Individual Curiosity Modulates Exploration in Sequential Book Selection

Xuanjun Gong
xjgong@tamu.edu

Texas A&M University https://orcid.org/0000-0001-9642-8885
Cuihua Shen
UC Davis https://orcid.org/0000-0003-1645-8211
Erie Boorman
University of California, Davis https://orcid.org/0000-0002-8438-4499
Richard Huskey
University of California Davis https://orcid.org/0000-0002-4559-2439

Article

Keywords: Book selection, Exploration-Exploitation, Curiosity, Computational modeling

Posted Date: December 5th, 2024

DOI: https://doi.org/10.21203/rs.3.rs-5442881/v1

Additional Declarations: There is **NO** Competing Interest.

Individual Curiosity Modulates Exploration in Sequential Book Selection

Abstract

How do people choose what book to read? At times, people choose to exploit well-known options that are likely to lead to high enjoyment. However, readers must also effectively explore novel books in order to learn about less-known alternatives that might lead to high enjoyment. It is unknown precisely how and why readers make these sequential book selection decisions. By placing book options in a semantic embedding space, we show that people decide which book to explore using a structured generalization mechanism based on semantic similarities between known and unknown books and a directed exploration mechanism that incentivizes seeking books in high uncertainty. In addition, we demonstrate that people's directed and random book exploration patterns are modulated by individual differences in curiosity, which fosters reading enjoyment and promotes exploring unfamiliar books. In summary, our study demonstrates that these computational mechanisms generalize to a new and ecologically valid context in order to drive consequential exploratory decisions with important real-world implications.

Keywords: Book selection, Exploration-Exploitation, Curiosity, Computational modeling

1 Introduction

Books are one of the oldest[1] and most popular forms of mass media[2]. People read books to acquire knowledge and skills, seek information, and enjoy leisure and entertainment[3–6]. Despite books' long history, popularity, and importance, the mechanisms that explain what books people choose to read and how they make such decisions remain poorly understood.

The problem of book selection is complicated due to substantial uncertainty associated with searching for high-value books in a vast and diverse book space and correctly anticipating the value of available book options. In order to make effective reading decisions, readers face at least two challenging questions: (1) which book has the highest subjective value that results in an immediate and highly rewarding reading experience, and (2) which book has the highest informative value that aids in correctly anticipating the value of future book selection. To overcome these two challenges, readers must successfully navigate a tradeoff between exploring novel but informative options or exploiting known high-value options. Therefore, readers need to decide when and which options to explore for an optimal reading experience. This tradeoff is known as the exploration-exploitation dilemma.

The exploration-exploitation dilemma is a famous problem for human search decisions in a hypothetical option space, such as food foraging[7], memory search[8], information seeking[9], and knowledge acquisition[10]. There is no optimal solution to this dilemma given its intractable computational complexity[11]. Nevertheless, many heuristic-based algorithms, such as generalization[12], random exploration[13], directed exploration[14], and optimal foraging[15], have been proposed to solve this issue. These algorithms formally specify the sequential dependencies among prior choices in order to mechanistically explain and predict future choices[16]. However, there has been very little work testing how mechanisms revealed in the lab generalize to such real-world settings and how these play out in a particular real-world decision space[17].

Among these mechanisms, generalization and directed exploration strategies are especially appealing for studying book selection. Specifically, the generalization mechanism addresses readers' first challenge in estimating the rewards of novel options based on past reading experiences. It assumes options are embedded in a structured space such that reward estimations for different options are correlated with the feature similarities of other options in the embedding space. On the other hand, the directed exploration mechanism addresses readers' second challenge of effectively reducing uncertainty associated with book selection. It hypothesizes that people evaluate the uncertainty of reward estimations and are motivated to choose books with a high information bonus, which inflates the subjective value of options with high uncertainty.

These decision strategies are linked to individual differences such as age[18] and impulsivity[19] and an increasingly resolved brain-network architecture[13, 20, 21]. However, prior findings are mainly based on artificial laboratory experiments, where participants make choices within a finite option space repetitively to maximize payoffs returned as objective rewards (e.g., points, money). It still remains unclear if and how these mechanisms apply to people's real-world selection behaviors, especially for book selections, where (1) the option space is theoretically infinite, (2) each choice is rarely repeated, (3) experienced rewards are implicit and subjective[22], and (4) selection is driven by intrinsic motives, such as interests and curiosity[23], rather than external monetary motives[24]. These distinctions obfuscate the empirical applicability of the exploration-exploitation theories in real-world decision problems in general[17, 22], and specifically for people's book selection. In order to understand people's real-world book selection, we investigate two broad questions.

