Humans balance breadth and depth: Near-optimal performance in many-alternative decision making

Alice Vidal^{1,2}, Salvador Soto-Faraco^{1,3}
and Rubén Moreno-Bote^{1,4}

¹ Center for Brain and Cognition, and Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

² Department of Experimental and Health Sciences, Universitat Pompeu Fabra, Barcelona, Spain

³ Insitució Catalana de la Recerca i Estudis Avançats (ICREA), Barcelona, Spain

⁴ Serra Hùnter Fellow Programme, Universitat Pompeu Fabra, Barcelona, Spain

Corresponding author: Alice Vidal

email: alice.vidal@upf.edu

Acknowledgements: This work is supported by the Howard Hughes Medical Institute (HHMI, ref 55008742), MINECO (Spain; BFU2017-85936-P) and ICREA Academia (2016) to R.M.-B. S.S-F. is funded by Ministerio de Ciencia e Innovación (Ref: PID2019-108531GB-I00 AEI/FEDER) and the FEDER/ERFD Operative Programme for Catalunya 2014-2020. A.V. is supported by a FI fellowship from the AGAUR (2019FI_B 00302). We would like to thank Pierre Marais for his precious help in setting up the Flutter application and coding the experimental task.

Abstract

Many everyday life decisions require allocating finite resources, such as attention or time, to examine multiple available options, like choosing an online food supplier. In these cases, our search resources can be spread across many options (breadth) or focused on a few of them (depth). Whilst theoretical work has described how finite resources should be allocated to maximise utility in these problems, evidence about how humans balance breadth and depth is lacking. We introduce a novel experimental paradigm where humans make a many-alternative decision under finite resources. In an imaginary scenario, participants allocate a finite budget to sample amongst multiple apricot suppliers in order to estimate the quality of their fruits, and ultimately choose the best one. We found that at low budget capacity participants sample as many suppliers as possible, and thus prefer breadth, whereas at high capacities participants sample just a few chosen alternatives in depth, and intentionally ignore the rest. The number of alternatives sampled increases with capacity following a power law with an exponent close to 0.75. In richer environments, where good outcomes are more likely, humans further favour depth. Participants deviate from optimality and tend to allocate capacity amongst the selected alternatives more homogeneously than it would be optimal, but the impact on the outcome is small. Overall, our results undercover a rich phenomenology of close-to-optimal behaviour and biases in complex choices.

Introduction

When choosing an online food supplier as we settle into a new area, we need to trade off the number of alternatives that we check with the time or money that we want to invest in each of them to learn about the quality of their products. Distributing resources widely (breadth search) allows us to sample many suppliers but very superficially, thus limiting our ability to distinguish which one is best. Allocating our resources to check just a few suppliers (depth search) allow us to learn detailed information but only from a few, at the risk of neglecting relevant ones. Striking the right balance between breadth and depth is critical in countless daily life endeavours such as when selecting which courses to register in college (Schwartz et al., 2009) or developing marketing strategies (Turner et al., 2002). Its implications are far reaching when it comes to understand behaviours in the internet, where breadth and depth-search has been related to navigation through lists of results (Klöckner et al., 2004) or web site menus (Miller, 1981).

Despite its relevance to understand how humans make decisions under finite resources, it is remarkable that the *breadth-depth* dilemma has mostly been investigated outside cognitive neuroscience (Moreno-Bote et al., 2020), in contrast to other well-studied trade-offs like speed-accuracy and exploration-exploitation (J. D. Cohen et al., 2007; Costa et al., 2019; Daw et al., 2006; Ebitz et al., 2018; Wilson et al., 2014). The breadth-depth dilemma underlies virtually all cognitive problems, from allocating attention amongst multiple alternatives in multi-choice decision making (Busemeyer et al., 2019; Hick, 1952; Proctor & Schneider, 2018), splitting encoding precision to items in working memory (Joseph et al., 2016; Ma et al., 2014), to dividing cognitive effort into several ongoing subtasks (Feng et al., 2014; Musslick & Cohen, 2021; Shenhav et al., 2013). In all these problems, finite resources, such as attention, memory precision or amount of

control, need to be allocated amongst many potential options simultaneously, making the efficient balancing between breadth and depth a fundamental computational conundrum.

The dynamics of resource allocation can be very complex, and thus it has been studied in decision making in simplified cases with few alternatives (Callaway et al., 2021; Jang et al., 2021; Krajbich et al., 2010) or in multi-tasking using a low number of simultaneously active tasks (Musslick & Cohen, 2021; Sigman & Dehaene, 2005). In these and other cases, resource allocation can be changed on the fly if feedback is immediate or is available within very short delays. In some real-life situations, however, feedback about the quality of the allocation is necessarily delayed. This happens for instance in problems such as investing (Blanchet-Scalliet et al., 2008; Reilly et al., 2016), choosing college (Schwartz et al., 2009) or, in ant colonies, sending scouts for exploration (Pratt et al., 2002); feedback can come after seconds, days or even years. As resources should then be allocated beforehand, the dynamic aspect of the allocation is less relevant. One-shot resource allocation is important in cognition as well, as building-in stable attentional or control strategies that work in a plethora of situations could relieve the burden of solving a taxing breadth-depth dilemma for optimal resource allocation every time (Moreno-Bote & Mastrogiuseppe, 2021).

Previous theoretical work has shown how to optimally trade-off breadth and depth over multi-alternative problems in the situations described above, where resources are allocated all at once before feedback is received (Moreno-Bote et al., 2020; Moreno-Bote & Mastrogiuseppe, 2021; Ramírez-Ruiz & Moreno-Bote, 2021). A central result is that the optimal trade-off depends on the search capacity of the agent: while at low capacity resources should be split in as many alternatives as capacity permits (breadth), at high capacity resources should be focused on a relatively small number of selected alternatives so that available resources are more focused (depth) (Moreno-Bote et al., 2020). In rich environments, where finding options rendering good outcomes is more likely, depth should be further favoured. Despite of the existence of precise predictions describing ideal behaviours in these scenarios, how humans solve breadth-depth trade-offs in many-alternative decision making is largely unknown (Brown et al., 2009; Callaway et al., 2021; Chau et al., 2014; A. L. Cohen et al., 2017; Hawkins et al., 2012; Moreno-Bote et al., 2020; Roe et al., 2001; Usher & McClelland, 2004; Vul et al., 2014).

To fill this gap, we designed a novel, many-alternative task where search capacity was parametrically controlled on a trial-by-trial basis. Human participants are immersed in a context where they are asked to allocate finite search capacity over a large number (several dozens) of apricot suppliers with the goal to choose the best one. We compare human sampling behaviour to optimality (Moreno-Bote et al., 2020) by using two different ranges of capacities to zoom in relevant regimes. Concerning the adaptation of the sampling strategy as a function of capacity, we expected a pure-breadth behaviour at low capacity followed by a sharp transition towards a trade-off between breadth and depth, characterised by an increase of the number of alternatives sampled with the square root of capacity. We also addressed the effect of environment richness (the overall probability of alternatives rendering good outcomes) on the sampling strategy. We used three different environments with either a majority of poor, neutral, or rich alternatives and expected sampling strategy to shift towards depth as the environment became richer. Finally, we looked at whether systematic deviations from optimality can be observed

regarding how samples are distributed amongst each sampled alternative. We employed model comparison to adjudicate between optimal and other heuristic models of behaviour.

Results

Humans switch from breadth to a breadth-depth balance as capacity increases

We developed a novel experimental approach, the ‘BD apricot task’, to study many-alternative decisions under uncertainty and limited resources by confronting participants with a breadth-depth dilemma. Participants played a gamified version of the task in which they were presented with virtual apricot suppliers (10 or 32) having different, and unknown, probabilities of good-quality apricots (Methods; [here](#) task demo video). These probabilities were independently generated for each supplier in each trial from a beta prior distribution. At the beginning of each trial, participants received a few coins, which defined their sampling capacity in that trial. Each coin was used to buy one apricot from a selected supplier (**Fig. 1a**). Once all coins have been spent (**Fig. 1b**), participants discovered which of the sampled apricots were of good quality, and which ones were bad (**Fig. 1c**), whose outcomes followed a Bernoulli process with the unknown probabilities of good-quality apricots for each supplier. These probabilities were sampled independently for each vendor and trial from a beta distribution. Based on the observed outcomes, participants could estimate the probabilities from the sampled suppliers, and finally make a final purchase of 100 apricots from the one they considered to be the best (**Fig. 1d**). Their goal was to maximise the number of good-quality apricots collected throughout the experiment through the implementation of an informative sampling strategy, adapted both to the sampling capacity and the environment richness. The task was intuitive and easily grasped by most participants. No instructions about the underlying probability generative model was provided as the context was informative enough to aid task understanding (Schustek et al., 2019).

Figure 1. The BD apricot task. Human participants allocate a finite search capacity (coins) to learn about the quality of good apricots in different suppliers (sampling phase) and then make a final purchase of 100 apricots from one of the sampled suppliers (purchase phase). Each black section of the wheel represents a different supplier. The number of coins represents the search capacity of the participants on each trial and varies randomly from trial to trial within a finite range (see Methods). The available coins (panels a; yellow green dots) at any time during the trial are displayed within the centre of the wheel. To allocate the coins to suppliers, participants have first to click on the designated active coin displayed at the centre (green dot) and then select

the supplier to sample from (panels a) –both touch screen events are indicated by a large grey dot. One of the inactive (yellow) coins is then automatically activated and displayed, in green, at the centre. This sequence repeats until all coins are allocated. Then, each of the allocated samples turn either orange, representing a good-quality apricot, or purple, representing a bad-quality apricot (panel c). Finally, after this information is revealed, the participant selects one of the sampled suppliers for the final purchase of 100 apricots (with a touch screen, indicated by a large grey dot) and the choice outcome is immediately displayed (panel d).

Two different ranges of sampling capacity, narrow (2 to 10 samples per trial) and wide (2 to 32), were tested to zoom-in relevant behavioural regimes (see Methods and **Table.1**). For each range, we used three environments differing in the parameters of the beta distribution used to generate probabilities of good-quality apricots for each supplier. By increasing the average probability of good alternatives, we can increase the richness of the environment to move from a ‘poor’, to a ‘neutral’, up to a ‘rich’ environment (beta prior means: poor, 0.25; neutral, 0.50; and rich, 0.75. While environments and capacity ranges were run in a block fashion, capacity in each was chosen randomly in each trial from the underlying range in that block. Environments were tested within and between subjects in two different sets of experiments, to address potential learning effects.

	Design	Capacity	N suppliers	N trials	N subject
W10	Within-subject	2 to 10	10	216	18
B10	Between-subject	2 to 10	10	72	45
W32	Within-subject	2,4,8,16,32	32	120	18
B32	Between-subject	2,4,8,16,32	32	40	45

Table 1. Summary of the experimental designs.

