Competing cognitive pressures on human exploration in the absence of trade-off with exploitation

Clémence Alméras^{1,3,*}, Valérian Chambon^{2,3}, and Valentin Wyart^{1,3,*}

¹ Laboratoire de Neurosciences Cognitives et Computationnelles, Institut National de la Santé et de la Recherche Médicale (Inserm), Paris, France

² Institut Jean Nicod, Centre National de la Recherche Scientifique (CNRS), Paris, France

³ Département d'Études Cognitives, École Normale Supérieure, Université PSL, Paris, France

*corresponding authors: clemence.almeras@gmail.com, valentin.wyart@inserm.fr

Exploring novel environments through sequential sampling is essential for efficient decision-making under uncertainty. In the laboratory, human exploration has been studied in situations where exploration is traded against reward maximisation. By design, these ‘explore-exploit’ dilemmas confound the behavioural characteristics of exploration with those of the trade-off itself. Here we designed a sequential sampling task where exploration can be studied and compared in the presence and absence of trade-off with exploitation. Detailed model-based analyses of choice behaviour revealed specific exploration patterns arising in situations where information seeking is not traded against reward seeking. Human choices are directed toward the most uncertain option available, but only after an initial sampling phase consisting of choice streaks from each novel option. These findings outline competing cognitive pressures on information seeking: the repeated sampling of the current option (for hypothesis testing), and the directed sampling of the most uncertain option available (for structure mapping).

Introduction

Making decisions often requires choosing between exploiting a familiar option whose rewarding value is known and exploring an unfamiliar option whose value is not known^{1–3}. Exploration is essential to discovering the structure of a novel environment (e.g., mapping possible itineraries from home to work in a new city), optimizing behaviour (finding the fastest or most pleasant itinerary) or adapting to changes in the environment (looking for an alternative itinerary when the familiar one is unavailable). Exploration is thus adaptive to acquire information about available options, but also risky: choosing an unfamiliar option means foregoing a known reward for an uncertain one. When the known reward can be obtained without delay (e.g., choosing between two snacks at a vending machine), humans arbitrate between the costs and benefits of exploration⁴.

However, not all choices bring immediate rewards. For example, when browsing restaurant reviews online, exploring different restaurant reviews before ordering yields information at virtually no cost. In such situations, exploration is not traded against immediate reward maximization, but used to acquire information about available options⁵. In some cases, like browsing restaurant reviews, the information acquired concerns the rewarding values of available options and can be readily used when choosing which restaurant to order from. But in other cases, like orienting oneself in a new city, the acquired information does not even concern value – whose definition depends on one's particular goal at a later point in time (e.g., going to work on a weekday, or finding a new friend's place on a Sunday).

In the laboratory, exploration is usually studied using bandit tasks in which agents are presented with two (or more) options to sample from^{6–8}. Each option is associated with a specific reward distribution, and the agents are incentivized to maximize rewards through sequential sampling. Across successive choices, the agents learn to choose the most rewarding option – and thus trade exploration against immediate reward maximization. Variants of these tasks have been developed to dissociate exploration from reward maximization by delaying the acquisition of rewards after an initial sampling phase^{9–11}. However, as it is the case when browsing restaurant reviews, the information acquired during the sampling phase concerns the rewarding values of available options, and choices are still attracted toward sampling the option associated with the highest rewarding value. This observation means that the exploration patterns observed in these tasks with delayed rewards are still biased by the rewarding values of available options.

Here we sought to fully dissociate exploration from reward-guided decision-making by designing a novel two-armed bandit task where the outcomes associated with the two options – shades of colour instead of points or monetary rewards – provide information about option-outcome associations but have no intrinsic value (Figure 1a). In one condition, we asked tested participants to draw one particular colour. This instruction confers a rewarding value to each outcome, and requires participants to trade exploration against exploitation to choose the option associated with the instructed colour, as in regular bandit tasks. By contrast, in the other condition, we asked the same participants to sample freely the two options to learn option-outcome associations – i.e., the task structure. In this condition, participants were rewarded based on their ability to guess correct option-outcome associations using a single two-alternative forced choice at the end of each game. This instruction fully dissociates exploration from any form of immediate or prospective reward maximisation, by making participants choose between options with no rewarding value. Except for this difference, the two conditions were matched in every possible aspect so that sampling patterns could be directly compared between them (Figure 1bc; see Methods).

We collected a first discovery dataset to derive task-based variables to compare the two conditions, including a computational model of human exploration patterns. We found important differences between human exploration in the two conditions. We validated the existence of these effects in two additional datasets, whose data were analysed using the exact same procedures (see Methods). Together, the findings reveal competing forms of information seeking in human exploration when it is not traded against reward maximization: a 'local' form of information seeking, corresponding to the continued sampling of the current option to generate and test a hypothesis regarding its associated outcome, and a 'global' form of information seeking, corresponding to the directed sampling of the most uncertain option available to map the overall task structure.

Results

Important differences in human exploration patterns between conditions

We started by assessing participants' performance in each condition. In the condition where participants were asked to draw a target colour (DRAW), we assessed performance by plotting the fraction of games where participants chose the option associated with the target colour – focusing on games where only one option was associated with the target colour (Figure 1de). As expected, participants ($N = 27$ after exclusion in the discovery and replication datasets; see Methods) chose the option associated with the target colour more often than the other option on the last choice of each game (discovery: median = 81.2%, IQR = [57.8% 93.8%]; replication: median = 84.4%, IQR = [53.9% 93.8%], signed-rank test against chance, $z = 3.9$, $p < 0.001$, $r = 0.75$). In the condition where participants were asked to guess the colour associated with each option (GUESS), we assessed performance as the fraction of correct option-outcome associations provided in the two-alternative forced choice at the end of each game (Figure 1f). As expected, participants guessed better than chance the correct option-outcome associations across all games (median = 81.2%, IQR = [70.3% 92.2%]; replication: median = 81.2%, IQR = [75% 90.6%], signed-rank test against chance, $z = 4.4$, $p < 0.001$, $r = 0.86$).

We then studied participants' choice patterns in the two conditions. First, we computed participants' sampling imbalance as the difference in number of choices between the most- and least-chosen options over the course of each game (Figure 2a). In the DRAW condition, sampling imbalance grew linearly over time (linear coefficient; discovery: mean = 0.43, CI = [0.37 0.48]; replication: mean = 0.43, CI = [0.35 0.49]). This linear growth suggests that participants converged progressively on a preferred option. By contrast, in the GUESS condition, sampling imbalance showed a slower growth (linear coefficient; discovery: mean = 0.20, CI = [0.18 0.23]; replication: mean = 0.26, CI = [0.19 0.32], paired t -test between conditions: $t_{26} = 3.9$, $p = 0.001$, $d = 0.97$), but also a non-linear profile. Indeed, sampling imbalance grew very rapidly in the first trials of each game (deviation from linear fit; discovery: peak at trial 5, mean = +1.72, CI = [+1.35 +2.10]; replication: mean = +1.19, CI = [+0.76 +1.61], t -test against zero, $t_{26} = 5.7$, $p < 0.001$, $d = 1.28$). This early sampling imbalance in the GUESS condition suggests that participants had the tendency to sample repeatedly from the same option in the first trials.

We confirmed this choice pattern using two additional metrics. First, we computed the fraction of re-sampling from the same option as a function of the number of previous samples acquired from this option (Figure 2b). This metric confirmed participants' tendency to sample repeatedly from options for which few samples have been acquired in the GUESS condition, a choice pattern absent from the DRAW condition (replication: repeated-measures ANOVA, interaction between number of acquired samples from the current option and condition, $F_{9,234} = 11.10$, $p < 0.001$, $\eta^2_G = 0.048$). Second, we computed 'choice similarity' matrices – the average similarity of a choice to all later choices in the same game (Figure 2c). We confirmed that participants converged progressively on a preferred option in the DRAW condition. In the GUESS condition, however, the analysis of choice similarity matrices revealed that participants not only sampled repeatedly from the same option in the first trials (choice similarity among trials 1-4; discovery: median = 88.2%, IQR = [67.5% 91.9%]; replication: median = 85.6%, IQR = [51.6% 92.3%], signed-rank test against random sampling, $z = 3.8$, $p < 0.001$), but also sampled preferentially from the other option in subsequent trials (choice similarity between trials 1-4 and 8-11; discovery: median = 39.5%, IQR = [19.9% 45.9%]; replication: median = 46.1%, IQR = [28.8%

51.2%], signed-rank test against random sampling, $z = -2.4$, $p = 0.019$). Both conditions also showed an elevated similarity between consecutive choices.