First, we study how different exploration-exploitation mechanisms, with a focus on generalization and directed exploration, characterize people's real-world book selection. We approach this question in parallel by investigating people's learning and selection sequences among two empirical book selection datasets. The first consists of large-scale real-world book selection digital trace data comprised of nearly thirty-five thousand readers and more than two million choices. These observational data were experimentally confirmed in a second empirical dataset that resembles real-world book selection. Convergent evidence across two datasets shows that people learn to select more favorable books following a generalization mechanism and are biased toward books with uncertain rewards following a directed exploration strategy.

Second, we clarify that the underlying mechanisms of book exploration are governed by curiosity, referred to as an intrinsic drive for information and learning[25]. People's preference for fiction is attached to their preference for exploration[26]. Trait curiosity impacts people's visual attention and information-seeking behavior[9, 27–29]. And curiosity promotes choices to gather information[30] and explains people's browsing choices on Wikipedia[10]. Therefore, we hypothesized and found that individual differences in curiosity traits regulate people's book selection. To approach this question, we draw on state-of-the-art conceptualizations that treat curiosity as a personality trait that varies in multiple dimensions [31] in order to explain people's exploratory book selection behavior in the book selection.

The main contributions of this paper are threefold. First, we introduce a novel behavioral modeling paradigm to analyze real-world digital trace and experimental data in parallel. This approach addresses a complex real-world domain-specific decision problem - how people choose books - by applying well established domain-general decision theories to both high-control and naturalistic settings. Second, we found that book selection is regulated by books' semantic features via a reward generalization mechanism that leads people to select more similar and more favorable books over time. Additionally, people relax semantic feature constraints in their book choices through a directed exploration mechanism in order to deliberately seek unfamiliar books with high uncertainty. Jointly, these selection strategies describe the way people navigate the exploration-exploitation dilemma when choosing what book to read. Finally, we demonstrate that individual differences in random and directed book exploration patterns are explained by the thrill seeking and joyous exploration dimensions of curiosity. In summary, curiosity serves as an intrinsic incentive that boosts book reading enjoyment and encourages reading books with high reward uncertainty.

2 Results

We gathered people's book selection and rating sequences from a real-world book review dataset[32] collected on Amazon, one of the world's largest book purchase and review databases. This dataset consists of 2,083,630 book rating records from 35,478 readers. In addition, we conducted a book selection experiment that resembles the real-world using a highly controlled decision making task. This experiment asked participants to make sequential preferential choices among 225 possible book options in a structured grid space. In what follows, we first report the descriptive characteristics of book selections both in real-world and experimental environments embedded in a semantic space. Then, we provide behavioral signatures and computational modeling evidence supporting reward generalization and directed exploration as mechanisms that govern book selection in both real-world and experimental datasets. Finally, we demonstrate that curiosity modulates these exploration mechanisms in book selection and facilitates reading enjoyment in book exploration.

2.1 Book selection as exploration in a semantic embedding space

In real-world book selection, the number of book options is enormous, and these books differ from each other in many ways. Conceptually, books are embedded in a semantic space such that each book can be represented by its semantic features. To reconstruct this semantic space for books included in our dataset, we scrapped the synopsis for each book option from the GoodReads website, one of the largest book metadata databases. We then transformed the summary for each book into a multi-dimensional semantic embedding vector using an advanced natural language processing technique (Figure 1A). This embedding represented all books in the real-world Amazon dataset, and a subset of 225 books for the experimental dataset (Supplemental Section 1). To test the validity of this semantic embedding method, we collected perceived pairwise dissimilarity ratings among a subsample of the 225 book options selected for the experimental dataset ($n = 22$ unique options; 10%) from 248 participants recruited from Prolific. A Mantel test[33] shows a significant positive association between human-perceived dissimilarities and machine-evaluated semantic dissimilarities ($r = 0.535, Z = 7.110, p < 0.001$; all significance tests were two-tailed; (Supplemental Section 2).

