 1 + \text{HPI} + \text{OPI} + (1 + \text{HPI} + \text{OPI} \mid \text{sample}) + (1 + \text{HPI} + \text{OPI} \mid \text{trial}) + (1 + \text{HPI} + \text{OPI} \mid \text{Exp. cond.}) + (1 + \text{HPI} + \text{OPI} \mid \text{participant})$. HPI: Habitual Policy Index; OPI: Optimal Policy Index.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +0.169$ $SE = 0.0079$ $t_{4546} = +21.42$ $p < 0.0001$	$\beta = +0.163$ $SE = 0.0139$ $t_{1821} = +11.72$ $p < 0.0001$	$\beta = +0.194$ $SE = 0.0124$ $t_{1864} = +15.62$ $p < 0.0001$	$\beta = +0.158$ $SE = 0.0129$ $t_{1873} = +12.21$ $p < 0.0001$	$\beta = +0.185$ $SE = 0.00641$ $t_{1794} = +28.88$ $p < 0.0001$
R_0	$\beta = -0.00339$ $SE = 0.00169$ $t_{4546} = -2.01$ $p = 0.045$	$\beta = -0.00328$ $SE = 0.00265$ $t_{1821} = -1.24$ $p = 0.22$			
η_s	$\beta = +0.00759$ $SE = 0.00301$ $t_{4546} = +2.52$ $p = 0.012$		$\beta = +0.0119$ $SE = 0.00335$ $t_{1864} = +3.55$ $p = 0.00039$		
$\eta_s : R_0$	$\beta = +0.000625$ $SE = 0.00221$ $t_{4546} = +0.28$ $p = 0.78$				
η_e				$\beta = -0.00125$ $SE = 0.00383$ $t_{1873} = -0.33$ $p = 0.74$	
t_{max}					$\beta = -0.00114$ $SE = 0.0027$ $t_{1794} = -0.42$ $p = 0.67$
$adj - R^2$	0.20	0.23	0.20	0.21	0.04
N_{obs}	4550	1823	1866	1875	1796
BIC	-6754.24	-2726.22	-2771.43	-2806.02	-2617.09

Supplementary Table 13: **Fixed effect of the analysis of task variables on the Optimal Policy Index.** Models were specified as follows. Exp. 1b: $\text{OPI} \sim 1 + \eta_s * R_0 + (1 + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $\text{OPI} \sim 1 + R_0 + (1 + R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $\text{OPI} \sim 1 + \eta_s + (1 + \eta_s \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $\text{OPI} \sim 1 + \eta_e + (1 + \eta_e \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $\text{OPI} \sim 1 + t_{max} + (1 + t_{max} \mid \text{participant}) + (1 \mid \text{trial})$. OPI: Optimal Policy Index, averaged across the first 5 samples of each trial.

	Exp. 1b	Exp. 2	Exp. 3	Exp. 4	Exp. 5
(Intercept)	$\beta = +0.0238$ $SE = 0.00422$ $t_{4546} = +5.64$ $p < \mathbf{0.0001}$	$\beta = +0.032$ $SE = 0.00775$ $t_{1821} = +4.14$ $p < \mathbf{0.0001}$	$\beta = +0.0259$ $SE = 0.00698$ $t_{1864} = +3.72$ $p = \mathbf{0.00021}$	$\beta = +0.0321$ $SE = 0.00635$ $t_{1873} = +5.06$ $p < \mathbf{0.0001}$	$\beta = +0.00908$ $SE = 0.000765$ $t_{1794} = +11.87$ $p < \mathbf{0.0001}$
R_0	$\beta = +0.00107$ $SE = 0.000539$ $t_{4546} = +1.98$ $p = \mathbf{0.048}$	$\beta = -0.000312$ $SE = 0.000885$ $t_{1821} = -0.35$ $p = 0.72$			
η_s	$\beta = -0.00461$ $SE = 0.00137$ $t_{4546} = -3.36$ $p = \mathbf{0.0008}$		$\beta = -0.00312$ $SE = 0.00178$ $t_{1864} = -1.75$ $p = 0.08$		
$\eta_s:R_0$	$\beta = -0.000121$ $SE = 0.000757$ $t_{4546} = -0.16$ $p = 0.87$				
η_e				$\beta = +0.00186$ $SE = 0.000938$ $t_{1873} = +1.98$ $p = \mathbf{0.047}$	
t_{max}					$\beta = +0.000657$ $SE = 0.000307$ $t_{1794} = +2.14$ $p = \mathbf{0.033}$
$adj - R^2$	0.52	0.47	0.58	0.48	0.08
N_{obs}	4550	1823	1866	1875	1796
BIC	-18501.94	-6663.06	-7957.60	-7747.23	-10997.81

Supplementary Table 14: **Fixed effect of the analysis of task variables on the Habitual Policy Index.** Models were specified as follows. Exp. 1b: $HPI \sim 1 + \eta_s * R_0 + (1 + \eta_s * R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 2: $HPI \sim 1 + R_0 + (1 + R_0 \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 3: $HPI \sim 1 + \eta_s + (1 + \eta_s \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 4: $HPI \sim 1 + \eta_e + (1 + \eta_e \mid \text{participant}) + (1 \mid \text{trial})$; Exp. 5: $HPI \sim 1 + t_{max} + (1 + t_{max} \mid \text{participant}) + (1 \mid \text{trial})$. HPI: Habitual Policy Index, averaged across the first 5 samples of each trial