We observed in all environments that participants’ sampling behaviour over the suppliers follows a pure-breadth strategy at low capacity (**Fig. 2A-B**), whereby they sample close to as many suppliers as coins they have. Indeed, at low capacity ($C \sim 2,3$) the number M of sampled suppliers averaged over trials and participants is close to C (**Fig. 2A**) and the ratio M/C reaches very close to one (**Fig. 2B**). At larger capacities, sampling progressively evolves towards a trade-off between breadth and depth. Therefore, although participants could have sampled more suppliers to learn about, they preferred to focus sampling capacity on a rather small fraction of suppliers, as shown by the decline of the ratio M/C towards values of around 0.4 at the largest capacities tested (**Fig. 2B**; right panels). Although these are predicted features and thus signs of optimal behaviour, we did not observe a fast transition between the low and high-capacity regimes, as previously predicted (Moreno-Bote et al., 2020). This could simply result from averaging over participants having different transition points, or by having allocation noise, which would smooth out fast transitions. Indeed, as we show below, models with allocation noise account for the smoothness of this transitions and are better models in predicting the behaviour than the noise-free optimal model.

Figure 2. Number of sampled suppliers increases with capacity and is strongly sensitive to the richness of the environment, consistent with theoretical predictions. **A.** Number of alternatives sampled M as a function of capacity averaged across participants (points), for each of the three different environments (colours), for the low-capacity between-subject design (B10). Dashed lines indicate unit slope line. Optimal observer predictions are displayed in black and grey lines represents individual data. Error bars correspond to s.e.m. **B.** Number of alternatives sampled M divided by capacity as a function of capacity. Colour code as in panel (A).

Richer environments promote depth

The effect of environmental richness was also clearly visible in our data, with richer environments typically causing a stronger preference for depth sampling (**Fig. 2A-B**; different colours), consistent with predictions. Therefore, as it becomes easier for the participant to find good options, they prefer to neglect an even larger number of suppliers and focus capacity on fewer ones. We also observed systematic deviations from optimally (**Fig. 2A**; black lines). In particular, despite participants' strategy was overall strongly dependent on environment richness, as predicted, this was mostly due to a strong effect of the poor environment (**Fig. 2A-B**; red lines), whereas they were not sensitive to the difference between neutral and rich environments (green and blue lines).

To test quantitatively the effect of environmental richness on the breadth-depth trade-offs, we fitted participant's individual data in each environment using three models. These models were chosen based both on our predictions and previous observations. First, a piece-wise power law model (W), having a fast transition at some arbitrary capacity value, was selected as it is the one anticipated by the ideal observer. Second, as previously said, we visually noticed that for a majority of participants the transition between pure-breadth and BD trade-off was progressive (see **Fig. 2**), so we decided to capture their sampling strategy using two simpler models: a linear model (L) and a power law model (P). Indeed, it has been predicted that once the BD trade-off established, the

number of alternatives sampled M approximately increases with a power law behaviour (Moreno-Bote et al., 2020). We observed that the power law model ($R_{adj}^2 = 0.96 \pm 0.04$ mean $\pm$ s.d.) showed significantly better fits than the linear model ($R_{adj}^2 = 0.92 \pm 0.06$ mean $\pm$ s.d.; $V = 404, p < 2.2 \times 10^{-16}$). Moreover, the linear piece-wise model was not significantly better than the power law model for all participants in all environments individually (ANOVAs, with $\alpha = .05$).

We compared the exponent estimated from the power law fits in each environment and experimental design (**Fig. 3**) using within- (W10 & 32) or between-subjects (B10 & 32) designs. The results demonstrated a significant effect of the environment on the exponent in all designs (one-way ANOVAs, B10: $F_1 = 13.17, p = 7.51 \times 10^{-4}$; W10: $F_1 = 31.05, p = 3.37 \times 10^{-5}$; W32: $F_1 = 15.12, p = .0012$), except for the between-subjects design and the larger range of capacities tested (B32: $F_1 = 0.82, p = .37$). Overall, effect sizes were larger in designs with narrow (B10: $\omega_G^2 = .213$, W10: $\omega_G^2 = .121$) compared to wide ranges of capacities (B32: $\omega_G^2 = -.004$, W32: $\omega_G^2 = .072$) and the non-result observed may stem from an insufficient sample size (achieved power in B32 is 14.7%, against 97.8% in W32, 95.2% in B10 and 94.7% in W10). As suggested by visual inspection above, there is no significant difference in the exponents between the neutral (median values, W10: $m_{neutral} = 0.72$, B10: $m_{neutral} = 0.68$, W32: $m_{neutral} = 0.70$) and rich (W10: $m_{rich} = 0.71$, B10: $m_{rich} = 0.66$, W32: $m_{rich} = 0.66$) environments in any of the designs (t-tests with Bonferroni correction, W10: $t_{17} = 0.57, p_{adj} = 1$, B10: $t_{23.9} = 0.24, p_{adj} = 1$, W32: $t_{17} = 2.04, p_{adj} = .17$), although strong and significant differences were typically observed between rich and poor (W10: $m_{poor} = 0.82$, B10: $m_{poor} = 0.84$, W32: $m_{poor} = 0.75$) environments (W10: $t_{17} = 5.57, p_{adj} = 1.01 \times 10^{-4}$, B10: $t_{20.8} = 3.51, p_{adj} = .006$, W32: $t_{17} = 3.89, p_{adj} = .004$). Significant differences were also observed between poor and neutral environments in the designs with smaller capacities (W10: $t_{17} = 3.25, p_{adj} = .014$, B10: $t_{26.7} = 4.45, p_{adj} = 4.11 \times 10^{-4}$). Overall, we observed in the three designs mentioned a decrease of the exponent as the environment gets richer. This suggests an adaptation of participants' sampling behaviour to the environment, with sampling depth increasing with the richness of environment.

Figure 3. Participants' strategy is modulated by the richness of the environment. Distribution of power factors extracted from fitting a linear model to values M vs capacity in a log-log scale. Colour dots represent subjects, black dots represent means across participants and bars are 95% confidence intervals. Results of post-hoc comparisons are displayed according to adjusted p-values ('ns': $p > .05$, '**': $p < .01$, '***': $p < .001$).

Sequential effects

In the within-subject designs (W10 and W32), where participants were successively presented with the three different environments, we wondered whether their sampling strategy depended on the previous environment presented (**Fig. S1**). We observe that in several cases sampling behaviours seems to shift toward breadth in environments directly following the presentation of a poorer environment, or towards depth if a richer environment was presented immediately before. For instance, participants sampling strategy in the neutral environment (**Fig. S1** middle panels) seems to differ depending on whether it was presented first or not. If presented after the poor environment, we observe a clear deviation towards breadth, whereas a shift towards depth is observed if presented after the rich environment. To statistically test the presence of a sequential, or contamination effect, on participants' sampling strategy, we chose to compare the exponent estimated from fitting the number of alternatives samples M depending on the capacity using the power-law model (P) previously described. We calculated the averaged exponent for each environment and experiment design (W10 and W32) based only on the first environment presented and use it as baseline. We then subtracted it to the estimated exponents of the second and third environments presented, for each environment and design separately, and obtained the deviations from the means depending on the previous environment presented. We observed on average positive deviations after the presentation of a poor environment (exponent mean $\pm$ s.d.: 0.05 ± 0.12) suggesting that participants sample with more breadth when previously presented with a poor environment. On the contrary, we observed on average negative deviations in the exponent after the presentation of a neutral (-0.029 ± 0.15) or rich environment (-0.23 ± 0.13) suggesting that participants sample more deeply when previously presented with those environments. Overall, we do observe a weak but significant effect of the previous environment presented on participants sampling strategy (measured by the difference between the observed and the average exponent) (ANOVA, $F_2 = 3.403, p = .045$). Post-hoc analyses do not reveal significant

differences in participants' sampling strategies following the presentation of a particular environment (t-tests with Bonferroni corrections, poor vs. rich: $t_{45.8} = 2.03, p_{adj} = .14$, poor vs. neutral: $t_{44.7} = 2.08, p_{adj} = .13$, neutral vs. rich: $t_{45.5} = -0.16, p_{adj} = 1$). Although based on exploratory analyses, these observations raise the possibility that participants sampling strategy was affected by the previous environment presented and that adaptation to a new environment requires to overcome the behaviour implemented earlier.

Biases relative to optimality

We demonstrated that, as expected, participants' sampling strategy is modulated by the environment richness. We now investigate whether the strategy used coincides with the optimal sampling behaviour or if specific deviations are observable. Pooling the data of the four experimental designs together, we confirmed our previous results and found a significant effect overall of the environment on the exponent of the power law (ANOVA, $F_1 = 21.03, p = 8.23 \times 10^{-6}$). Importantly there was no significant effect of the experimental design (ANOVA, $F_3 = 0.76, p = .52$), nor a significant interaction between the environment and the experimental design (ANOVA, $F_3 = 1.40, p = .24$) confirming that the effect of the environment can be studied on the whole data set, independently of the experimental design. In order to quantify deviations of participants' sampling strategies from the optimal strategy, we fitted the optimal values of the number M of sampled suppliers for all capacities ($C = \{2 - 10, 16, 32\}$) using the power-law model previously described (see Methods) and extracted the power law exponent. Comparing the observed values of the exponent to the optimal ones (see **Table. 2**), we observed that participants' sampling strategy significantly shifted toward excessive depth in the poor environment (t-test, $t_{65} = -5.66, p_{adj} = 3.72 \times 10^{-7}$), and a similar tendency occurred in the neutral environment (t-test, $t_{65} = -1.74, p_{adj} = .087$). In contrast, participants' sampling strategy deviated significantly from optimality in the direction of excessive breadth in the rich environment (t-test, $t_{65} = 4.06, p_{adj} = 1.36 \times 10^{-4}$).

environment value of power factor a	poor	neutral	rich
optimal (predicted)	.877	.732	.612
data (observed)	.795 $\pm$.118	.705 $\pm$.127	.688 $\pm$.152
comparison	$t_{65} = -5.66,$ $p_{adj} = 3.72 \times 10^{-7}$	$t_{65} = -1.74,$ $p_{adj} = .087$	$t_{65} = 4.06,$ $p_{adj} = 1.36 \times 10^{-4}$

Table 2. Participants' sampling strategy significantly deviates from optimality. Values of the factor a predicted (first row) or observed (averaged across participants $\pm$ s.d., second row) depending on capacity using a power law function with free exponent and fixed intercepted (see methods for more details). Results of comparisons between factors a using one-sample t-tests with Bonferroni corrections (third row) show that participants' sampling strategy extracted from fitting the number of alternatives sampled is significantly tilted toward depth in the poor and neutral (tendency) environments compared to optimality, while in the rich environment participants are sampling in a breather way than predicted.