In both conditions, the colour associated with one option could not be predicted by the colour associated with the other option (Figure 1b). In the DRAW condition, this aspect of the task design means that different games required different amounts of switches between the two

a task conditions

b option-color associations

c color mixture distributions

d sampling choices of DRAW games

e last sampling choice of DRAW games

f final question of GUESS games

Fig. 1 | Task design and performance. **a**, Task conditions. Left: game instructions. In DRAW games, participants were asked to draw outcomes from a target colour (e.g., blue). In GUESS games, the same participants were asked to identify option-colour associations. Centre: sequential sampling phase. Each game consisted of 8, 12, 16 or 20 choices, varying pseudo-randomly across games. The sequential sampling phase was identical across conditions. Right: final question. In GUESS games, participants were asked to report the colour associated with either of the two options. **b**, Colour-option associations. In both conditions, participants were informed that there was no relation between individual option-colour associations, such that the two options could be associated with the same colour or with different colours. **c**, Distribution of colour mixtures. Outcome colours were drawn from a continuum ranging from blue to orange. Blue bags were associated to 66% of blue outcomes and 33% of orange outcomes, and vice versa for orange bags. **d**, Sampling choices of DRAW games (left: discovery dataset, right: replication dataset). The fraction of choices towards the actual target option is plotted against choice position for single-target games (group means $\pm$ 95% CI). The thick horizontal line indicates a cluster of significant difference from chance in the replication dataset. **e**, Last sampling choice of DRAW games. Fraction of target-option choices on the last choice of single-target games (medians $\pm$ inter-quartile ranges). **f**, Final question of GUESS games. Fraction of correct responses to the final question (medians $\pm$ inter-quartile ranges).

options. Indeed, games where both options were associated with the target colour required few switches, whereas games where both options were associated with the non-target colour required frequent switches between the two options. We thus compared choice patterns in these ‘target-absent’ games of the DRAW condition to the choice patterns observed in the GUESS condition. Choice imbalance grew at similar rates in target-absent games of the DRAW condition (linear coefficient; discovery: mean = 0.26, CI = [0.22 0.29]; replication: mean = 0.24, CI = [0.20 0.28]) and in the GUESS condition (replication: paired t -test between conditions, $t_{26} = 0.4$, $p = 0.716$). However, choice imbalance deviated only weakly from a linear growth in the first trials of target-absent games of the DRAW condition (deviation from linear fit; discovery: peak at trial 5, mean = +0.53, CI = [+0.31 +0.76]; replication: mean = +0.54, CI = [+0.30 +0.77]) – much less than the deviation observed in the GUESS condition (Figure 2a; replication: paired t -test between conditions at trial 5, $t_{26} = -3.2$, $p = 0.003$).

Modelling differences in human exploration patterns between conditions

To characterize the policies used by participants to choose which option to sample in the two conditions, we computed – for each choice of each game – the accumulated information regarding option-outcome associations based on all previous choices made in the game. We could then label the direction of the choice (Figure 2d) as target-congruent (in the DRAW condition) and as uncertainty-congruent based on the accumulated information (in both conditions). We defined as target-congruent choices in direction of the option most likely to be associated with the target colour (the option with higher accumulated information for the target colour), and as uncertainty-congruent choices in direction of the option whose associated colour is most uncertain (the option with lower absolute accumulated information). As expected, a majority of participants’ choices were target-congruent in the DRAW condition (t -test against 50%, replication: peak $t_{26} = 9.2$, cluster-level $p < 0.001$), and more choices were uncertainty-congruent in the GUESS condition (paired t -test between conditions, replication: peak $t_{26} = 3.2$; cluster-level $p < 0.001$). But this was not true in the first trials of each GUESS game, where participants’ choices were strongly biased toward the more certain option (replication: peak $t_{26} = 3.7$; cluster-level $p < 0.01$).

The specific choice pattern observed in the GUESS condition – where outcomes have no rewarding nor prospective value – differs sharply from the behaviour of a rational agent which chooses the option whose outcome is most uncertain on each trial of each game (see Methods): unlike participants, this theoretical agent does not sample repeatedly from one option, and then from the other option in the first trials of each game (Supplementary Figure 1).

Based on the specific features of human exploration patterns, we designed a process model of sampling choices. This model was derived from the rational agent which makes target-congruent choices in the DRAW condition, and uncertainty-congruent choices in the GUESS condition (Supplementary Figure 1). Like the rational agent, the choices of the model are sensitive to the difference between the information already acquired from the two options (sensitivity parameter β_{inf}), and attracted toward the option most likely to be associated with the target colour in the DRAW condition (sensitivity parameter β_{tar}). In addition, unlike the rational agent, the model could include an initial sampling phase during which it acquires, in turn, a number of samples from each option at the beginning of each game. This initial sampling strategy aims at forming initial hypotheses regarding individual option-outcome associations in the first trials of each game. It is controlled by a number n_{ini} of initial samples from each option and a threshold strength θ_{ini} (see Methods). The model could also include a tendency to sample repeatedly from the same option

after the initial sampling phase. This tendency is controlled by a repetition bias parameter b_{rep} . We fitted this model to human sampling patterns in the DRAW and GUESS conditions.

Best-fitting parameter estimates obtained through maximum likelihood estimation (see Methods) confirmed behavioural observations (Figure 3). First, participants' choices were attracted toward the option most likely to be associated with the target colour in the DRAW condition (Figure 3a; sensitivity parameter β_{tar} , discovery: mean = 1.97, CI = [1.58 2.37], replication: mean =

Fig. 2 | Differences in exploration patterns. **a**, Sampling imbalance. Difference between the number of choices for the most and the least chosen option at each choice throughout the games (group means $\pm$ 95% CI). Dotted lines correspond to linear fits of the data, averaged over participants; thick horizontal lines indicate clusters of significant differences between participants' sampling imbalance and the linear fit of the same data. **b**, Option re-sampling. Fraction of re-sampling the same option as a function of the number of times the option was previously sampled (group means $\pm$ 95% within-subject CI). Thick horizontal lines illustrate clusters of significant differences between conditions. **c**, Choice-choice similarity matrices. Average similarity of each choice to the other choices in the same sequence. **d**, Choice direction. Left: fraction of target-congruent choices, defined as choices towards the currently most target-like option (group means $\pm$ 95% CI). The thick horizontal line indicates the cluster of significant difference from chance. Right: fraction of uncertainty-congruent choices, defined as choices towards the currently most uncertain option (group means $\pm$ 95% CI). The thick horizontal line indicates the cluster of significant difference between conditions.

1.86, CI = [1.34 2.38], t -test against zero, $t_{26} = 7.0$, $p < 0.001$, $d = 1.35$). Second, participants' choices were sensitive to the difference in information acquired from the two options, but in opposite directions between the two conditions (Figure 3b). The associated sensitivity parameter β_{inf} was positive in the DRAW condition (discovery: mean = 0.14, CI = [0.10 0.19]; replication: mean = 0.13, CI = [0.07 0.19], $t_{26} = 4.2$, $p < 0.001$, $d = 0.81$), but negative in the GUESS condition (discovery: mean = -0.45, CI = [-0.57 -0.33]; replication: mean = -0.29, CI = [-0.49 -0.08], $t_{26} = -2.8$, $p = 0.01$, $d = -0.53$). This means that participants' choices were attracted toward the option for which the associated colour is most uncertain only in the GUESS condition – where participants do not aim at maximising the immediate reward of each choice (difference in β_{inf} between conditions; discovery: mean = 0.59, CI = [0.47 0.71]; replication: mean = 0.42, CI = [0.23 0.61], $t_{26} = 4.3$, $p < 0.001$, $d = 0.83$). Unlike the sensitivity to information, participants' choices showed similar (and substantial) repetition biases b_{rep} between the DRAW condition (Figure 3c; discovery: mean = 0.77, CI = [0.56 0.97]; replication: mean = 0.78, CI = [0.52 1.04], $t_{26} = 5.9$, $p < 0.001$, $d = 1.13$) and the GUESS condition (discovery: mean = 1.02, CI = [0.81 1.22]; replication: mean = 0.95, CI = [0.67 1.22], $t_{26} = 6.8$, $p < 0.001$, $d = 1.30$).

Finally, we compared variants of the model with vs. without the initial sampling phase using Bayesian model selection (Figure 3d; see Methods). In the DRAW condition, the model with the initial sampling phase was much less prevalent than the model without the initial sampling phase (discovery: mean = 16.4%, s.d. = 6.8%; replication: mean = 14.6%, s.d. = 6.4%, exceedance $p < 0.001$). By contrast, the inclusion of an initial sampling phase improved model fits for a majority of participants in the GUESS condition (discovery: mean = 93.6%, s.d. = 4.5%; replication: mean = 67.9%, s.d. = 8.5%, exceedance $p = 0.976$). Accordingly, the number n_{ini} of initial samples from each option was larger in the GUESS condition (discovery: median = 3, IQR = [2 5]; replication: median = 3, IQR = [0 4]) than in the DRAW condition where $n_{\text{ini}} = 0$ for more than 75% of tested participants (signed-rank test, $z = 3.4$, $p < 0.001$, $r = 0.65$). Furthermore, when considering participants for which an initial sampling phase improved model fits, the estimated threshold strength θ_{ini} was close to a 'hard' threshold (discovery: median = 0.854, IQR = [0.698 0.876]; replication: median = 0.869, IQR = [0.736 0.879]).

To validate the notion that the initial sampling phase was required to reproduce the specific features of human exploration patterns in the GUESS condition, we performed a 'knock-in' simulation procedure in which we tested the effects of each aspect of the model on behavioural measures in the two conditions (discovery: Supplementary Figures 2-3; replication: Supplementary Figures 4-5; see Methods). The inclusion of an initial sampling phase was necessary to reproduce – both qualitatively and quantitatively – the sampling patterns observed in the GUESS condition. A repetition bias was necessary to explain participants' overall tendency to sample repeatedly from the same choice option in both conditions.