	Sample	ISI (ln-transformed)	α (ln-transformed)	EE (ln-transformed)	score
(Intercept)	$\beta = +0.00327$ $SE = 0.148$ $t_{40} = +0.02$ $p = 0.98$	$\beta = -0.0112$ $SE = 0.153$ $t_{40} = -0.07$ $p = 0.94$	$\beta = +0.113$ $SE = 0.144$ $t_{40} = +0.79$ $p = 0.44$	$\beta = +0.0155$ $SE = 0.159$ $t_{40} = +0.10$ $p = 0.92$	$\beta = +0.00885$ $SE = 0.122$ $t_{40} = +0.07$ $p = 0.94$
AMI	$\beta = +0.221$ $SE = 0.158$ $t_{40} = +1.39$ $p = 0.17$	$\beta = -0.262$ $SE = 0.164$ $t_{40} = -1.59$ $p = 0.12$	$\beta = -0.0346$ $SE = 0.154$ $t_{40} = -0.22$ $p = 0.82$	$\beta = -0.146$ $SE = 0.17$ $t_{40} = -0.86$ $p = 0.40$	$\beta = -0.297$ $SE = 0.131$ $t_{40} = -2.27$ $p = 0.029$
BIS	$\beta = -0.00161$ $SE = 0.175$ $t_{40} = -0.01$ $p = 0.99$	$\beta = +0.222$ $SE = 0.181$ $t_{40} = +1.23$ $p = 0.23$	$\beta = -0.0133$ $SE = 0.17$ $t_{40} = -0.08$ $p = 0.94$	$\beta = +0.0379$ $SE = 0.188$ $t_{40} = +0.20$ $p = 0.84$	$\beta = -0.0223$ $SE = 0.144$ $t_{40} = -0.15$ $p = 0.88$
IUS	$\beta = +0.382$ $SE = 0.166$ $t_{40} = +2.30$ $p = 0.027$	$\beta = -0.263$ $SE = 0.172$ $t_{40} = -1.52$ $p = 0.14$	$\beta = -0.0654$ $SE = 0.162$ $t_{40} = -0.40$ $p = 0.69$	$\beta = -0.112$ $SE = 0.179$ $t_{40} = -0.63$ $p = 0.53$	$\beta = -0.554$ $SE = 0.137$ $t_{40} = -4.04$ $p = 0.00024$
Raven	$\beta = +0.0248$ $SE = 0.156$ $t_{40} = +0.16$ $p = 0.87$	$\beta = -0.078$ $SE = 0.161$ $t_{40} = -0.48$ $p = 0.63$	$\beta = -0.0794$ $SE = 0.151$ $t_{40} = -0.52$ $p = 0.60$	$\beta = -0.111$ $SE = 0.167$ $t_{40} = -0.66$ $p = 0.51$	$\beta = +0.013$ $SE = 0.129$ $t_{40} = +0.10$ $p = 0.92$
age	$\beta = +0.0769$ $SE = 0.165$ $t_{40} = +0.47$ $p = 0.64$	$\beta = +0.0582$ $SE = 0.171$ $t_{40} = +0.34$ $p = 0.74$	$\beta = -0.228$ $SE = 0.16$ $t_{40} = -1.42$ $p = 0.16$	$\beta = -0.0934$ $SE = 0.177$ $t_{40} = -0.53$ $p = 0.60$	$\beta = -0.113$ $SE = 0.136$ $t_{40} = -0.83$ $p = 0.41$
gender	$\beta = -0.281$ $SE = 0.157$ $t_{40} = -1.79$ $p = 0.08$	$\beta = -0.0561$ $SE = 0.163$ $t_{40} = -0.34$ $p = 0.73$	$\beta = +0.185$ $SE = 0.153$ $t_{40} = +1.21$ $p = 0.23$	$\beta = +0.237$ $SE = 0.169$ $t_{40} = +1.40$ $p = 0.17$	$\beta = +0.246$ $SE = 0.13$ $t_{40} = +1.90$ $p = 0.06$
logBDI	$\beta = -0.0884$ $SE = 0.178$ $t_{40} = -0.50$ $p = 0.62$	$\beta = +0.032$ $SE = 0.185$ $t_{40} = +0.17$ $p = 0.86$	$\beta = -0.105$ $SE = 0.173$ $t_{40} = -0.61$ $p = 0.55$	$\beta = +0.0706$ $SE = 0.192$ $t_{40} = +0.37$ $p = 0.71$	$\beta = -0.0881$ $SE = 0.147$ $t_{40} = -0.60$ $p = 0.55$
$adj - R^2$	0.21	0.13	0.12	0.10	0.43
N_{obs}	48	48	48	48	48

Supplementary Table 15: **Robust multiple linear regressions of the relationship between active task performance and questionnaire measures.** For each dependent variable, the robust multiple linear regression was specified as follows: $y \sim 1 + \text{AMI} + \text{BIS} + \text{IUS} + \ln(\text{BDI}) + \text{Raven} + \text{Age} + \text{Gender}$. The BDI score was ln-transformed for normalisation purposes.