Looking in greater detail at which capacities these deviations occur, we computed the differences between the optimal number M_{opt} of sampled suppliers and the observed M for each capacity and environment and observed again a significant effect of the environment on the differences (Scheirer-Ray-Hare test, $H_2 = 159.67, p < 2.2 \times 10^{-16}$), but also a significant effect of capacity ($H_{10} = 98.07, p = 2.2 \times 10^{-16}$) and a significant interaction between the environment and capacity ($H_{10} = 84.08, p = 7.89 \times 10^{-10}$), showing that participants bias toward excessive depth or breadth varies with the environment and the capacity. Indeed, we observed that at low capacity, the difference between M_{opt} and the observed M tended to be positive, which is especially visible in the poor and neutral environments (**Fig. S2** and exhaustive analyses presented in **Table. S1**). In contrast, at high capacity this difference tends to be negative, which is especially noticeable in the neutral and rich environments.

The results of this analysis indicate deviations from optimality and suggest that participants have a bias to sample more deeply than optimal at low capacities while this bias is reversed at high capacities. These overall biases could be accounted for by assuming that participants have a biased model of the richness of the environment, not trusting as much as they should the extreme nature of poor or rich environments (Schustek et al., 2019).

Model comparison

To closer characterize the behavioural patterns, we generated quantitative predictions from a variety of models, including the optimal model and several heuristics, and tested which are better able to predict the empirical observations. We chose to represent extreme behaviours, such as full-breadth (M is always equal to capacity C) or depth (M always equal to 2) as boundary models, as well as in-between behaviours which would represent the use of a trade-off between breadth and depth. Based on observations of the data and predictions (Moreno-Bote et al., 2020), we chose to model M as function of C using linear and power-law models. In the latter case, we used two versions of the power-law model: one where M increases with the square-root of capacity (exponent of 1/2) (Moreno-Bote et al., 2020), and another where the exponent was free to vary (see Methods for details of the models). To model realistic behavioural patterns, we injected noise into the model predictions that corrupts the chosen number of suppliers M . As noise models, we used binomial and discretised Gaussian distributed noise added to the model prediction of M . In the second case, the noise follows a normal distribution with a standard deviation that increased linearly with capacity. This choice was made based on the empirical observation that the variance of M increases linearly with capacity (Kendall's rank correlation test, $z = 11.26, \tau = .87, p < 2.2 \times 10^{-16}$, see **Fig. S3**). Data from the four experimental designs were analysed and are presented together. Four-fold cross-validated log-likelihoods (CVLL) are used to compare the models (Methods; qualitatively identical results hold when using AIC instead, see **Fig.S5** and **S6**). Using binomial noise, the free power-law model outperforms all five other models (**Fig.4A** and **Table.3**). The optimal model, even though less efficient than the power model, is better at predicting participants' sampling strategy than the pure breadth, depth and square root models. The outcomes of model comparison are mostly equivalent when using the Gaussian noise model (for the results of this model, see **Fig.S4** and **Table.S2**), and

therefore we restrict our discussion to the binomial noise model from now on. We just note that the only divergence between the noise models appeared in the pairwise comparisons, with the hierarchy of models' performance being more clearly statistically established in the binomial noise model compared to the Gaussian one.

We observed a significant interaction between the environment and the models on the CVLL (Scheirer-Ray-Hare test; $H_5 = 24.70, p = 1.59 \times 10^{-4}$). The depth model is better at predicting the data in the rich compared to the poor environment (Wilcoxon-test with Bonferroni correction, $W_{66} = 348, p_{adj} = 3.99 \times 10^{-6}$). The square root model is also shown to perform worst in the poor environment compared to the rich ($W_{66} = 379, p_{adj} = 1.06 \times 10^{-5}$) and the neutral ($W_{66} = 325, p_{adj} = 1.88 \times 10^{-6}$) environments. Additionally, the free power-law model once again outperforms any other models in all environments (**Fig.4B**).

As overall the free power-law model is best at predicting the empirical data, we chose to confirm the effect of the environment on participants' sampling strategy (see **Fig.3**) comparing the exponents parameters estimated from the power-law model between environments (**Fig.4C**). This new analysis confirmed the results reported earlier (see **Fig.3**), with a significant effect of the environment on the exponent w in all designs B10 (Kruskal-Wallis test, $K - W \chi^2_2 = 14.43, p = 7.34 \times 10^{-4}$), W10 (one-way ANOVA, $F_1 = 26.47, p = 8.11 \times 10^{-5}$) and W32 ($F_1 = 12.81, p = .002$), except for the design B32 ($F_1 = 0.84, p = .37$). In all three designs W10, B10 and W32, we observed significant differences in the power exponent between the poor and the rich environment (t-tests with Bonferroni corrections, W10: $t_{17} = 5.15, p_{adj} = 2.43 \times 10^{-4}$, B10: $W_{15} = 182, p_{adj} = .012$, W32: $t_{17} = 3.58, p_{adj} = .007$). We also observed significant differences between the poor and the neutral environment in the designs W10 ($t_{17} = 3.16, p_{adj} = .017$) and B10 ($t_{15} = 200, p_{adj} = 8.37 \times 10^{-4}$) but not in the design W32 ($t_{17} = 2.14, p_{adj} = .142$). Moreover, we did not find any significant difference in the power exponent w values between the neutral and rich environments in any of the designs (W10: $t_{17} = 0.42, p_{adj} = 1$, B10: $t_{15} = 127, p_{adj} = 1$, W32: $t_{17} = 1.62, p_{adj} = .37$).

Figure 4. The free power-law model is better at predicted participants' sampling strategy than both the optimal and other heuristic models (using binomial distributed noise). Averaged negative cross-validated log-likelihood (CVLL) across participants for each model overall (A) and in each environment (B). Results of pair-wise comparisons are displayed according to adjusted p-values ('*': $p < .05$, '**': $p < .01$, '***': $p < .001$) and error bars correspond to s.e.m. C. Distribution of the power factor w in the free power model. Each colour dot represents a subject, black dots represent distribution means and bars 95% confidence intervals. Results of post-hoc comparisons are displayed according to adjusted p-values ('ns': $p > .05$, '**': $p < .01$, '***': $p < .001$).

	pure breadth	square root	optimal	linear	power
depth	$V_{126} = 2796,$ $p_{adj} = .051$	$V_{126} = 2,$ $p_{adj} = 3.22 \times 10^{-21}$	$V_{126} = 176,$ $p_{adj} = 1.91 \times 10^{-19}$	$V_{126} = 65,$ $p_{adj} = 1.44 \times 10^{-20}$	$V_{126} = 2,$ $p_{adj} = 6.90 \times 10^{-21}$
pure breadth		$V_{126} = 2926,$ $p_{adj} = .134$	$V_{126} = 1154,$ $p_{adj} = 6.35 \times 10^{-11}$	$V_{126} = 4,$ $p_{adj} = 4.95 \times 10^{-21}$	$V_{126} = 2,$ $p_{adj} = 3.22 \times 10^{-21}$
square root			$V_{126} = 1243,$ $p_{adj} = 2.86 \times 10^{-10}$	$V_{126} = 746,$ $p_{adj} = 3.48 \times 10^{-14}$	$V_{126} = 14,$ $p_{adj} = 1.98 \times 10^{-20}$
optimal				$V_{126} = 1510,$ $p_{adj} = 2.01 \times 10^{-8}$	$V_{126} = 143,$ $p_{adj} = 8.92 \times 10^{-20}$
linear					$V_{126} = 259,$ $p_{adj} = 4.49 \times 10^{-18}$

Table 3. Summary of the pair-wise comparisons (Wilcoxon Matched Pairs Signed-Ranks test) of the 4-folds averaged CVLL between all six models using Binomial distributed noise. P-values are adjusted with Bonferroni corrections and significative differences ($p < .05$) are highlighted. Models are ordered from worst (depth) to best (free power law).

Participants tend to sample homogeneously amongst alternatives

Above, we have presented analyses regarding how participants allocate the samples within the alternatives selected during sampling and compared it to optimal allocation of samples across capacities and environments. Deviations from the optimal strategy are clear, as shown both in the empirical data (which often deviates from the optimal allocation), as well as in the outcome of model comparison to fit these empirical data. This deviation is interesting because the difference to optimality is not only accounted for by unbiased noise introduced in the model comparison above, which suggests the existence of systematic biases in human choice behaviour.

In order to better understand these deviations from optimality, we looked in more detail at the most frequent observed allocation of samples. We observed that participants frequently allocated the samples homogeneously across the selected alternatives, a sampling strategy which in most cases differs largely from optimal allocation (**Fig. 5A**). To characterize this bias toward homogenous sampling, we computed the standard deviation of the ordered counts for each allocation of samples averaged for each participant and environment (Methods) and compared it to the standard deviation of the optimal allocation (**Fig. 5B**). We observed that the standard deviation of the sample allocations measured in the poor (Wilcoxon-test with Bonferroni correction, $W_{66} = 529, p_{adj} = 7.01 \times 10^{-4}$), neutral ($W_{66} = 295, p_{adj} = 6.87 \times 10^{-7}$) and rich ($W_{66} = 48, p_{adj} = 4.38 \times 10^{-11}$) environments were all significantly lower than the standard deviation assuming optimal sample allocation. Moreover, we observed that the magnitude of these differences differed between environments (Kruskal-Wallis test, $K - W \chi^2_2 = 111.92, p < 2.2 \times 10^{-16}$). In fact, they increased with the environment richness, with larger differences observed in the rich compared to neutral (Wilcoxon-test with Bonferroni correction, $W_{66} = 4089, p_{adj} = 1.04 \times 10^{-17}$) and poor ($W_{66} = 4183, p_{adj} = 2.20 \times 10^{-19}$) environments and larger differences in the neutral compared to the poor environment ($W_{66} = 2922, p_{adj} = .002$). That is, there was an overall bias toward homogenous sampling, and this bias was stronger as the environment was richer.