Dissociating the effects of delayed exploitation and outcome valuation on exploration

The GUESS condition (in which participants sample freely the two options to learn option-outcome associations) differs from the DRAW condition (in which the same participants seek to sample a rewarded outcome) in two important ways. First, participants are rewarded based on a final two-alternative forced choice at the end of each game in the GUESS condition, whereas they are rewarded based on the outcomes they have sampled on each choice of each game in the DRAW condition. This means that exploitation is traded against exploration on each choice in one condition, whereas exploitation is delayed until the final 'money time' choice in the other condition. Second, outcomes have no rewarding value in the GUESS condition, whereas they are associated

with positive or negative value depending on the target colour in the DRAW condition. It is therefore impossible to know the extent to which each of these two factors (delayed exploitation and outcome valuation) contributes to the specific exploration patterns observed in the GUESS condition.

Fig. 3 | Model fits. **a**, Sensitivity to the target (β_{tar}). Participant estimates of choice sensitivity to the target colour in DRAW games (group means $\pm$ 95% CI). **b**, Sensitivity to information (β_{inf}). Participant estimates of choice sensitivity to the uncertainty of option-outcome associations (group means $\pm$ 95% CI). **c**, Repetition bias (b_{rep}). Model estimates of participants' repetition bias in each condition (group means $\pm$ 95% CI). **d**, Model comparison. Prevalence of the model with an initial sampling phase in each condition (mean $\pm$ s.d.). Exceedance probability of the model with vs. without an initial sampling phase is given for the validation dataset in each condition, as well as the probability that the model with an initial sampling phase is more prevalent in the GUESS condition. **e**, Number of initial samples from each novel option (n_{ini}) (individual estimates, group medians and inter-quartile ranges). **f**, Threshold strength (θ_{ini}) (individual estimates, group medians $\pm$ inter-quartile ranges).

To address this issue, we designed a third condition (FIND) where participants were asked to find an option associated with a target colour – provided at the beginning of each game (Figure 4a). In this new condition, as in the DRAW condition, each outcome is associated with a positive or negative value depending on the target colour. However, as in the GUESS condition, participants were rewarded not based on the outcomes they have sampled, but based on their answers to a final two-alternative forced choice at the end of each game where they have to select an option associated with the target colour. This new condition was akin to browsing restaurant reviews: choices yielded no immediate reward, but outcomes were still valued with respect to an

end goal. Consequently, the contrast between the FIND and DRAW conditions reflects the difference between delayed and immediate exploitation, whereas the contrast between the FIND and GUESS reflects the difference between the sampling of outcomes with and without rewarding value.

a task conditions

b sampling choices

last sampling choice

c final question

Fig. 4 | Task design and performance in the third dataset. a, Task conditions. An intermediate FIND condition was introduced in this dataset. Participants were instructed that they were rewarded on accuracy at the final question, but they were given the target colour in advance, from the beginning of the sequence: “you will have to identify a blue/orange bag”. In other words, rewards were delayed, as in GUESS games, but outcomes were valued with respect to a target colour as in DRAW games. **b**, Sampling choices. Left: fraction of target-option choices across DRAW and FIND games with a single-target option (group mean $\pm$ 95% CI and clusters of significant difference against chance). Participants sampled preferentially from the target option more than would be expected by chance in the FIND condition. Right: fraction of target-option choices at the last sampling choice of the sequence, for single-target games (medians $\pm$ inter-quartile ranges). **c**, Final question. Fraction of correct responses to the final question in single-target FIND and GUESS games (group medians $\pm$ inter-quartile ranges).

As in previous datasets, participants ($N = 31$ after exclusion; see Methods) selected the option associated with the target colour more often than the other option on the last choice of each game with a single target option in the DRAW condition (Figure 4b; median = 81.2%, IQR = [62.5% 93.8%], signed-rank test against chance, $z = 4.8$, $p < 0.001$, $r = 0.86$). This preference for the target option was also visible, albeit to a lesser extent, in the FIND condition (median = 62.5%, IQR = [50.0% 75.0%], $z = 4.0$, $p < 0.001$, $r = 0.71$). Likewise, as in previous datasets, participants responded accurately to the final two-alternative forced choice in the GUESS condition (Figure 4c; median = 87.5%, IQR = [81.2% 93.8%], $z = 4.9$, $p < 0.001$, $r = 0.87$). Participants also responded very accurately to the final choice in the FIND condition (median = 93.8%, IQR = [87.5% 100%], $z = 4.9$, $p < 0.001$, $r = 0.88$). In this condition, the fact that participants selected the option associated

with the target color much more often on the final choice (93.8%) than on the last sampling choice of the game (62.5%) shows that participants effectively delayed exploitation until the final choice.

Participants' sampling imbalance showed the same characteristics as in previous datasets (Figure 5a). Sampling imbalance grew linearly over time in the DRAW condition (linear coefficient: mean = 0.40, CI = [0.33 0.46]). This growth was again slower and non-linear in the GUESS condition (mean = 0.19, CI = [0.14 0.23]; paired t -test between conditions: $t_{30} = -7.7$, $p < 0.001$, $d = -1.40$). As before, sampling imbalance grew rapidly in the first trials of each game (deviation from linear fit; peak at trial 5, mean = +1.24, CI = [+0.89 +1.59], t -test against zero, $t_{30} = 7.2$, $p < 0.001$, $d = 1.61$). Sampling imbalance also grew non-linearly in the FIND condition (deviation from linear fit; peak at trial 5, mean = +0.79, CI = [+0.48 +1.10], $t_{30} = 5.2$, $p < 0.001$, $d = 0.94$), but less so than in the GUESS condition (paired t -test between conditions: $t_{30} = 4.0$, $p < 0.001$, $d = 0.72$).

The fraction of re-sampling from the same option as a function of the number of previous samples acquired from this option differed significantly between conditions (Figure 5b; interaction between number of acquired samples from the current option and condition, $F_{18,540} = 6.2$, $p < 0.001$, $\eta^2_G = 0.027$). As for the GUESS condition, participants tended to re-sample more from options for which few samples have been acquired in the FIND condition than in the DRAW

Fig. 5 | Differences in exploration patterns in the third dataset. **a**, Choice direction. Left: fraction of target-congruent choices, defined as choices towards the currently most target-like option (group means $\pm$ 95% CI). Thick horizontal lines indicate clusters of significant difference from chance (plain colour) or between conditions (hatched). Right: fraction of uncertainty-congruent choices, defined as choices towards the currently most uncertain option (group means $\pm$ 95% CI). Thick horizontal lines indicate clusters of significant difference between conditions. **b**, Sampling imbalance. Difference between the number of choices for the most and the least chosen option at each choice throughout the games (group means $\pm$ 95% CI). Dotted lines correspond to linear fits of the data, averaged over participants; thick horizontal lines indicate clusters of significant differences between participants' sampling imbalance and the linear fit of the same data. **c**, Option re-sampling. Fraction of re-sampling the same option as a function of the number of times the option was previously sampled (group means $\pm$ 95% within-subject CI). Thick horizontal lines illustrate clusters of significant differences between conditions. **d**, Choice-choice similarity matrices. Mean similarity of each choice to the other choices in the same sequence (group means).

condition, whereas the converse was true for options for which several samples have been acquired. Choice similarity matrices (Figure 5c) confirmed the specific sampling patterns observed in the GUESS condition: participants sampled repeatedly from the same option in the first trials (choice similarity among trials 1-4; median = 67.7%, IQR = [56.6% 88.3%]; signed-rank test against random sampling, $z = 4.2$, $p < 0.001$, $r = 0.76$), and then sampled preferentially from the other option in subsequent trials (choice similarity between trials 1-4 and 8-11; median = 47.4%, IQR = [35.7% 50.6%], $z = -2.7$, $p = 0.006$, $r = -0.49$). This second effect was absent in the FIND condition (median = 49.5%, IQR = [46.6% 53.9%], $z = -0.0$, $p = 0.977$, $r = -0.01$; signed-rank test between conditions: $z = 3.7$, $p < 0.001$, $r = 0.66$).

Labelling individual choices as target-congruent and as uncertainty-congruent, based on the accumulated information, replicated known effects (Figure 5d). Participants made a majority of target-congruent choices in the DRAW condition (t -test against 50%: peak $t_{30} = 10.5$, cluster-level $p < 0.001$). Uncertainty-congruent choices were more frequent in the GUESS condition (paired t -test between conditions: peak $t_{30} = 2.5$, cluster-level $p < 0.01$), after the first trials of each game during which choices were biased toward the more certain option (peak $t_{30} = 3.1$, cluster-level $p < 0.01$). Participants also made more target-congruent than target-incongruent choices in the FIND condition (peak $t_{30} = 6.3$, cluster-level $p < 0.001$), but less than in the DRAW condition (paired t -test between conditions: peak $t_{30} = 5.4$, cluster-level $p < 0.001$). Furthermore, choices made in the FIND condition were more uncertainty-congruent than in the DRAW condition (peak $t_{30} = 2.4$, cluster-level $p < 0.05$) – and uncertainty congruence showed similar time courses in the FIND and GUESS conditions. In the FIND condition, participants were thus making uncertainty-congruent choices as in the GUESS condition, but their choices were also biased towards the more target-congruent option as in the DRAW condition.