Finally, we investigated whether the observed deviations from the optimal sample allocation impacted the average reward obtained by the participants in our task. In each trial, the reward or outcome is defined as the number of good quality apricots among the 100 apricots bought. When computing the differences between the observed and optimal outcomes (the averaged outcome obtained when following the ideal observer), we observed significantly negative deviations from optimality in all three environments (mean $\pm$ s.d. in poor: $-8.74 \pm 11.60\%$, neutral: $-2.56 \pm 4.23\%$ and rich: $-2.39 \pm 2.74\%$; two-way Wilcoxon-tests with Bonferroni correction, poor: $W_{66} = 311, p_{adj} = 1.18 \times 10^{-6}$, neutral: $W_{66} = 457, p_{adj} = 1.04 \times 10^{-4}$, rich: $W_{66} = 228, p_{adj} = 6.33 \times 10^{-8}$) (**Fig. 5C**), showing that the bias observed in the sample allocation did have a deleterious impact on the outcome. The magnitude of these differences between optimal

and actual outcome was modulated by the environment (Kruskal-Wallis test, $K - W \chi^2_2 = 12.29, p = .002$), and while the harmful effect of a suboptimal strategy was moderate in the poor environment, it was very small in the remaining two. Surprisingly, the direction of the effect of environment on outcome was opposite to the one observed on standard deviations, as more deviation from optimality in the rich environment did not affect performance as much as in the poor environment. This inverted correlation is due to the nature of the environment itself. Indeed, a small deviation from the optimal strategy in the poor environment has a much larger impact on the outcome than in the rich environment, where a great majority of alternatives have a high probability of success. Indeed, we observed that the difference between the optimal and the observed outcomes seemed to decrease as the environment richness increased with differences in the poor environment being significantly higher than the ones in the neutral (Wilcoxon-tests with Bonferroni correction, $W_{66} = 1549, p_{adj} = .013$) and rich environment ($W_{66} = 1475, p_{adj} = .004$). No significant difference was observed in between the neutral and rich environments ($W_{66} = 2185, p_{adj} = 1$).

Figure 5. Participants tend to allocate homogeneously the samples across the selected alternatives, which largely differs from the optimal allocation, and impacts the outcome. **A.** Number of samples allocated to each sampled alternative depending on the sampling capacity C . Upper panels: most frequent allocation of samples observed across participants in the rich environment (design B32) as a function of capacity. Lower panels: allocation of samples maximising the reward (optimal). **B.** Distributions of the differences between observed and optimal standard deviations of the distribution of samples among the selected alternatives in each

environment (e.g. if $C = 4$ and 2 samples are allocated in a first alternative while the last 2 samples are each allocated in a second and third alternative, the standard deviation of this sample allocation would correspond to $sd(\{2,1,1\}) \approx 0.577$). Note that more homogeneous distributions tend to lead to lower standard deviations. C. Distributions of the differences between observed and optimal outcomes in each environment. In the last two panels, dots represent participants and includes all trials for which the optimal number of alternatives sampled was inferior to capacity ($M_{opt} < C$ – see methods for more details). Below each distribution is presented results of one-sample Wilcoxon-tests (‘**’: $p_{adj} < .01$, ‘***’: $p_{adj} < .001$) and above are presented results of Wilcoxon tests between each environment (‘ns’: $p_{adj} > .05$, ‘*’: $p_{adj} < .05$, ‘**’: $p_{adj} < .01$, ‘***’: $p_{adj} < .001$). All p -values have been adjusted with Bonferroni corrections.

Participants’ tendency to sample homogeneously may also be one of the sources of deviation from optimality observed in the number of alternatives sampled M . Indeed, we observe that among the trials where participants did not use a pure breadth nor a pure depth strategy ($M < C$ and $M > 1$), they first allocate all the samples in an alternative before sampling another one on 39.9% of the trials. In contrast, they chose to first sample once all the alternatives M on 21.8% of the trials. While the first strategy seems to prioritise the number of samples per alternative, the second one focuses on the number of alternatives sampled.

Discussion

Finding the best way to balancing opposite strategies is a ubiquitous problem in decision making. Dilemmas such as speed or accuracy (Fitts, 1966; Wickelgren, 1977), exploration or exploitation (March, 1991), breadth or depth, are zero-sum situations. Increasing one necessary comes to the detriment of the other, and maximising overall expected utility requires to strike an optimal trade-off. For instance, in the classic EE dilemma exploiting a known reward option precludes exploring potentially better ones. The BD dilemma, studied here, is related to EE but has the crucial difference that a limited resource can be divided into several options simultaneously. Therefore, both exploration and exploitation could a priori be performed in parallel. The trade-off that arises in BD is two-fold. First, given limited capacity, how much of it should be used for exploration and how much for exploitation. And second —and as we chose to focus on in this paper—, during exploration the agent can further divide finite resources to learn much about few alternatives or little about many of them. These trade-offs do not naturally arise in the EE, as it is the divisibility of the agent’s resource, and thus a new degree of freedom, what results in a qualitatively new problem. By focusing on the latter, here we provide answers on how human participants manage the allocation of finite, but potentially large, resources.

Previous research has formalized the BD dilemma using rational decision theory under capacity constraints (Moreno-Bote et al., 2020; Moreno-Bote & Mastrogiuseppe, 2021; Ramírez-Ruiz & Moreno-Bote, 2021). This research has described the optimal allocation of capacity over multiple alternatives as a function of both agent’s capacity and environmental richness. At low capacities, it is optimal to draw one sample per alternative (pure-breadth strategy), while for larger capacities, some alternatives are

ignored (balancing breadth and depth), and the number of alternatives sampled roughly increases with the square root of capacity, independently of the richness of the environment. Additionally, the passage from pure breadth to a BD trade-off occurs with a sharp transition at a specific capacity which directly depends on the overall success probability of the alternatives. Indeed, the transition is shifted towards larger capacity values as the environment gets poorer. Finally, it is optimal to favour uneven allocations of samples among the considered alternatives, especially as capacity and environment richness increases. Rational decision theory provides us with a normative hypothesis about how humans ought to behave.

Experimentally, BD trade-offs have been studied outside cognitive neuroscience (Halpert, 1958; Schwartz et al., 2009; Turner et al., 2002), mostly by using choice menus of different nature. However, this research does not address whether allocation policies used by humans are rational nor they parametrically manipulate an agent's capacity or the difficulty of the environment. To fill this gap, here we have introduced a novel experimental paradigm that allows us to compare human empirical data with normative predictions under a BD dilemma. The BD apricot task is intuitive and easily grasped by the participants. Consistent with theoretical predictions, we observe that, at low capacity, a pure-breadth strategy is favoured (sampling as many alternatives as possible), while as capacity increases, participants progressively sample relatively fewer of them in more depth. Additionally, our results reveal that participants' sampling strategy adapts to the environment with richer environments promoting depth over breadth.

We observe that human behaviour is close to, but systematically deviated, from optimality. Regarding the number of alternatives sampled as a function of capacity, we observe that participants' sampling strategy does not follow the predicted optimal model but is better captured by a power law. This strategy has the advantage to be continuous (no break point) which might contribute to reduce computational (hence, cognitive) cost, within a controlled loss of efficiency. Most likely is however an explanation based on the presence of sample allocation noise, which largely smooths the predicted sharp transitions between low and high capacity, noise for which we found some evidence (see Fig. S3). We also demonstrated that, at low capacity, participants often sample deeper than optimal, whilst at high capacities the reverse pattern is observed, with participants sampling more shallowly than optimal. Additionally, the deviations towards depth at low capacities are especially observed in poorer environments, while deviations towards breadth at high capacities are more prevalent in richer environments. Given that environment richness is not known precisely by participants, such results may be explained by errors in the estimation of the environment richness (Drugowitsch et al., 2016; Schustek et al., 2019; Wyart & Koechlin, 2016). However, they could also have more complicated origin. Vul and colleagues (Vul et al., 2014) showed that the optimal number of samples allocated per sampled alternative directly depends on the action/sample cost ratio in a way that a low action/sample cost ratio promotes depth over breadth. In our paradigm, the ratio between the gain in outcome (outcome received after sampling compared to chance) and the sampling cost (linearly increasing with the number of samples) is not constant over all capacities and environments. As a result, this change in ratio could also play a role in explaining why the amplitude of the deviations observed vary depending on the environment richness.

Regarding how samples are allocated across alternatives, we observe a participants' tendency to sample alternatives homogeneously. This distribution of samples deviates significantly from optimal behaviour, but caused only a small reduction of outcome obtained compared to optimal (mean $\pm$ s.d. deviations from optimality: $-3.5 \pm 6.24\%$). At a cognitive level, such bias toward an even distribution of samples may have various origins. It could emerge to simplify computations during sampling (even division of samples) or choice (easier to compare fractions with a common denominator), or could have other sources, such as the human preference for symmetry (Attneave, 1955). We tried to infer the thought processes that participants use while sampling through the identification of patterns in the sampling sequence and thus have a better understanding of the bias towards symmetric allocation. Our results suggest that participants focus more on the number of samples per alternative over the number of sampled alternatives. Consequently, participants' bias to sample homogeneously may drive, at least in part, deviations from optimality observed in the number of sampled alternatives.

In general, even though participants engage in behaviours which deviate significantly from optimality, the heuristics identified - continuous power law sampling strategy and homogeneous sampling - are associated with a lower cognitive effort compared to the optimal behaviour. Additionally, if following these biases impact significantly the outcome, the loss remains relatively small in our task. It is not clear whether these biases reflect participants' computational limitations - bounded rationality (Simon, 1955) - or are the result of a compromise between the error in the sampling strategy and the cost associated with brain computations - bounded optimality (Russell & Subramanian, 1995) or resource-rationality (Griffiths et al., 2015; Lieder et al., 2012). Importantly, none of the reported deviations can easily be explained by the presence of motor biases, as samples are allocated one by one and participants always must go to the screen centre to drag the next sample to a supplier, so that the distance to any supplier is identical and any motion-direction asymmetry is the same sample to sample.

Further, despite the large number of alternatives presented in the apricot task (either 10 or 32), we did not find evidence that participants considered much fewer options than the optimal number, and therefore our study cannot be explained by –and offers a counter example of– choice overload (Iyengar & Lepper, 2000; Kuksov & Villas-Boas, 2010; Sethi-Iyengar et al., 2003). First, we observe that participants consider many alternatives, often more than 5 (43.3% of the trials with $C > 5$) and up to 32 in some rare cases. Secondly, although we do observe that the number of sampled alternatives is sometimes inferior to the optimal number, this is predominant at rather low capacities ($M < 8$, see Fig.S2) when few alternatives are actually considered. In the contrary, at larger capacities, participants tend to consider more alternatives than optimal. In our experimental set-up, no information about any alternative was directly accessible to the participant before the outcome of the allocation was revealed, which may prevent perturbations and biases in the decision process by facilitating the a priori consideration of all alternatives on an equal footing.

Our novel experimental framework has demonstrated to be appropriate for studying human decisions in the context of the DB dilemma, but it also raises some issues concerning the difficulty to test particular conditions. Concerning the effect of

environment richness, we suspect contamination to happen in participants' sampling strategy when presented to different environments consecutively. The use of between-subjects designs should be favoured to thwart this effect. We also observed that the use of very large capacities ($C = 32$) is very costly for the participants and may have a negative impact in their motivation to engage in the task and the development of an optimal strategy. Taking into consideration these observations, our framework has the advantage to be easily adapted to investigate future questions. We are particularly curious to test whether the results observed here stand when using a, more realistic, continuous rather than a discrete capacity (Ramírez-Ruiz & Moreno-Bote, 2021) or investigate how participants manage to spend a limited capacity not over a single choice but over multiple ones. In addition, this design can be transposed to solve the BD dilemma in more complex choices, such as inside large decision trees (Moreno-Bote & Mastrogiuseppe, 2021).

To conclude, we have developed the first thorough experimental study of the BD dilemma, which, benefiting from a large adaptability and capability to translate real-life situations, offers a promising framework to study many-alternative human decision making under finite resources. Already, our results reveal the use of close-to-optimal choice behaviours which are both sensitive to the capacity and the environment and identify some of the heuristics used to simplify computations at the cost of optimality.

Methods

Experimental design

We developed a protocol – the Breadth-Depth Apricot Task – to study how humans sample the environment using limited resources, to make economic choices. The task was programmed using Flutter development software (flutter.dev) as an application for smartphones or tablets. This allowed testing participants outside lab premises. Participants were initially introduced with a realistic narrative that provided a concrete context to aid understanding the task goals and constraints (a demonstration video of the task is available [here](#)). According to this narrative, in each trial the participant had to buy an order of apricots in bulk from one specific supplier, out of many available. The goal was to maximize the amount of good quality apricots accumulated throughout the experiment. Because suppliers vary in the proportion of good quality apricots they serve, participants were given the opportunity to learn about the suppliers' overall quality by sampling: prior to the purchase, participants were asked to sample apricots from various suppliers of their choice using a fixed number of 'free' samples given to them. Based on this, they were to choose the supplier for the final purchase in the trial.

Specifically, each trial in the task was divided into a sampling phase and a final purchase phase. In the sampling phase (Fig. 1a), the participant was given a number of coins (yellow dots) that varied from trial to trial. The number of coins determines the search capacity of the participant on each trial. We ran different designs, which varied in the range of available capacity throughout trials, with each capacity selected randomly and equiprobably, from the pre-established range. In low-capacity designs, the range of possible capacities ranged between 2 and 10, in steps of one, while in high-capacity designs capacity took one of the values in the set $\{2,4,8,16,32\}$. In each trial, the coins could be freely allocated one by one to any of the suppliers by clicking the active coin in

the middle of the display and then by clicking the desired supplier to sample from. Participants could arbitrarily allocate the coins in a given trial (i.e., all coins to just one supplier, or each coin to a different supplier, or anything in between). The number of possible suppliers was always fixed to the maximum capacity of the design being tested. Once all the coins had been allocated (Fig. 1b), the samples were revealed (Fig. 1c) as either of good (orange) or bad quality (purple) apricots. The sampling outcomes X_i at each supplier i (given the range 1 to 10, or 1 to 32) followed a binomial distribution $X_i \sim B(n_i, p_i)$ where n_i is the number of samples allocated in supplier i and p_i is the fraction of good quality apricots in that supplier. While n_i is chosen by the participants, p_i is unknown to them. Based on the information collected, participants could estimate p_i , and based on the estimation could choose amongst the sampled suppliers to perform a final bulk purchase of 100 apricots. The number of good quality apricots contained in the purchase was revealed (Fig. 1d), and the next trial (purchase cycle) started. The cumulative sum of good-quality apricots collected, as well as a bar informing participants about their progress, were displayed at the bottom of the screen throughout the experiment.

Independently in each trial and for each supplier i , the fraction p_i of good apricots was randomly drawn from a beta distribution (with parameters α, β). We considered three different environments, varying in the relative abundance of good apricots (p_i), denoted poor ($\alpha=1/3, \beta=1$), neutral ($\alpha=3, \beta=3$) and rich ($\alpha=1, \beta=1/3$). Participants were either presented to one environment throughout the experiment (between-subject designs) or to the three environments in different blocks of the same experiment with block order counterbalanced between participants (within-subject designs – see below for more details).

Four experimental designs were run varying in the number of environments presented and the sampling capacity (see Table 1 for summary). In the two between-subject (B) designs, participants were presented only with one environment (rich, neutral or poor), and with one of two capacities, narrow capacity (up to 10 – design B10) or wide capacity (up to 32 – design B32). In the within-subject (W) designs, participants were presented with all three environments (in blocks), either with narrow capacity (up to 10 – design W10) or with wide capacity (up to 32 – design W32). In all designs, participants were presented with 8 repetitions of each sampling capacity for each environment (one, or three depending on the design). All participants were first presented with a practice block composed of 10 trials covering the whole range of sampling capacities in their design (these trials were excluded from the analyses). In the middle of the practice block, participants had to take a short quiz to assess their understanding of the task. They were then provided with feedback and in case of an insufficient score (less than 6 correct answers out of 8 questions), they had to redo the quiz for their participation to be accepted. The whole experiment was self-paced, and opportunities were given to participants to rest in halfway through and after each block.

Participants

Participants were recruited through the online platform Prolific (© 2021 Prolific, <https://www.prolific.co/>), with the criteria of being fluent in English, and with age between 18-52 years old. They were asked to perform the task using a touch screen

device (smartphone or tablet for designs with 10 alternatives, only tablet for designs with 32 alternatives). Participants received a fixed monetary compensation of 8€ per hour and, to increase their motivation, an additional reward was attributed depending on their final score at the task. In between-subject designs, participants who obtained the first and second top scores in each environment received respectively 20€ and 10€. In within-subject designs, the players with the three top scores overall were rewarded with 20, 10 and 5€, respectively.

Participants were recruited until completing a valid final sample size of 45 participants in each of the between-subject designs (B10 and B32 – 15 participants for each environment), and 18 participants in each of the within-subject designs (W10 and W32). The final sample size included in the study was 126 (84 males, mean age $\pm$ sd: 25.7 $\pm$ 6.9). Sample size was decided before starting data collection. In between-subject designs, sample size was calculated (using GPower, Erdfelder, Faul, Buchner, & Lang, 2009) to achieve a 95% power in discriminating participants strategy between environments and was based on the expected effect size (estimated at Cohen's $d=.467$) of the environment on the sampling strategy ($M \sim C$). This effect size was calculated based on the expected effect of the environment following the ideal observer (see below for more details) and the variance observed within each group (s.d. = 0.3) in a previous pilot study. In within-subjects design, the expected effect size was harder to estimate due to a likely contamination in participants' strategies between environments. As we needed to counterbalance the presentation order of the three environments, the sample size had to be a multiple of 6 and thus simply we opted for 18 participants.

Data from an additional 26 participants were discarded before analysis based on pre-established criteria: the use of a wrong device (e.g., computer), insufficient score at the instructions quiz (less than 6 out of 8), a score not significantly higher than chance in the task, and the selection of the 'worst supplier' (the one presenting, after sampling, the smaller proportion of good-quality apricots) on 10% or more of the trials. We considered that such mistakes could occur when participants were confused and based their choice on the opposite colour or when they simply did not pay enough attention. As a result, 3 participants (in B32) were excluded because they used a wrong device, one participant (in W32) because s/he did not pass the instructions quiz, 8 because of a score not significantly higher than chance (3 in B10, 1 in W10, 2 in B32 and 2 in W32), 9 because of more than 10% 'worst choice' (4 in B10, 3 in W10 and 2 in W32) and 5 because they did not satisfy the two last criteria (1 in W10, 4 in B32).

Analyses

Analyses were run using R and MATLAB. Normality of the data was tested using Shapiro tests and homoscedasticity was tested using F tests or Bartlett tests (for more than 2 samples). In cases where it was possible, parametric tests were preferred, otherwise non-parametric tests were used. One sample Wilcoxon tests against the environment averaged outcome (25, 50 and 75 respectively for the poor, neutral and rich environment) were used to test whether participant's final score was significantly different than chance.

Sampling strategy

Our objective is to investigate how humans allocate limited resources to gather information about alternatives whose probability of success is unknown a priori. The resources are represented in our study by the number of samples available (sampling capacity, or coins) on a trial-by-trial basis. We first focus on how this limited capacity influences the number of alternatives participants chose to sample. To do so we introduce the notion of ‘sampling strategy’ characterised as the number of alternatives sampled (M) depending on the sampling capacity (C) available (Figure 2). The lower the ratios of M/C the more the strategy tends to depth, whereas a ratio of 1 indicates pure breadth. We observe how closely empirical behaviour relates to the optimal strategy predicted theoretically (see Moreno-Bote, Ramírez-Ruiz, Drugowitsch, & Hayden, 2020) and whether it is affected by contextual parameters such as the overall probability of success (environment richness).

Effect of environment on sampling strategy

In order to characterize how participants’ sampling strategies might differ depending on the richness of the environment, we decided to test three models to fit M as a function of capacity C separately to data from each participant and environment:

1. a piece-wise power law model (W): $M(C) = \begin{cases} C^{a_1} & \text{if } C \leq B \\ C^{a_2} + b & \text{if } C > B \end{cases}$
where B corresponds to the breakpoint with $B \in \{3, 4, \dots, 9\}$ in narrow-capacity designs, and $B \in \{3, 4, \dots, 16\}$ in wide-capacity designs.
2. a linear model (L): $M(C) = aC + b$
3. a power-law model (P): $M(C) = C^a$

Linear and power law models were compared using a paired Wilcoxon test on individual R-squared adjusted while power law and piece-wise power law models were compared using ANOVA (with $\alpha = .05$) at the participant and environment levels. The power law model captured the empirical relationship between C and M best.

Given the results of the model comparisons, the effect of the environment on participants’ sampling strategies was therefore tested using the power law model. We compared the power factor a extracted from power law fits in each environment using within (designs W10 and W32) or between (designs B10 and B32) one-way ANOVA. Post-hoc comparisons between environments were conducted using t-tests with Bonferroni correction for multiple comparisons. The four experimental designs were analysed individually, in order to account for the structure of the data (within- or across-subjects).

Observed versus optimal sampling strategy

Fitting M as a function of capacity, we previously observed that even when participants’ sampling strategy falls close to the predicted optimal one (see Fig.2), they do not coincide completely. To better characterise participants’ deviations from optimality we first describe them overall, independently from the capacity, by fitting the optimal number of alternatives sampled (M_{opt}) as a function of capacity using the power-law model (described above) as it was the one better accounting for our data. We compared both power-law factors a using one-sample t-tests in each environment individually,

correcting for multiple comparisons, and observed significant differences, going in counter-directions depending on the environment. To further investigate these deviations, we investigate if capacity could be a cause of modulations. To do so, we computed the differences between the number of alternatives observed M and M_{opt} for each capacity and environment (Fig.S1) and tested for marginal effects of the capacity and the environment, as well as for an interaction between capacity and environment using a Sheirer-Ray-Hare test. All effects were found significant and post-hoc analyses were performed using one-sample Wilcoxon tests (against the null hypothesis $\mu = 0$, see Table.S1) in order to better understand how both capacity and environment richness affect the nature and directions of participants' sampling deviations from optimality.