Model fits provided additional insights. First, the sensitivity of participants' choices to the target outcome β_{tar} was lower in the FIND condition (Figure 6a; mean = 1.29, CI = [0.78 1.79], t -test against zero: $t_{30} = 4.9$, $p < 0.001$, $d = 0.89$) than in the DRAW condition (mean = 1.95, CI = [1.49 2.41], t -test against zero: $t_{30} = 8.4$, $p < 0.001$, $d = 1.50$; paired t -test between conditions: $t_{30} = 3.0$, $p = 0.005$, $d = 0.55$). Participants' choices were also directed toward the more uncertain option in the FIND condition (Figure 6b; sensitivity parameter β_{inf} , mean = -0.11, CI = [-0.21 -0.02], t -test against zero: $t_{30} = -2.4$, $p = 0.024$, $d = -0.43$), as in the GUESS condition (mean = -0.15, CI = [-0.27 -0.02], $t_{30} = -2.4$, $p = 0.021$, $d = -0.44$) but unlike the DRAW condition (mean = 0.12, CI = [0.07 0.17], $t_{30} = 4.7$, $p < 0.001$, $d = 0.84$). And as in previous datasets, participants' choices showed similar repetition biases b_{rep} across all three conditions (Figure 6c; t -test against zero: $p < 0.001$, $d > 0.5$ in all conditions).

The prevalence of the initial sampling phase differed across conditions (Figure 6d). The initial sampling phase had low prevalence in the DRAW condition (mean = 29.4%, s.d. = 7.8%, exceedance $p < 0.001$), but high prevalence in the GUESS condition (mean = 74.4%, s.d. = 7.5%, exceedance $p > 0.999$). In the FIND condition, the initial sampling phase improved the fits for half of tested participants (mean = 51.9%, s.d. = 8.6%, exceedance $p = 0.589$) – more than in the DRAW condition (exceedance $p = 0.971$), but less than in the GUESS condition (exceedance $p = 0.027$). As in previous datasets, the number n_{ini} of initial samples from each option varied across participants (Figure 6e), and the estimated threshold strength θ_{ini} was close to a 'hard' threshold when considering participants showing an initial sampling phase (Figure 6f). The inclusion of an initial sampling phase was necessary to reproduce the sampling patterns observed in the GUESS and FIND conditions. By contrast, as in previous datasets, a repetition bias was necessary to explain participants' overall tendency to sample repeatedly from the same choice option in all three conditions (Supplementary Figures 6-7). In summary, delaying exploitation results in less target-

congruent choices and more uncertainty-congruent choices, whereas stripping outcomes from their rewarding values triggers initial choice streaks from each novel option.

Fig. 6 | Model fits in the third dataset. **a**, Model comparison. Prevalence of the model with an initial sampling phase in each condition (mean $\pm$ s.d.). Exceedance probability of the model with vs. without an initial sampling phase is given in each condition, as well as the probability that the model with an initial sampling phase is more prevalent in the GUESS condition vs. each other condition. **b**, Sensitivity to the target (β_{tar}). Participant estimates of choice sensitivity to the target colour in DRAW and FIND games (group means $\pm$ 95% CI). **c**, Sensitivity to information (β_{inf}). Participant estimates of choice sensitivity to the uncertainty of option-outcome associations (group means $\pm$ 95% CI). **d**, Repetition bias (b_{rep}). Model estimates of participants' repetition bias in each condition (group means $\pm$ 95% CI). **e**, Number of initial samples from each novel option (n_{ini}) (individual estimates, group medians $\pm$ inter-quartile ranges). **f**, Threshold strength (θ_{ini}) (individual estimates, group medians $\pm$ inter-quartile ranges).

Positive covariations between exploration patterns across individuals

Human exploration patterns in the GUESS condition show two specific features absent from the DRAW condition: a tendency to choose the option whose associated outcome is currently more uncertain (feature 1), which follows an initial sampling phase where sampling proceeds in streaks of choices from each option (feature 2). In contrast to the first feature, the second feature is statistically suboptimal: indeed, in the GUESS condition, a rational agent would always choose the more uncertain option to maximise response accuracy during the final choice (Supplementary Figure 1). But is this initial sampling phase detrimental to accuracy, and expressed in individuals who do not make directed choices toward the more uncertain option? Or do these features of exploration patterns reflect two forms of information seeking that are not traded against each other, but rather co-expressed in the same individuals?

To address this issue, we studied covariations between exploration patterns across tested participants, by pooling the three datasets together (total $N = 85$). We split participants in two groups: a first group with participants for whom an initial sampling phase improved model fits in the GUESS condition ($N = 62$), and a second group with participants for whom an initial sampling phase did not improve model fits in the same condition ($N = 23$). Unlike other participants, and as expected, participants with an initial sampling phase showed a clear non-linear increase in sampling imbalance in the GUESS condition (Figure 7a), a strong tendency to sample repeatedly

Fig. 7 | Individual differences in exploration patterns. Participants were split in two groups based on whether a model with (ini.) or without (no ini.) an initial sampling phase better accounted for their exploration patterns in GUESS games. **a**, Sampling imbalance. Difference between the number of choices for the most and the least chosen option at each choice throughout the games (group mean $\pm$ 95% CIs). **b**, Option re-sampling. Fraction of re-sampling the same option as a function of the number of times the option was previously sampled (group means $\pm$ 95% within-subject CI). **c**, Choice direction. Fraction of target-congruent choices (left, DRAW games) or uncertainty-congruent choices (right, GUESS games; group mean $\pm$ 95% CI). **d**, Performance. Left: fraction of target-congruent choices on the last sampling choice, for single-target DRAW games (medians and inter-quartile ranges). Right: Fraction of correct responses at the final question, for single-target games (medians $\pm$ inter-quartile ranges). **e**, Model parameters. Individual estimates, group means $\pm$ 95% CI. Left: sensitivity to target (β_{tar}). Middle: sensitivity to information (β_{inf}). Right: repetition bias (b_{rep}). The scatter plot on the right depicts the correlation of repetition bias estimates between DRAW and GUESS games.

from options for which few samples have been acquired (Figure 7b), and an initial preference for the more certain option in the first trials of each GUESS game (Figure 7c). But these participants also appeared to converge more rapidly on a preferred option in the DRAW condition (Figure 7a), to repeat more often their previous choice (Figure 7b), and to select more often the more uncertain option than other participants in the second half of GUESS games (Figure 7c). Regarding performance, we found that participants with an initial sampling phase not only made more accurate responses to the final two-alternative forced choice in the GUESS condition (Figure 7d; rank-sum test between groups of participants: $z = 3.7$, $p < 0.001$, $r = 0.40$), but also chose the target option more often than other participants on the last choice of each DRAW game ($z = 2.7$, $p < 0.01$, $r = 0.29$).

The analysis of model fits consolidated these observations (Figure 7e). Participants with an initial sampling phase (in the GUESS condition) showed higher sensitivity to the target outcome β_{tar} than other participants in the DRAW condition (two-sample t-test between groups of participants: $t_{83} = 4.13$, $p < 0.001$, $d = 1.01$). The choices of these participants were also more sensitive to the difference in information acquired from the two options in the GUESS condition (sensitivity parameter β_{inf} , mixed-effects ANOVA, interaction between group and condition: $F_{1,83} = 29.2$, $p < 0.001$, $\eta^2_{\text{G}} = 0.122$). This means that participants with an initial sampling phase were the ones whose choices were most attracted toward the more uncertain option after the end of the initial sampling phase. Finally, participants with an initial sampling phase showed higher repetition bias b_{rep} than other participants, in both conditions (main effect of group: $F_{1,83} = 13.8$, $p < 0.001$, $\eta^2_{\text{G}} = 0.128$; interaction: $F_{1,83} = 3.08$, $p = 0.083$, $\eta^2_{\text{G}} = 0.004$). The strength of this repetition bias was highly correlated between conditions across all participants ($N = 85$, linear correlation: $r^2 = 0.62$, $p < 0.001$). Together, these results show strong positive covariations between exploration patterns across tested participants and conditions.

Discussion

Exploration is central to efficient decision-making: discovering and monitoring the current structure of one's environment allow agents to choose goal-relevant actions and to adapt to changes in its state. In the laboratory, however, exploration is typically studied in terms of its trade-off with reward maximisation using tasks that involve explore-exploit dilemmas (e.g., multi-armed bandits). Here, in contrast to this large body of work, we sought to study human exploration patterns in situations where choices are not biased by reward maximisation – as is often the case outside the laboratory. For this purpose, we compared sequences of choices where exploration conflicts with the exploitation of a rewarding option, and sequences where choice options are not associated with reward – and exploration therefore does not conflict with reward maximisation. Across three datasets, we obtained replicable evidence of specific exploration patterns arising in situations where exploration is released from its trade-off with reward maximisation. These specific patterns outline competing cognitive pressures on information seeking – between the continued sampling of the current option to generate and test the hypothesis regarding its associated outcome (a 'local' form of information seeking), and the directed sampling of the most uncertain option available to map the overall task structure (a 'global' form of information seeking).