Model comparison

The breadth-depth (BD) dilemma can be solved using extreme behaviours (pure-depth or pure-breadth) but also trade-offs which establish a balance between breadth and depth and often lead to more optimal choices. In the current study, we quantitatively assess which strategy describes best the individual participants' sampling behaviour by comparing the ability of the optimal and heuristic models (characterising both extreme and trade-off behaviours) to predict the data. The number of vendors M that participant k chose in trial j , denoted M_j^k , was predicted by using one of three models, corrupted by either Gaussian or Binomial noise. Each model predicts the mean $M(C_j^k)$ number of selected vendors given that capacity was C_j^k . We use the following models:

1. *Optimal model*: $M(C)$ follows the ideal observer $M_{opt}(C)$

$$M(C) = M_{opt}(C)$$

The ideal observer corresponds to the sample allocation that would maximize expected reward for each capacity in each environment. It describes the optimal number of suppliers to be sample M , as well as the distribution of samples per supplier. The ideal observer was described in a previous theoretical paper (Moreno-Bote et al., 2020) and here has been calculated with Monte-Carlo simulations.

2. *Linear models*. We used three variants of linear model:

- a) *Depth model*: $M(C)$ is *constant* and independent of capacity C , with M following an almost pure-depth strategy,

$$M(C) = 2 .$$

- b) *Pure-breadth model*: $M(C)$ always equals to C ,

$$M(C) = C .$$

- c) *Free-slope linear model*: $M(C)$ increases with capacity C with a free factor d ,

$$M(C) = dC .$$

3. *Power Law* model. We used two variants of this model:

a) *Square root* model: $M(C)$ increases with the square-root of capacity C .

$$M(C) = \sqrt{C}.$$

b) *Free power-law* model: $M(C)$ increases with capacity C to the w -th power,

$$M(C) = C^w.$$

For each model, to compute the likelihood of the data M_j^k for each trial j and participant k , we assume a noise model that is either Gaussian or binomial. In the Gaussian noise model, we assume that $M_j^k = M(C_j^k) + \sigma_j^k z_j^k$, where z_j^k is independently and identically distributed standard normal noise, and $\sigma_j^k = a^k + b^k C_j^k$. In all cases we obtained that the best fit parameters obeyed $a^k \neq 0, b^k \neq 0$, so that $\sigma_j^k \neq 0$ and $(\sigma_j^k)^2 > 0$ for all values of C . The per-participant likelihood of the model is the sum of individual trial-by-trial likelihoods

$$\mathcal{L}_N(a^k, b^k) = \sum_j \left(-\log(\sqrt{2\pi}) - \log(\sigma_j^{k^2}) - \frac{(M_j^k - M(C_j^k) - \mu^k)^2}{2\sigma_j^{k^2}} \right)$$

In the binomial case, we assume that the observed M_j^k follows a binomial distribution $M_j^k \sim B(n^k, p^k)$ with $p^k \in [0,1]$ and

$$n^k = \frac{M(C_j^k) + \mu^k}{p^k}, \quad n^k = 1, 2, \dots, N \text{ and } n^k \geq M_j^k$$

so that the expectation of M equals the desired model's expectation $M(C)$. In this case, the per-participant likelihood of the data under the model and binomial noise is

$$\mathcal{L}_B(p^k) = \sum_j \left(\log \left(\frac{n^k!}{M_j^k! (n^k - M_j^k)!} \right) + M_j^k \log(p^k) + (n^k - M_j^k) \log(1 - p^k) \right)$$

The overall log likelihood (LL) of each model was the sum of log likelihoods across trials, participants and environments.

Each model (1-3 above) with each noise model (Gaussian or binomial) was fit and tested to each participant and environment (in the case of within-subjects designs) individually using four-fold cross-validation. Specifically, in one-fold, 75% of the data in each participant was used to fit the per-participant model parameters by maximizing their individual log likelihoods. The remaining 25% of the data was used to measure the quality of the fit. This was quantified as the log likelihood of the model with its fitted parameters in the test set, and summed across participants, and averaged over the four

folds (4 different ways of taking 25% non-overlapping trials for testing and the remaining for training). We call the resulting average cross-validated log-likelihood (CVLL). In the binomial noise model, the log likelihood of the test set becomes minus infinity in trials when $n < M$, which is an impossible event. This concerned a small set of trials in 27 participants in the square root model, 31 in the optimal model, 16 in the depth model, 69 in the free power law model and 48 in the linear model. However, in all cases it affected only one out of the of the four folds, so the averaged log likelihood could be computed across the three remaining folds. Model comparison was performed by comparing the CVLL across models and noise models using Wilcoxon matched pairs signed-ranks tests with Bonferroni correction. In order to ascertain the results of the LL-based model comparisons, the models were also compared using Akaike Information Criteria (AIC, Akaike, 1977), which were computed on the whole dataset. Both methods provide qualitatively identical results (compare Fig.4 with Fig.S5 for Binomial noise models and Fig.S4 with Fig.S6 for Gaussian noise models). Free parameter a was optimized on a grid from -3 to 1 and b on a grid from -1 to 1, both with a step of 0.05 in the Gaussian noise model and the parameter p in the Binomial noise model which was optimized on grid from 0 to 1 with a step of 0.01. The free parameter w was optimized on a grid from 0.05 to 1.1 with a step of 0.05 in both Gaussian and Binomial noise models.

Exploring how samples are dispatched within the alternatives sampled (splits)

To descript more precisely participants' sampling strategy, we studied how the samples are allocated inside the selected suppliers. Indeed, for an identical number of alternatives sampled M and capacity C , we can observe various allocations of samples. For example, $M = 2$ and $C = 4$ can result in the allocation of 2 samples in 2 different alternatives each; $\{2,2\}$, or in the allocation of 3 samples in a first alternative and 1 sample in a second; $\{3,1\}$. Visually, we observe that participants have a tendency to allocate the samples homogeneously across the sampled alternatives (e.g. $\{2,2\}$) (see Fig.5A). In order to statistically capture this putative bias, we computed the standard deviations of the sample allocations (SA). An homogenous allocation of samples would result in a standard deviations close to zero (e.g. $\text{sd}(\{2,2\}) = 0$) whereas a more heterogeneous allocation of samples would be associated with higher standard deviations (e.g. $\text{sd}(\{1,3\}) \approx 1.41$). We computed the standard deviation of each sample allocation and compared it to the standard deviation of the optimal allocation of samples (predicted by Moreno-Bote et al., 2020) in each environment using Wilcoxon tests with Bonferroni corrections to adjust for the multiple comparisons. We tested the effect of the environment on the magnitude of these differences using a Kruskal-Wallis test and performed post-hoc analyses using Wilcoxon tests with Bonferroni corrections.

We also computed the differences between the outcomes observed in our data with the ones obtained when following the ideal observer strategy. A Kruskal test was performed to test the significant effect of the environment on the outcome differences and one-sample Wilcoxon tests with Bonferroni corrections were then used to assess whether observed outcomes were significantly different than optimal outcomes inside each environment condition.

These analyses were performed for all trials (from all experimental designs) where the optimal number of alternatives sampled was strictly inferior to the capacity ($M_{opt} < C$

(see Fig. 5B and 5C). This condition has been introduced because when $M_{opt} = C$, the optimal SA follows a pure-breadth strategy, and its associated standard deviation is zero. Therefore, the standard deviation of the SA observed in the data can only be superior to the optimal SA which impairs the observation of participants' bias to sample homogenously the alternatives.

Bibliography

- Akaike, H. (1977). On entropy maximization principle. *Undefined*.
- Attneave, F. (1955). Symmetry, information, and memory for patterns. *The American Journal of Psychology*, 68(2), 209–222. <https://doi.org/10.2307/1418892>
- Blanchet-Scalliet, C., El Karoui, N., Jeanblanc, M., & Martellini, L. (2008). Optimal investment decisions when time-horizon is uncertain. *Journal of Mathematical Economics*, 44(11), 1100–1113. <https://doi.org/10.1016/J.JMATECO.2007.09.004>
- Brown, S., Steyvers, M., & Wagenmakers, E. J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology*, 53(6), 453–462. <https://doi.org/10.1016/j.jmp.2009.09.002>
- Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-based Decisions. *Trends in Cognitive Sciences*, 23(3), 251–263. <https://doi.org/10.1016/J.TICS.2018.12.003>
- Callaway, F., Rangel, A., & Griffiths, T. L. (2021). Fixation patterns in simple choice reflect optimal information sampling. *PLoS Computational Biology*, 17(3), e1008863. <https://doi.org/10.1371/JOURNAL.PCBI.1008863>
- Chau, B. K. H., Kolling, N., Hunt, L. T., Walton, M. E., & Rushworth, M. F. S. (2014). A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nature Neuroscience*, 17(3), 463–470. <https://doi.org/10.1038/nn.3649>
- Cohen, A. L., Kang, N., & Leise, T. L. (2017). Multi-attribute, multi-alternative models of choice: Choice, reaction time, and process tracing. *Cognitive Psychology*, 98, 45–72. <https://doi.org/10.1016/J.COGLPSYCH.2017.08.001>
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942. <https://doi.org/10.1098/rstb.2007.2098>
- Costa, V. D., Mitz, A. R., & Averbeck, B. B. (2019). Subcortical Substrates of Explore-Exploit Decisions in Primates. *Neuron*, 103(3), 533-545.e5. <https://doi.org/10.1016/j.neuron.2019.05.017>
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879. <https://doi.org/10.1038/nature04766>
- Drugowitsch, J., Wyart, V., Devauchelle, A. D., & Koechlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron*, 92(6), 1398–1411. <https://doi.org/10.1016/j.neuron.2016.11.005>

- Ebitz, R. B., Albarran, E., & Moore, T. (2018). Exploration Disrupts Choice-Predictive Signals and Alters Dynamics in Prefrontal Cortex. *Neuron*, 97(2), 450-461.e9. <https://doi.org/10.1016/j.neuron.2017.12.007>
- Erdfelder, E., Faul, F., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41(4), 1149–1160. <https://doi.org/10.3758/BRM.41.4.1149>
- Feng, S. F., Schwemmer, M., Gershman, S. J., & Cohen, J. D. (2014). Multitasking versus multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cognitive, Affective, & Behavioral Neuroscience* 2014 14:1, 14(1), 129–146. <https://doi.org/10.3758/S13415-013-0236-9>
- Fitts, P. M. (1966). Cognitive aspects of information processing: III. Set for speed versus accuracy. *Journal of Experimental Psychology*, 71(6), 849–857. <https://doi.org/10.1037/h0023232>
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational Use of Cognitive Resources: Levels of Analysis Between the Computational and the Algorithmic. *Topics in Cognitive Science*, 7, 217–229. <https://doi.org/10.1111/tops.12142>
- Halpert, H. (1958). Folklore: Breadth versus Depth. *The Journal of American Folklore*, 71(280), 97–103. <https://www.jstor.org/stable/537679?seq=1>
- Hawkins, G., Brown, S. D., Steyvers, M., & Wagenmakers, E. J. (2012). Context Effects in Multi-Alternative Decision Making: Empirical Data and a Bayesian Model. *Cognitive Science*, 36(3), 498–516. <https://doi.org/10.1111/j.1551-6709.2011.01221.x>
- Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, 4(1), 11–26. <https://doi.org/10.1080/17470215208416600>
- Iyengar, S. S., & Lepper, M. R. (2000). PERSONALITY PROCESSES AND INDIVIDUAL DIFFERENCES When Choice is Demotivating: Can One Desire Too Much of a Good Thing? <https://doi.org/10.1037/0022-3514.79.6.995>
- Jang, A. I., Sharma, R., & Drugowitsch, J. (2021). Optimal policy for attention-modulated decisions explains human fixation behavior. *ELife*, 10. <https://doi.org/10.7554/eLife.63436>
- Joseph, S., Teki, S., Kumar, S., Husain, M., & Griffiths, T. D. (2016). Resource allocation models of auditory working memory. *Brain Research*, 1640, 183–192. <https://doi.org/10.1016/J.BRAINRES.2016.01.044>
- Klöckner, K., Wirschum, N., & Jameson, A. (2004). *Depth-and Breadth-First Processing*

of Search Result Lists.

- Krajovich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10), 1292–1298. <https://doi.org/10.1038/nn.2635>
- Kuksov, D., & Villas-Boas, J. M. (2010). When more alternatives lead to less choice. *Marketing Science*, 29(3), 507–524. <https://doi.org/10.1287/mksc.1090.0535>
- Lieder, F., Griffiths, T. L., & Goodman, N. D. (2012). Burn-in, bias, and the rationality of anchoring. *Advances in Neural Information Processing Systems*, 4, 2690–2698.
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, 17(3), 347. <https://doi.org/10.1038/NN.3655>
- March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science*, 2(1), 71–87.
- Miller, D. P. (1981). The Depth/Breadth Tradeoff in Hierarchical Computer Menus. *Proceedings of the Human Factors Society Annual Meeting*, 25(1), 296–300. <https://doi.org/10.1177/107118138102500179>
- Moreno-Bote, R., & Mastrogiuseppe, C. (2021). *Deep imagination is a close to optimal policy for planning in large decision trees under limited resources*. 1–24.
- Moreno-Bote, R., Ramírez-Ruiz, J., Drugowitsch, J., & Hayden, B. Y. (2020). Heuristics and optimal solutions to the breadth:depth dilemma. *Proceedings of the National Academy of Sciences of the United States of America*, 117(33), 19799–19808. <https://doi.org/10.1073/PNAS.2004929117>
- Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. *Trends in Cognitive Sciences*, 25(9), 757–775. <https://doi.org/10.1016/J.TICS.2021.06.001>
- Pratt, S. C., Mallon, E. B., Sumpter, D. J. T., & Franks, N. R. (2002). Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *Leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 52(2), 117–127. <https://doi.org/10.1007/s00265-002-0487-x>
- Proctor, R. W., & Schneider, D. W. (2018). Hick’s law for choice reaction time: A review. *Quarterly Journal of Experimental Psychology (2006)*, 71(6), 1281–1299. <https://doi.org/10.1080/17470218.2017.1322622>
- Ramírez-Ruiz, J., & Moreno-Bote, R. (2021). *Optimal allocation of finite sampling capacity in accumulator models of multi-alternative decision making*. <http://arxiv.org/abs/2102.01597>

- Reilly, G., Souder, D., & Ranucci, R. (2016). Time Horizon of Investments in the Resource Allocation Process: Review and Framework for Next Steps. *Journal of Management*, 42(5), 1169–1194. <https://doi.org/10.1177/0149206316630381>
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative Decision Field Theory: A Dynamic Connectionist Model of Decision Making. *Psychological Review*, 108(2), 370–392. <https://doi.org/10.1037/0033-295X.108.2.370>
- Russell, S. J., & Subramanian, D. (1995). Provably Bounded-Optimal Agents. *Journal of Artificial Intelligence Research*, 2, 575–609. <https://doi.org/10.1613/jair.133>
- Schustek, P., Hyafil, A., & Moreno-Bote, R. (2019). Human confidence judgments reflect reliability-based hierarchical integration of contextual information. *Nature Communications*, 10(1), 1–15. <https://doi.org/10.1038/s41467-019-13472-z>
- Schwartz, M. S., Sadler, P. M., Sonnert, G., & Tai, R. H. (2009). Depth versus breadth: How content coverage in high school science courses relates to later success in college science coursework. *Science Education*, 93(5), 798–826. <https://doi.org/10.1002/sce.20328>
- Sethi-Iyengar, S., Huberman, G., Jiang, W., Sethi-Iyengar, S. ;, & Huberman, G. ; (2003). *How Much Choice is Too Much?: Contributions to 401(k) Retirement Plans*. https://repository.upenn.edu/prc_papers/426
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron*, 79(2), 217–240. <https://doi.org/10.1016/J.NEURON.2013.07.007>
- Sigman, M., & Dehaene, S. (2005). Parsing a cognitive task: A characterization of the mind's bottleneck. *PLoS Biology*, 3(2), 0334–0349. <https://doi.org/10.1371/journal.pbio.0030037>
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>
- Turner, S., Bettis, R., & Burton, R. (2002). Exploring Depth Versus Breadth in Knowledge Management Strategies. *Computational and Mathematical Organization Theory*, 8(1), 49–73. <https://doi.org/10.1023/A:1015180220717>
- Usher, M., & McClelland, J. L. (2004). *Loss Aversion and Inhibition in Dynamical Models of Multialternative Choice*. <https://doi.org/10.1037/0033-295X.111.3.757>
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599–637. <https://doi.org/10.1111/cogs.12101>
- Wickelgren, W. A. (1977). Speed-Accuracy Tradeoff and information processing

dynamics. *Acta Psychologica*, 41, 67–85.
https://doi.org/10.1007/springerreference_183986

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore-exploit dilemma. *Journal of Experimental Psychology: General*, 143(6), 2074–2081.
<https://doi.org/10.1037/a0038199>

Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. In *Current Opinion in Behavioral Sciences* (Vol. 11, pp. 109–115). Elsevier Ltd. <https://doi.org/10.1016/j.cobeha.2016.07.003>

Supplementary Figures

Figure S1. Participants' sampling strategy seems to be influenced by the previous environment presented. Number of alternatives sampled M as a function of capacity averaged across participants (points) for the 3 different environments (left: poor, centre: neutral, right: rich) and depending on its presentation order (colours; first environment presented or presented after a poor, neutral or rich environment), for within-subject design with narrow (W10, panel A) or wide capacities (W32, panel B).

Figure S2. Participants' sampling strategy deviates from optimality and tends to be tilted toward depth at low capacity and breadth at high capacity. Difference between the optimal (M_{opt}) and the observed (M) number of alternatives sampled averaged across all participants for each capacity and environment (coloured lines). Error bars represent the standard error of the mean and significant deviations from 0 (dashed line) are marked as follows: '*': $p_{adj} < .05$, '**': $p_{adj} < .01$, '***': $p_{adj} < .001$, '****': $p_{adj} < .0001$.

environment capacity	poor	neutral	rich
2	$V_{66} = 231.0$ $p_{adj} = 1.56 \times 10^{-3}$	$V_{66} = 276.0$ $p_{adj} = 7.39 \times 10^{-4}$	$V_{66} = 171.0$ $p_{adj} = 5.97 \times 10^{-3}$
3	$V_{33} = 190.0$ $p_{adj} = 3.99 \times 10^{-3}$	$V_{33} = 435.0$ $p_{adj} = 8.38 \times 10^{-5}$	$V_{33} = 19.5$ $p_{adj} = 3.80 \times 10^{-4}$
4	$V_{66} = 1128.0$ $p_{adj} = 7.89 \times 10^{-8}$	$V_{66} = 1075.5$ $p_{adj} = 1$	$V_{66} = 1196.0$ $p_{adj} = 1$
5	$V_{33} = 435.0$ $p_{adj} = 8.65 \times 10^{-5}$	$V_{33} = 410.0$ $p_{adj} = .00881$	$V_{33} = 152.0$ $p_{adj} = 1$
6	$V_{33} = 300.0$ $p_{adj} = 6.10 \times 10^{-4}$	$V_{33} = 331.5$ $p_{adj} = 1$	$V_{33} = 134.0$ $p_{adj} = 1$
7	$V_{33} = 473.0$ $p_{adj} = .0197$	$V_{33} = 443.5$ $p_{adj} = .0268$	$V_{33} = 267.5$ $p_{adj} = 1$
8	$V_{66} = 1930.0$ $p_{adj} = 6.96 \times 10^{-7}$	$V_{66} = 1424.0$ $p_{adj} = .340$	$V_{66} = 808.5$ $p_{adj} = 1$
9	$V_{33} = 419.5$ $p_{adj} = .123$	$V_{33} = 287.0$ $p_{adj} = 1$	$V_{33} = 199.0$ $p_{adj} = 1$
10	$V_{33} = 428.0$ $p_{adj} = .0736$	$V_{33} = 420.5$ $p_{adj} = .417$	$V_{33} = 118.0$ $p_{adj} = .366$
16	$V_{33} = 365.0$ $p_{adj} = 1$	$V_{33} = 236.0$ $p_{adj} = 1$	$V_{33} = 64.5$ $p_{adj} = 6.53 \times 10^{-3}$
32	$V_{33} = 266.0$ $p_{adj} = 1$	$V_{33} = 43.0$ $p_{adj} = 7.52 \times 10^{-4}$	$V_{33} = 28.0$ $p_{adj} = 2.21 \times 10^{-4}$

Table S1. Participants' sampling strategy deviates from optimality and tends to be tilted toward depth at low capacity and breadth at high capacity. Results of one-sample Wilcoxon tests against the null hypothesis ($\mu=0$, no bias towards depth nor breadth) of the differences between the optimal (M_{opt}) and the observed (M) number of alternatives sampled averaged per participant for each capacity and environment. V statistics and p-values adjusted with Bonferroni corrections for multiple comparisons are reported and significant results ($\alpha < .05$) are highlighted in bold.

Figure S3. The standard deviation of the number of alternatives sampled M increases linearly with the capacity. Each point represents the standard deviation of M averaged across participants for each experimental design, environment and capacity.