First, we found that participants made exploratory decisions directed at seeking information about uncertain choice options only when they were not incentivized to maximise the reward associated with each choice. Specifically, participants' choices were significantly directed toward the most uncertain option in the GUESS and FIND conditions where participants were rewarded based on a single two-alternative forced choice at the end of each game. This was not the case in the DRAW condition despite matched objective levels of uncertainty across all three conditions. This first finding is consistent with previous work showing that longer temporal 'horizons' allow participants to engage in more information seeking – often described as 'directed' exploration in contrast to 'random' exploration^{10,11}. This type of exploration corresponds to a global form of information seeking¹², which reduces the overall uncertainty about the task structure (i.e., option-outcome associations). Note that participants' choices also displayed significant randomness in all conditions – which we captured using a 'softmax' selection policy as in earlier work^{2,3,6}. Recent work has shown that some of this randomness can be triggered by random variance ('noise') in the learning of option values and not only by stochasticity in the selection policy¹³. Future work should therefore dissociate the relative contributions of learning noise and choice stochasticity to the overall variability of participants' choices.

We observed that participants' exploration patterns diverged significantly from directed exploration (i.e., choosing the most uncertain option at each choice) in ways that cannot be explained by random exploration (or learning noise) alone. Indeed, participants repeated their previous choice much more than was predicted by normative accounts. This 'repetition' bias was structured temporally in two stages over the course of each game (i.e., each new environment to sample): 1. when sampling from an option for the first time, participants were repeatedly sampling from the same option before switching to the other, and 2. when returning to a previously sampled option, participants sustained a significant bias toward sampling repeatedly from the same option. The immediate consequence of this repetition bias is that participants sampled the two options in temporal 'streaks' that were longer than the ones predicted by normative accounts in the same conditions – especially at the beginning of each new delayed reward game.

In sequential sampling tasks (including bandits), sustained repetition biases are typically attributed to a form of 'confirmation' bias – the tendency to validate a previous choice, either

because of an asymmetric learning process favouring choice-supportive evidence^{14,15,16} or a biased choice process that assigns value to previous choices^{17,18}. In either case, confirmation biases are described as affecting option/action values. This account, however, cannot explain alone our findings. Indeed, participants are incentivised to maximise immediate rewards only in the DRAW condition, whereas the sustained repetition bias is equally strong across all three conditions – and even reliably correlated across conditions. In the GUESS condition, options and actions have no immediate or delayed rewarding value, thus repeating a previous choice does not – in itself – validate it in any way. Participants may instead be biased toward repeating their previous choice to test their current hypothesis about the colour associated with the currently sampled option – a local (option-specific) form of information seeking. But it could also reflect a strategy for mitigating the large cognitive cost associated with updating beliefs about two (or more) choice options in parallel. By chunking samples into temporal streaks of each option, participants are effectively minimizing such ‘switch costs’ which are associated with mental effort^{19,20} and even subjective pain²¹. Furthermore, hypothesis testing strategies have been argued to be more cognitively efficient in self-directed learning contexts (as it is the case in our task), because such strategies provide evidence perceived as more relevant at each choice^{22,23}.

From this perspective, the initial streaks of samples from each choice option at the beginning of each game in the ‘delayed reward’ conditions (GUESS and FIND) could aim at generating hypotheses about each option, before testing them in subsequent choices. Participants behave as if they probed each option in turn until they were reasonably convinced of its associated colour before moving on to the other option. This choice pattern was entirely absent in the ‘immediate reward’ condition (DRAW) despite the tight match of task variables across all conditions. This selectivity explains why it was not described in earlier work using tasks involving explore-exploit dilemmas^{6,10,11}. This selectivity also means that, in the absence of any explore-exploit trade-off at each choice (i.e., no pressure to maximise immediate reward), participants prefer to seek information about a single choice option at a time. This exploration pattern corresponds to a local form of information seeking, which aims at forming and testing the hypothesis tied to a specific option (rather than the overall task structure). A more general conclusion of our study is that human exploration patterns are very different when participants seek to maximise immediate rewards compared to when they can freely acquire information that is not associated with reward. Our findings suggest that restricting the study of exploration to explore-exploit dilemmas may have taught us more about how humans arbitrate this trade-off – and its neural correlates, than about exploration itself²⁴. The two-stage exploration strategy we have identified is also reminiscent of information search algorithms used to describe the behaviour of flying moths and birds²⁵.

Here we modelled these initial streaks of choices as a ‘soft’ (probabilistic) threshold on the number of samples to draw from an option when it is sampled for the first time. In statistical terms, the number of samples obtained from a given option reflects the precision of the mean outcome (colour) associated with the option^{26,27}. The best-fitting strength of this threshold was very high (close to a ‘hard’ threshold) in all three datasets, which suggests that it reflects a strongly preferred strategy during the initial sampling of uncertain choice options. It is nevertheless possible that the threshold is not set directly on the number of samples drawn from an option, but on the subjective amount of evidence acquired for an option – which reflects the confidence in the hypothesis associated with the current option²⁸. Future research should further characterize the precise cognitive variable on which this threshold is set, either by implementing and comparing variants of our model, or by testing additional task conditions. As an example, making option-outcome associations more uncertain should lead to higher thresholds in terms of number of samples. Indeed, it should lead participants to take more samples in their initial sampling bouts, because each sample provides less evidence regarding the mean outcome associated with the sampled

option. Similarly, informing participants about the total number of choices at the beginning of each game, and varying this number across games, should prompt adaptations of the threshold: higher thresholds for games with more samples and lower thresholds for games with less samples.

A specific feature of our laboratory-based, controlled study of human exploration is the use of environments with only two choice options. This task feature facilitates the comparison of our findings to the existing literature using variants of the bandit task with small numbers of choice options^{6,9-11}. However, the ecological validity of such a choice problem is indisputably limited²⁹. Indeed, outside the laboratory, humans (and other animals) choose between much more than two options, leading to structured patterns of exploration across these large option spaces⁵. Our observation of a patterned sampling of choice options in novel environments (i.e., at the beginning of each game), however, should not depend critically on the number of available choice options. Indeed, our findings suggest that participants structure their sampling patterns in streaks as a way to generate and test hypotheses about each choice option in a cost-minimizing fashion. The costs which we believe drive participants to structure their sampling patterns in this way should only increase for choice problems with more than two options, and we therefore predict that these initial streaks of choices should also be visible in large option spaces – in the absence of any explore-exploit trade-off. Furthermore, a similar patterned sampling of choice options has previously been described as an efficient adaptation to large hypothesis spaces²³. The fact that we observe such patterned sampling in a two-option choice setting provides some indication that it reflects cognitive limitations (such as the cost of updating multiple hypotheses in parallel) more than an adaptive optimization of choice to multi-alternative environments. In any case, we expect participants to display similar streaks in their sampling patterns when confronted with environments with more choice options to explore.

We observed significant individual differences in our study: most but not all participants displayed a patterned sampling of choice options at the beginning of games which did not elicit an immediate explore-exploit trade-off – in the GUESS and FIND conditions. And even when this patterned sampling was present, its associated threshold varied in strength and number of samples. These individual differences afforded us to relate the patterned sampling of choice options to other aspects of participants' behaviour. Participants who displayed patterned sampling were more accurate at identifying the overall structure of each game (i.e., the option-outcome associations), and also more sensitive to task-relevant decision variables, even in games requiring to maximise immediate reward where participants did not display patterned sampling – in the DRAW condition. This positive relation between patterned sampling in the first trials of each game (a suboptimal strategy in normative terms) and task performance is important, because it shows that patterned sampling does not reflect task disengagement nor inattention. However, the interpretation of this positive relation remains two-fold: it could either mean that this patterned sampling reduces cognitive effort and switch costs^{19,20}, or that the participants who deploy this sampling strategy have different psychological traits than participants who do not deploy it. Future research should investigate whether and which psychological traits affect participants' tendency to display this sampling strategy. Differences in impulsivity, orientation toward future events or attitude towards uncertainty could potentially explain individual differences in the patterned sampling of choice options observed in this study³⁰⁻³².

Taken together, our findings delineate specific features of human exploration patterns in the absence of trade-off with immediate reward maximisation. These specific exploration patterns suggest competing cognitive pressures on information seeking that are neither described nor present in conditions eliciting explore-exploit dilemmas: 1. a 'local' form of information seeking that aims at generating and testing a hypothesis about the currently sampled option, and 2. a 'global'

form of information seeking that aims at reducing the overall uncertainty about the task structure. These results call for further research into the neural mechanisms of these exploration patterns outside of any explore-exploit dilemma, and into their ecological/psychiatric validity. These results also suggest that the exploration patterns described in explore-exploit dilemmas may reflect the arbitration of this trade-off more than exploration itself. Outside the laboratory, exploration is often not traded against immediate reward maximisation, and more work should investigate the specific drivers of exploration in these ecological contexts.

Methods

Participants

For the first (discovery) dataset, we recruited $N = 30$ adult participants (15 females, mean age = 25 years, s.d. = 4.3 years). Because we had no specific expectation regarding effect sizes, sample size was determined to match or exceed earlier studies of human exploration patterns in the literature^{6,10}. Participants were recruited through a local, open mailing list advertising psychology and neuroscience experiments in the Paris area. Participants were paid 22 euros for 90 min of testing. Recruited participants were aged between 18 and 35 years, right-handed, with normal or correct-to-normal vision, and had no history of psychiatric or neurological disorder. For the second (replication) dataset, we also recruited $N = 30$ adult participants (15 females, mean age = 25 years, s.d. = 3.7 years). Participants were recruited through the local mailing list used for the first dataset, and through a national website (*L'Etudiant*) advertising short-term jobs for adult students. The other recruitment criteria were identical to those used for the first dataset. For the third (three-condition) dataset, we balanced the lower number of games per condition by aiming to reach $N = 30$ adult participants after exclusion. We recruited a total of $N = 38$ participants on the website *l'Etudiant* (21 females, mean age = 24 years, s.d. = 4.4 years). The other recruitment criteria were identical to those used for the first two datasets. Participants gave written informed consent before taking part in the study. Ethics approval was obtained from the relevant authorities (Comité de Protection des Personnes Ile-de-France VI, ID RCB: 2007-A01125-48, 2017-A01778-45).