Figure S4. The free power-law model is better at predicting participants' sampling strategy than both the optimal and other heuristics models using Gaussian distributed noise. Averaged negative cross-validated log-likelihood (CVLL) across participants for each model overall (A) and in each environment (B). Results of pair-wise comparisons are displayed according to adjusted p-values ('*': $p < .05$, '**': $p < .01$, '***': $p < .001$) and error bars correspond to s.e.m. C. Distribution of the power factor w in the free power model. Each colour dot represents a subject, black dots represent distribution means and bars 95% confidence intervals. Results of post-hoc comparisons are displayed according to adjusted p-values ('ns': $p > .05$, '**': $p < .01$, '***': $p < .001$).

	depth	square root	linear	optimal	power
pure breadth	$V_{126} = 3748,$ $p_{adj} = 1$	$V_{126} = 3582,$ $p_{adj} = 1$	$V_{126} = 1443,$ $p_{adj} = 7.32 \times 10^{-4}$	$V_{126} = 991,$ $p_{adj} = 3.56 \times 10^{-12}$	$V_{126} = 802,$ $p_{adj} = 5.31 \times 10^{-12}$
depth		$V_{126} = 4039,$ $p_{adj} = 1$	$V_{126} = 2501,$ $p_{adj} = .004$	$V_{126} = 1252,$ $p_{adj} = 3.33 \times 10^{-10}$	$V_{126} = 1025,$ $p_{adj} = 6.56 \times 10^{-12}$
square root			$V_{126} = 2910,$ $p_{adj} = .119$	$V_{126} = 2347,$ $p_{adj} = 8.54 \times 10^{-4}$	$V_{126} = 799,$ $p_{adj} = 4.23 \times 10^{-13}$
linear				$V_{126} = 3485,$ $p_{adj} = 1$	$V_{126} = 480,$ $p_{adj} = 6.08 \times 10^{-15}$
optimal					$V_{126} = 2630,$ $p_{adj} = .013$

Table S2. Summary of the pair-wise comparisons (Wilcoxon Matched Pairs Signed-Ranks test) of the 4-folds averaged CVLL between all six models using Gaussian distributed noise. P-values are adjusted with Bonferroni corrections and significative differences ($p < .05$) are highlighted. Models are ordered from worst (pure breadth) to best (free power-law).

Designs tests	B10	W10	B32	W32
ANOVA	$K - W \chi^2_2 = 16.65,$ $p = 2.43 \times 10^{-4}$	$F_1 = 23.21,$ $p = 1.61 \times 10^{-4}$	$F_1 = 1.62,$ $p = .21$	$F_1 = 8.68,$ $p = .009$
Poor vs. Neutral	$W_{15} = 206,$ $p_{adj} = 3.63 \times 10^{-4}$	$t_{17} = 3.81,$ $p_{adj} = .004$	-	$t_{17} = 1.37,$ $p_{adj} = .56$
Poor vs. Rich	$W_{15} = 188,$ $p_{adj} = .005$	$t_{17} = 4.82,$ $p_{adj} = 4.83 \times 10^{-4}$	-	$t_{17} = 2.95,$ $p_{adj} = .027$
Neutral vs. Rich	$W_{15} = 126,$ $p_{adj} = 1$	$t_{17} = 0.42,$ $p_{adj} = 1$	-	$t_{17} = 2.32,$ $p_{adj} = .10$

Table S3. Summary of the comparisons between the averaged individual 4-fold CVLL in each environment and experimental design using Gaussian distributed noise. In the first row are displays results of the effect of environment on the value of the averaged CVLL using ANOVA or non-parametric equivalent. Underneath are displayed the results of post-hoc comparisons between each environment (if applicable). P-values are adjusted with Bonferroni corrections and significative differences ($p < .05$) are highlighted.

Figure S5. The free power-law model is better at predicting participants' sampling strategy than both the optimal and other heuristics models using Binomial distributed noise. Averaged negative AIC across participants for each model overall (**A**) and in each environment (**B**). Results of pairwise comparisons are displayed according to adjusted p-values (*: $p < .05$, **: $p < .01$, ***: $p < .001$) and error bars correspond to s.e.m. **C.** Distribution of the power factor w in the free power model. Each colour dot represents a subject, black dots represent distribution means and bars 95% confidence intervals. Results of post-hoc comparisons are displayed according to adjusted p-values (ns: $p > .05$, **: $p < .01$, ***: $p < .001$).

	pure breadth	sqrt	optimal	linear	power
depth	$V_{126} = 12619,$ $p_{adj} = .009$	$V_{126} = 19671,$ $p_{adj} = 7.29 \times 10^{-33}$	$V_{126} = 19177,$ $p_{adj} = 1.09 \times 10^{-29}$	$V_{126} = 19480,$ $p_{adj} = 1.28 \times 10^{-31}$	$V_{126} = 19691,$ $p_{adj} = 5.37 \times 10^{-33}$
pure breadth		$V_{126} = 12519,$ $p_{adj} = .014$	$V_{126} = 16912,$ $p_{adj} = 3.32 \times 10^{-17}$	$V_{126} = 19681,$ $p_{adj} = 6.26 \times 10^{-33}$	$V_{126} = 19673,$ $p_{adj} = 7.06 \times 10^{-33}$
sqrt			$V_{126} = 16429,$ $p_{adj} = 5.55 \times 10^{-15}$	$V_{126} = 17505,$ $p_{adj} = 3.80 \times 10^{-20}$	$V_{126} = 19560,$ $p_{adj} = 3.88 \times 10^{-32}$
optimal				$V_{126} = 15310,$ $p_{adj} = 2.04 \times 10^{-10}$	$V_{126} = 19295,$ $p_{adj} = 1.96 \times 10^{-30}$
linear					$V_{126} = 16459,$ $p_{adj} = 6.61 \times 10^{-27}$

Table S4. Summary of the pair-wise comparisons (Wilcoxon Matched Pairs Signed-Ranks test) of the individual AIC between all six models using Binomial distributed noise. P-values are adjusted with Bonferroni corrections and significative differences ($p < .05$) are highlighted. Models are ordered from worst (depth) to best (power).

Designs tests	B10	W10	B32	W32
ANOVA	$K - W \chi^2_2 = 14.18,$ $p = 8.35 \times 10^{-4}$	$F_1 = 25.56,$ $p = 9.75 \times 10^{-6}$	$F_1 = 0.92,$ $p = .34$	$F_1 = 13.62,$ $p = .0018$
Poor vs. Neutral	$W_{15} = 200,$ $p_{adj} = 8.88 \times 10^{-4}$	$t_{17} = 3.22,$ $p_{adj} = .015$	-	$t_{17} = 2.48,$ $p_{adj} = .071$
Poor vs. Rich	$W_{15} = 181,$ $p_{adj} = .013$	$t_{17} = 5.06,$ $p_{adj} = 2.92 \times 10^{-4}$	-	$t_{17} = 3.69,$ $p_{adj} = .005$
Neutral vs. Rich	$W_{15} = 126,$ $p_{adj} = 1$	$t_{17} = 0.218,$ $p_{adj} = 1$	-	$t_{17} = 1.67,$ $p_{adj} = .342$

Table S5. Summary of the comparisons between the individual AIC in each environment and experimental design using Binomial distributed noise. In the first row are displays results of the effect of environment on the value of AIC using ANOVA or non-parametric equivalent. Underneath are displayed the results of post-hoc comparisons between each environment (if applicable). P-values are adjusted with Bonferroni corrections and significative differences ($p < .05$) are highlighted.

Figure S6. The free power-law model is better at predicting participants' sampling strategy than both the optimal and other heuristics models using Gaussian distributed noise. Averaged negative AIC across participants for each model overall (**A**) and in each environment (**B**). Results of pairwise comparisons are displayed according to adjusted p-values (*: $p < .05$, **: $p < .01$, ***: $p < .001$) and error bars correspond to s.e.m. **C**. Distribution of the power factor w in the free power model. Each colour dot represents a subject, black dots represent distribution means and bars 95% confidence intervals. Results of post-hoc comparisons are displayed according to adjusted p-values ('ns': $p > .05$, **: $p < .01$, ***: $p < .001$).

	pure breadth	depth	optimal	linear	power
sqrt	$V_{126} = 9657,$ $p_{adj} = 1$	$V_{126} = 12670,$ $p_{adj} = .007$	$V_{126} = 16059,$ $p_{adj} = 2.22 \times 10^{-13}$	$V_{126} = 16878,$ $p_{adj} = 4.80 \times 10^{-17}$	$V_{126} = 19532,$ $p_{adj} = 5.85 \times 10^{-32}$
pure breadth		$V_{126} = 10829,$ $p_{adj} = 1$	$V_{126} = 16996,$ $p_{adj} = 1.31 \times 10^{-17}$	$V_{126} = 18221,$ $p_{adj} = 4.26 \times 10^{-24}$	$V_{126} = 19531,$ $p_{adj} = 5.94 \times 10^{-32}$
depth			$V_{126} = 16421,$ $p_{adj} = 6.03 \times 10^{-15}$	$V_{126} = 17137,$ $p_{adj} = 2.78 \times 10^{-18}$	$V_{126} = 19408,$ $p_{adj} = 3.74 \times 10^{-31}$
optimal				$V_{126} = 9909,$ $p_{adj} = 1$	$V_{126} = 15828,$ $p_{adj} = 2.00 \times 10^{-12}$
linear					$V_{126} = 16511,$ $p_{adj} = 1.99 \times 10^{-29}$

Table S6. Summary of the pair-wise comparisons (Wilcoxon Matched Pairs Signed-Ranks test) of the individual AIC between all six models using Gaussian distributed noise. P-values are adjusted with Bonferroni corrections and significant differences ($p < .05$) are highlighted. Models are ordered from worst (square root) to best (power).

Designs tests	B10	W10	B32	W32
ANOVA	$K - W \chi^2_2 = 17.05,$ $p = 1.99 \times 10^{-4}$	$F_1 = 20.22,$ $p = 3.18 \times 10^{-4}$	$F_1 = 1.52,$ $p = .23$	$F_1 = 11.32,$ $p = .0037$
Poor vs. Neutral	$W_{15} = 206,$ $p_{adj} = 3.00 \times 10^{-4}$	$t_{17} = 3.48,$ $p_{adj} = .009$	-	$t_{17} = 2.08,$ $p_{adj} = .16$
Poor vs. Rich	$W_{15} = 190,$ $p_{adj} = .004$	$t_{17} = 4.50,$ $p_{adj} = 9.54 \times 10^{-4}$	-	$t_{17} = 3.36,$ $p_{adj} = .011$
Neutral vs. Rich	$W_{15} = 123,$ $p_{adj} = 1$	$t_{17} = 0.362,$ $p_{adj} = 1$	-	$t_{17} = 2.29,$ $p_{adj} = .11$

Table S7. Summary of the comparisons between the individual AIC in each environment and experimental design using Gaussian distributed noise. In the first row are displays results of the effect of environment on the value of AIC using ANOVA or non-parametric equivalent. Underneath are displayed the results of post-hoc comparisons between each environment (if applicable). P-values are adjusted with Bonferroni corrections and significative differences ($p < .05$) are highlighted.