We excluded participants whose sampling strategy was stereotyped or did not depend on presented outcomes. We computed the number of stereotyped games for each participant: games where the participant systematically alternated between left and right responses, or between the two options, and games where the participant sampled a single option throughout the whole game. We excluded participants who produced more than 50% of such stereotyped games in any condition ($N = 3$ participants in the discovery dataset, $N = 3$ participants in the replication dataset, and $N = 6$ participants in the third dataset). In the third experiment, one additional participant was sent home before taking part in the study due to COVID-19 symptoms.

Experimental design

Participants made sequences of choices (games) between two options depicted by shapes, and each choice yielded an outcome. Outcomes were colour mixtures, drawn from a continuum between blue and orange (through grey in the middle of the range). The outcomes of each option were drawn from one of two probability distributions: an 'orange' distribution with 67% of dominantly orange outcomes (and 33% of dominantly blue outcomes), and a 'blue' distribution with 67% of dominantly blue outcomes (and 33% of dominantly orange outcomes). The most frequent colour associated with an option was independent of the most frequent colour associated with the other option, such that participants needed to sample each individual option to learn its associated colour (Figure 1b). Each game consisted of a sequence of choices whose number varied pseudo-randomly and uniformly across games from 8, 12, 16, or 20 choices.

Each game was preceded by condition-specific instructions. In the DRAW condition, participants were asked to draw a maximum of outcomes from a target colour (orange or blue, counterbalanced across games). Participants were informed that the colour intensity of each outcome was translated into points: the more the outcome was of the target colour, the higher the number of points. This task feature rendered each outcome intrinsically rewarding. In this first

condition, each game ended at the end of the sequence of choices: No final question was asked and participants immediately moved on to the instructions screen for the next game. In the GUESS condition, participants were asked to guess the dominant colour associated with each of the two options. Each game ended with a final question, where participants were asked about the dominant colour associated with one of the two options. In this second condition, participants were rewarded solely based on the accuracy of their responses to this final question. Each outcome was not intrinsically rewarding in this condition, but it was informative regarding the dominant colour associated with the sampled option. In the FIND condition, participants were asked to find an option associated with a target colour (orange or blue, counterbalanced across games). As in the GUESS condition, participants were rewarded solely based on the accuracy of their responses to a final question, which in this case asked them to select an option associated with the target colour at the end of the game. To further enhance the symmetry between the FIND and GUESS conditions in the third experiment, we asked the same final question at the end of GUESS games. The single difference between the two conditions is that the question (i.e., the target colour) was known in advance in the FIND condition, but unknown in the GUESS condition.

Eight geometric shapes were used to instantiate options, arranged into twenty-eight possible pairs. Shape-colour associations changed pseudo-randomly between games. The colour intensities of outcomes were pre-defined in advance. Colour intensity indexes the amount of information provided by the outcome about its associated probability distribution: a more intense colour is more informative than a less intense one. This advance pre-definition of colour intensities allowed us to make sure that, for each choice, the amount of information that could be acquired from the two options was exactly the same. Because the colour associated with an option could not be identified from the colour associated with the other option, different games used options associated either with the same or with different colours. This means that in the DRAW condition, some games had one single target option, whereas other games had either zero (or two) target option(s). We included only games with a single target option when measuring performance in the DRAW condition. Each unique game was duplicated to be used once in each condition, so that each condition was tightly matched with the other(s) – except for the instructions provided to the participants. In particular, the uncertainty associated with each condition was therefore matched across conditions. Conditions were interleaved pseudo-randomly using blocks of games of the same condition, with a self-paced break between each block.

Experimental procedures

Participants were instructed to perform the task using a step-by-step, self-paced tutorial, followed by a practice session of eight games. Participants were then asked to describe the instructions to the experimenter, and it was emphasised again for all participants that: 1. the colour intensity of each outcome was relevant, 2. each option would not always draw the same colour, so the task was to identify its most frequent colour, 3. the two options were not always associated with different colours (they could be both associated with orange, or blue), 4. each game would end after a variable, unpredictable number of choices, and 5. performance is measured by the colours of all sampled outcomes in the DRAW condition, and by the final question in the GUESS and FIND conditions.

For the discovery dataset, the experiment (excluding practice games) consisted of 48 games of each condition (96 in total, divided in 8 blocks of 12 games). For the replication dataset, the experiment consisted of 64 games of each condition (128 in total, divided in 8 blocks of 16 games). For the third dataset (with the additional FIND condition), the experiment consisted of 32 games of each condition (96 in total, divided in 4 blocks of 24 games).

The task was run using the Psychtoolbox-3 toolbox³³ (version 3.0.14) for MATLAB 2017b (The Mathworks). The testing conditions were slightly different across datasets. For the first (discovery) dataset, eye position was recorded during the task using an EyeLink-1000 Plus eye-tracker system (SR Research), and participants' head movements were constrained by a chin rest. For the second (replication) and third datasets, eye position was not recorded and no chin rest was used. Participants were seated with their eyes at approximately 70 cm from a 24-inch LCD screen with a resolution of $1,920 \times 1,080$ pixels and a refresh rate of 60 Hz.

The trial structure is depicted in Fig.1a. After reading the instructions and initiating a game by pressing the spacebar, each sampling decision began with the two shapes appearing on the left and right sides of a fixation circle - sides were pseudo-randomly assigned. This choice screen was displayed until participants selected an option, by pressing the corresponding [A] (left shape) or [P] (right shape) key on a French keyboard, after which the outcome was displayed for 500 ms. The outcome corresponds to the filling of the selected shape with the sampled colour mixture. A fixation circle was then displayed alone for 500 ms before the trial began. In the DRAW condition, each game ended as soon as all the sampling choices were made (8, 12, 16, or 20) and their outcomes seen. The instructions for the following game then appeared on screen. In the GUESS (and FIND) conditions, each game ended with a final question screen. In the GUESS condition, one shape (pre-selected at random) was displayed and participants were asked to report its associated (dominant) colour. Participants selected a colour with a key press, and a selection box was displayed around the selected colour for 400 ms. Feedback on the accuracy of this final response was provided by filling the shape with its associated colour during 500 ms. The instructions for the following game then appeared on screen.

Computational model

The computational model of sampling behaviour featured the same learning module across all three conditions. The model tracks the mean colour mixture x associated with each option by accumulating the observed coloured outcomes acquired from each option (ranging from -1 for pure blue to +1 for pure orange). Based on these mean colour mixture estimates $\hat{x}$, the decision variable DV_t is computed at trial t of each game as the sum of three components. A first component, defined in the DRAW and FIND conditions, corresponds to the difference between the mean colour mixture estimates for each option $\hat{x}_{t,1}$ and $\hat{x}_{t,2}$, signed by the target colour c_{tar} . This first component is scaled by a sensitivity parameter β_{tar} , which controls the sensitivity of model choices to the target. A second component, defined in all conditions, corresponds to the difference between the information already acquired from each option, $|\hat{x}_{t,1}|$ and $|\hat{x}_{t,2}|$. This second component is scaled by a second sensitivity parameter β_{inf} , which controls the sensitivity of model choices to the uncertainty about individual option-outcome associations. The third component, also defined in all conditions for all choices except the first choice of each game, corresponds to an additive bias in direction of the previous choice, controlled by a repetition bias parameter b_{rep} . The choice process was modelled by a standard softmax process on the decision variable DV_t .

The computational model could include an initial sampling phase during which sampling behaviour is not controlled by the decision variable described above. Instead, during this initial phase at the beginning of each game, the model acquires, in turn, a number of samples from each option. This initial sampling phase is controlled by a number n_{ini} of initial samples from each option and a threshold strength θ_{ini} that controls the probability that the constraint on the number of initial samples n_{ini} is enforced before sampling from the other option. A threshold strength $\theta_{ini} = 1$ means that the corresponding participant samples repeatedly the current option until the number of initial

samples n_{ini} has been acquired. By contrast, a threshold strength $\theta_{\text{ini}} = 0$ means that the corresponding participant does not use an initial sampling phase.

Best-fitting parameter estimates for the model, which could either include or not include an initial sampling phase, were obtained through maximum likelihood estimation (MLE) through gradient descent on the negative model log-likelihood, using the ‘interior point’ algorithm implemented by the `fmincon` function in MATLAB. Bayesian model selection for comparing variants of the model with vs. without the initial sampling phase was conducted using a random-effects approach, which assumes that different participants may rely on different models. It consists in estimating the distribution over models that participants are drawn from. We used the Dirichlet parameterization of the random-effects approach implemented in SPM12 (Wellcome Trust Center for Human Neuroimaging; <http://www.fil.ion.ucl.ac.uk/spm>).

Note that the computational model described above nests the rational sampling agent for the DRAW and GUESS conditions. The rational agent does not include an initial sampling phase in either condition. In the DRAW condition, the rational agent chooses based solely on the first component: the difference between the mean colour mixture estimates signed by the target colour, with $\beta_{\text{tar}} \rightarrow \infty$ (i.e., an argmax choice process). In the GUESS condition, the rational agent chooses based solely on the difference between the information already acquired from each option, with $\beta_{\text{inf}} \rightarrow \infty$. The rational agent was simulated for each participant and each condition to compare its behavioural measures to those of tested participants.

A ‘knock-out’ simulation procedure based on best-fitting parameter estimates from the ‘full’ model to test the effects of the two suboptimal aspects of the model (the initial sampling phase and the repetition bias) on behavioural measures in the different conditions. For this purpose, separate model simulations were conducted after inactivating the initial sampling phase and the repetition bias independently of each other (by setting $\theta_{\text{ini}} = 0$, or $b_{\text{rep}} = 0$, none, or both).

Statistical analyses

Performance In all conditions, performance is reported in single-target games only as computing performance in target-absent or both-target games was not interpretable. Because performance is typically not distributed normally across participants, we performed Wilcoxon signed-ranks tests to compare participants’ performance to chance level (50%), or between conditions. Effect sizes are reported as $r = z/\sqrt{N}$, where N corresponds to the sample size. To compare performance between groups of participants with and without an initial sampling phase, we performed rank-sum tests for a difference in median between independent samples. Effect sizes are reported as $r = z/\sqrt{N_1 + N_2}$. We used a cluster-level permutation-based analysis to identify time periods where the fraction of target-congruent choices was significantly different from chance performance. We first identified clusters of contiguous trials where choice direction was significantly different from chance using a two-tailed p -value of 0.05 at the trial-wise level (t -tests against 50%). We extracted the statistic of each identified cluster as its sum of t -values. We then computed ‘null’ distributions for the cluster statistic by extracting the highest absolute cluster statistic for simulated random choices ($N = 1,000$ simulations). We then computed the empirical cluster-level p -value as the fraction of ‘null’ cluster statistics exceeding the observed cluster statistic.

Sampling imbalance The mean growth rate of sampling imbalance over the course of each game was estimated through linear regression as a linear coefficient. We performed t -tests against zero on these coefficients, and paired t -tests to compare the growth rate of sampling imbalance

between conditions. To test whether the increase in sampling imbalance throughout each game deviated from linearity, we compared participants' sampling imbalance profiles to the ones predicted by linear regression, and we tested these deviations using t -tests against zero, and paired t -tests between conditions. Effect sizes are reported as Cohen's d . We used a cluster-level permutation based analysis to identify time periods when participant's choice imbalances deviated significantly from the linear fits. We first identified clusters of contiguous trials where choice imbalance was significantly different from linear fits using a two-tailed p -value of 0.05 at the trial-wise level (paired t -tests). We extracted the statistic of each identified cluster as its sum of t -values. Second, we computed 'null' distributions for the cluster statistic by extracting the highest absolute cluster statistic for simulated random choices ($N = 1,000$ simulations). We then computed the empirical cluster-level p -value as the fraction of 'null' cluster statistics exceeding the observed cluster statistic.

Fraction of re-sampling the same option We performed repeated-measures ANOVAs on the fraction of re-sampling from the same option as a function of the number of previous samples acquired from this option. We tested the interaction between the condition (DRAW and GUESS in the discovery and replication datasets; DRAW, GUESS and FIND in the third dataset) and the number of previous samples acquired from the current option (from 1 to 10). Effect sizes are reported as $\eta^2_G = SS_{\text{eff}} / (SS_{\text{eff}} + \text{sum}(SS_{\text{res}}))$, where SS_{eff} corresponds to the sum of squares of the effect (interaction) of interest, and SS_{res} is the aggregate sum of squares of the residual errors. We used cluster-level permutation based tests to identify for which numbers of acquired samples the fraction of resampling choices was significantly different between conditions. First, we identified clusters of contiguous numbers of previous trials where the fraction resample was significantly different between conditions (p -value obtained from a paired-sample t -tests < 0.05). We extracted the statistic of each identified cluster as its sum of t -values. Then we computed the 'null' distribution for the cluster sum of t -value by shuffling the condition labels of individual games within participants. For each participant, we re-assigned the game choices to each condition randomly and recomputed the cluster statistic in the same way ($N = 1,000$ permutations). We then computed the empirical cluster-level p -value as the fraction of 'null' cluster statistics exceeding the observed cluster statistic.

Choice similarity Because choice similarity is typically not distributed normally across participants, we performed Wilcoxon signed-rank tests to compare participants' choice similarity to chance level (50%). For each game, each sampling choice was compared to all the other sampling choices in the same game, resulting in a similarity score for each pair of choice positions within each game, which was then averaged over games of the same condition.

Choice direction We used a cluster-level permutation-based analysis to identify time periods where choices were significantly target- or uncertainty-congruent. We first identified clusters of contiguous trials where choice direction was significantly different from chance using a two-tailed p -value of 0.05 at the trial-wise level (t -tests against 50%). We extracted the statistic of each identified cluster as its sum of t -values. We then computed 'null' distributions for the cluster statistic by extracting the highest absolute cluster statistic for simulated random choices ($N = 1,000$ simulations). We then computed the empirical cluster-level p -value as the fraction of 'null' cluster statistics exceeding the observed cluster statistic. We used a similar procedure to compare the fractions of uncertainty-congruent choices between conditions. In this case, we computed 'null' distributions by permuting randomly the condition labels of individual games and extracting cluster statistics for these permuted conditions.

Model parameter estimates We performed t -tests against zero and paired t -tests between conditions for unbounded sensitivity parameters β_{tar} and β_{inf} , and for the repetition bias b_{rep} . Effect sizes are reported as Cohen's d . For bounded parameters n_{ini} and θ_{ini} , we performed Wilcoxon signed-rank tests between conditions. Effect sizes are reported as $r = z/\sqrt{N}$. We compared parameter estimates between groups of participants with and without an initial sampling phase using two-sample t -tests for unbounded sensitivity parameters, and rank-sum tests for bounded parameters. In this case, effect sizes are reported as $d = t\sqrt{1/N_1 + 1/N_2}$ for unbounded sensitivity parameters, and as $r = z/\sqrt{N_1 + N_2}$ for bounded parameters.

Data and code availability

The code used to run the experiments described in the study, the behavioural data, and the main analysis pipeline are available on a public repository (<https://gitlab.com/cle-a/colpub>). The main analysis pipeline has been uploaded prior to the collection of the replication dataset on a dedicated repository (<https://gitlab.com/cle-a/colrep>). The model fitting code has been updated during the analysis of the third dataset, and the final version is available on the main repository.

References

1. Gittins, J. C. Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society: Series B (Methodological)* **41**, 148–164 (1979).
2. Sutton, R. S. & Barto, A. G. *Introduction to reinforcement learning*. vol. 135 (Cambridge: MIT Press, 1998).
3. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction*. (MIT press, 2018).
4. Mehlhorn, K. *et al.* Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision* **2**, 191–215 (2015).
5. Schulz, E. *et al.* Structured, uncertainty-driven exploration in real-world consumer choice. *Proc Natl Acad Sci USA* **116**, 13903–13908 (2019).
6. Daw, N. D., O’Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879 (2006).
7. Cohen, J. D., McClure, S. M. & Yu, A. J. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Phil. Trans. R. Soc. B* **362**, 933–942 (2007).
8. Schulz, E., Franklin, N. T. & Gershman, S. J. Finding structure in multi-armed bandits. *Cognitive Psychology* **119**, 101261 (2020).
9. Hills, T. T. & Hertwig, R. Information Search in Decisions From Experience: Do Our Patterns of Sampling Foreshadow Our Decisions? *Psychol Sci* **21**, 1787–1792 (2010).
10. Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A. & Cohen, J. D. Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General* **143**, 2074–2081 (2014).
11. Wilson, R. C., Bonawitz, E., Costa, V. D. & Ebitz, R. B. Balancing exploration and exploitation with information and randomization. *Current Opinion in Behavioral Sciences* **38**, 49–56 (2021).
12. Gottlieb, J., Oudeyer, P.-Y., Lopes, M. & Baranes, A. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in Cognitive Sciences* **17**, 585–593 (2013).
13. Findling, C., Skvortsova, V., Dromnelle, R., Palminteri, S. & Wyart, V. Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nat Neurosci* **22**, 2066–2077 (2019).
14. Sharot, T., Korn, C. W. & Dolan, R. J. How unrealistic optimism is maintained in the face of reality. *Nat Neurosci* **14**, 1475–1479 (2011).
15. Sharot, T. & Garrett, N. Forming Beliefs: Why Valence Matters. *Trends in Cognitive Sciences* **20**, 25–33 (2016).
16. Lefebvre, G., Lebreton, M., Meyniel, F., Bourgeois-Gironde, S. & Palminteri, S. Behavioural and neural characterization of optimistic reinforcement learning. *Nat Hum Behav* **1**, 0067 (2017).
17. Lau, B. & Glimcher, P. W. Dynamic response-by-response models of matching behaviour in rhesus monkeys. *Journal of the Experimental Analysis of Behavior* **84**, 555–579 (2005).
18. Padoa-Schioppa, C. Neuronal Origins of Choice Variability in Economic Decisions. *Neuron* **80**, 1322–1336 (2013).

19. Shackman, A. J. *et al.* The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nat Rev Neurosci* **12**, 154–167 (2011).
20. Kool, W. & Botvinick, M. Mental labour. *Nat Hum Behav* **2**, 899–908 (2018).
21. Vogel, T. A., Savelson, Z. M., Otto, A. R. & Roy, M. Forced choices reveal a trade-off between cognitive effort and physical pain. *eLife* **9**, e59410 (2020).
22. Gureckis, T. M. & Markant, D. B. Self-Directed Learning: A Cognitive and Computational Perspective. *Perspect Psychol Sci* **7**, 464–481 (2012).
23. Markant, D. B. & Gureckis, T. M. Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General* **143**, 94–122 (2014).
24. Gottlieb, J. & Oudeyer, P.-Y. Towards a neuroscience of active sampling and curiosity. *Nat Rev Neurosci* **19**, 758–770 (2018).
25. Vergassola, M., Villermaux, E. & Shraiman, B. I. ‘Infotaxis’ as a strategy for searching without gradients. *Nature* **445**, 406–409 (2007).
26. Wald, A. & Wolfowitz, J. Optimum Character of the Sequential Probability Ratio Test. *Annals of Mathematical Statistics* **19**, 326–339 (1948).
27. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review* **113**, 700–765 (2006).
28. Desender, K., Boldt, A. & Yeung, N. Subjective Confidence Predicts Information Seeking in Decision Making. *Psychol Sci* **29**, 761–778 (2018).
29. Hunt, L. T. *et al.* Formalizing planning and information search in naturalistic decision-making. *Nat Neurosci* **24**, 1051–1064 (2021).
30. Locander, W. B. & Hermann, P. W. The Effect of Self-Confidence and Anxiety on Information Seeking in Consumer Risk Reduction. *JOURNAL OF MARKETING RESEARCH* **7** (1979).
31. Sadeghiyeh, H. *et al.* Temporal discounting correlates with directed exploration but not with random exploration. *Sci Rep* **10**, 4020 (2020).
32. Schulz, L., Rollwage, M., Dolan, R. J. & Fleming, S. M. Dogmatism manifests in lowered information search under uncertainty. *Proc Natl Acad Sci USA* **117**, 31527–31534 (2020).
33. Kleiner, M., Brainard, D. & Pelli, D. What’s new in Psychtoolbox-3? 89 (2007).

Acknowledgments

This work was supported by a starting grant from the European Research Council (ERC-StG-759341) awarded to V.W. V.C. was supported by the Agence Nationale de la Recherche (ANR-16-CE37-0012-01 and ANR-19-CE37-0014-01). All authors benefited from an institutional grant from the Agence Nationale de la Recherche (ANR-17-EURE-0017).

Author contributions

V.W. acquired funding and supervised the study. C.A., V.C. and V.W. designed the experiments. C.A. acquired the data. C.A. and V.W. designed the model. C.A. and V.W. analysed the data. C.A., V.C. and V.W. discussed the results. C.A. and V.W. drafted the manuscript. C.A., V.C. and V.W. reviewed and edited the manuscript.

Competing interests

The authors declare no competing interests.

Supplementary information

Competing cognitive pressures on human exploration in the absence of trade-off with exploitation

Clémence Alméras, Valérian Chambon, and Valentin Wyart

Supplementary Fig. 1 | Rational sampling agent. Simulations of the behaviour of a rational sampling agent (full colours, plain lines, group means $\pm$ 95% CI) alongside participants' behaviour (depicted in lighter colours, dashed lines, group means $\pm$ 95% CI). Each panel features data for the discovery (left) and validation (right) sets. **a**, Sampling imbalance. Difference between the number of choices for the most and the least chosen option at each choice throughout the games. **b**, Option re-sampling. Fraction of re-sampling the same option as a function of the number of times the option was previously sampled (group means $\pm$ 95% within-subject CI for simulations and participants). **c**, Choice direction. Left: fraction of target-congruent choices, defined as choices towards the currently most target-like option. Right: fraction of uncertainty-congruent choices, defined as choices towards the currently most uncertain option.

a choice direction

b sampling imbalance

c option re-sampling

Supplementary Fig. 2 | Discovery dataset: knock-out model simulations. Model simulations based on fitted parameters in the discovery dataset (red curves, group means), against participants' behaviour (condition-coloured curves, group means $\pm$ 95% CI). Columns displays different simulated 'knock-out' versions of the model: the model without initial sampling phase nor repetition bias, featuring only sensitivity to the target colour and to the uncertainty of options (first column), a model without the initial sampling phase (second column), a model without the repetition bias (third column), and a full model featuring both sensitivity parameters as well as the initial sampling phase and repetition bias (fourth column). **a**, Choice direction. Top: fraction of target-congruent choices, defined as choices towards the currently most target-like option. Bottom: fraction of uncertainty-congruent choices, defined as choices towards the currently most uncertain option. **b**, Sampling imbalance. Difference between the number of choices for the most and the least chosen option at each choice throughout the games. **c**, Option re-sampling. Fraction of re-sampling the same option as a function of the number of times the option was previously sampled (group means $\pm$ 95% within-subject CI for participants).

Supplementary Fig. 3 | Discovery dataset: knock-out model simulations. Simulated choice-choice similarity matrix based on fitted parameters in the discovery dataset. Columns displays different simulated 'knock-out' versions of the model: the model without initial sampling phase nor repetition bias, featuring only sensitivity to the target colour and to the uncertainty of options (first column), a model without the initial sampling phase (second column), a model without the repetition bias (third column), and a full model featuring both sensitivity parameters as well as the initial sampling phase and repetition bias (fourth column). The rightmost column displays the average choice-choice similarity of participants' choices (group-level means). **a**, Simulations of DRAW games. **b**, Simulations of GUESS games. The similarity between simulated and participants' choice-choice similarity matrices is indicated next to each simulated matrix.

a choice direction

b sampling imbalance

c option re-sampling

Supplementary Fig. 4 | Replication dataset: knock-out model simulations. Model simulations based on fitted parameters in the discovery dataset (red curves, group means), against participants' behaviour (condition-coloured curves, group means $\pm$ 95% CI). Columns displays different simulated 'knock-out' versions of the model: the model without initial sampling phase nor repetition bias, featuring only sensitivity to the target colour and to the uncertainty of options (first column), a model without the initial sampling phase (second column), a model without the repetition bias (third column), and a full model featuring both sensitivity parameters as well as the initial sampling phase and repetition bias (fourth column). **a**, Choice direction. Top: fraction of target-congruent choices, defined as choices towards the currently most target-like option. Bottom: fraction of uncertainty-congruent choices, defined as choices towards the currently most uncertain option. **b**, Sampling imbalance. Difference between the number of choices for the most and the least chosen option at each choice throughout the games. **c**, Option re-sampling. Fraction of re-sampling the same option as a function of the number of times the option was previously sampled (group means $\pm$ 95% within-subject CI for participants).

Supplementary Fig. 5 | Replication dataset: knock-out model simulations. Simulated choice-choice similarity matrix based on fitted parameters in the discovery dataset. Columns displays different simulated 'knock-out' versions of the model: the model without initial sampling phase nor repetition bias, featuring only sensitivity to the target colour and to the uncertainty of options (first column), a model without the initial sampling phase (second column), a model without the repetition bias (third column), and a full model featuring both sensitivity parameters as well as the initial sampling phase and repetition bias (fourth column). The rightmost column displays the average choice-choice similarity of participants' choices (group-level means). **a**, Simulations of DRAW games. **b**, Simulations of GUESS games. The similarity between simulated and participants' choice-choice similarity matrices is indicated next to each simulated matrix.

a choice direction

b sampling imbalance

c option re-sampling

Supplementary Fig. 6 | Third dataset: knock-out model simulations. Model simulations based on fitted parameters in the discovery dataset (red curves, group means), against participants' behaviour (condition-coloured curves, group means $\pm$ 95% CI). Columns displays different simulated 'knock-out' versions of the model: the model without initial sampling phase nor repetition bias, featuring only sensitivity to the target colour and to the uncertainty of options (first column), a model without the initial sampling phase (second column), a model without the repetition bias (third column), and a full model featuring both sensitivity parameters as well as the initial sampling phase and repetition bias (fourth column). **a**, Choice direction. Top: fraction of target-congruent choices, defined as choices towards the currently most target-like option. Bottom: fraction of uncertainty-congruent choices, defined as choices towards the currently most uncertain option. **b**, Sampling imbalance. Difference between the number of choices for the most and the least chosen option at each choice throughout the games. **c**, Option re-sampling. Fraction of re-sampling the same option as a function of the number of times the option was previously sampled (group means $\pm$ 95% within-subject CI for participants).

Supplementary Fig. 7 | Third dataset: knock-out model simulations. Simulated choice-choice similarity matrix based on fitted parameters in the discovery dataset. Columns displays different simulated ‘knock-out’ versions of the model: the model without initial sampling phase nor repetition bias, featuring only sensitivity to the target colour and to the uncertainty of options (first column), a model without the initial sampling phase (second column), a model without the repetition bias (third column), and a full model featuring both sensitivity parameters as well as the initial sampling phase and repetition bias (fourth column). The rightmost column displays the average choice-choice similarity of participants’ choices (group-level means). **a**, Simulations of DRAW games. **b**, Simulations of GUESS games. **c**, Simulations of FIND games. The similarity between simulated and participants’ choice-choice similarity matrices is indicated next to each simulated matrix.