Previous works widely considers that people's exploration decisions resemble foraging behavior in a patchy environment[12, 34], which exhibits a clumpy spatial distribution of resources[35]. Consistent with this hypothetical idea, we found that book options are naturally clustered in a patchy format and grouped by their genre in semantic space (Figure 1A). In addition, we found that the book options have a higher than null clustering tendency[36] ($H = 0.754$; Supplemental Section 3), which indicates that book options are clustered rather than randomly dispersed in the embedding space.

In this embedding space, book selections can be ordered as a sequence of non-repetitive discrete choices represented by a set of numeric semantic features (Figure 1B). We found that people's book explorations were constrained by semantic distance among options. The observed distances between consecutive choices ($M = 1.170, SD = 0.155$) are significantly smaller ($Z = -1172.34, p < 0.001$; Figure 1C)) than semantic distances among randomly chosen book options ($M = 1.301, SD = 0.090$). Thus, readers are more likely to explore a book that is semantically similar to previously read books compared to randomly chosen alternatives, which is consistent with previous research on memory retrieval and purchase behaviors[17, 37].

These results depict people's book selections as a trajectory of explorations in a patchy environment embedded in a multidimensional semantic space, where people decide which book to read based on the semantic features of previously read books and available book options. We designed a multi-armed bandit task, a widely used experimental paradigm for studying exploration decisions[38], to recreate this real-world book selection scenario in a controlled experimental decision environment[12]. In this task, a 15x15 2-D grid was displayed, with each cell representing one of 225 unique book options selected from the real-world dataset (Figure 1D). The pairwise spatial distances among options in this grid space were designed to represent the semantic distance among the corresponding books in a way such that semantically

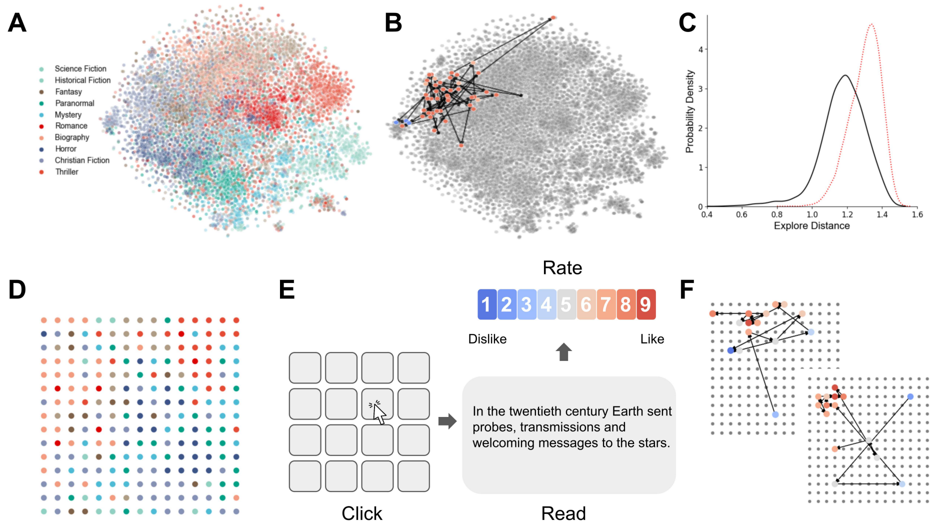

Fig. 1 We gathered two datasets of people's book selections from real-world and experimental settings. The real-world dataset comprises large-scale records of sequential reading choices and subjective reading experience ratings. These book choices can be arranged in a multidimensional semantic embedding space which numerically represents the semantic meaning of the book synopsis. The experimental dataset collects book selection records from a multi-armed bandit task, which simulates the real-world book selection environment. **A.** Real-world books are represented in a semantic embedding space. A sample of 10,000 books were plotted in a 2-D space, which is t-SNE transformed from a 384-dimensional semantic embedding space. Each point represents a unique book and is colored by its genre. Books are naturally clustered by their genre in this semantic embedding space, and together constitute a patchy book foraging environment. **B.** Readers make sequential book selection trajectories in the semantic book space. Colored points (redish color represents a rating greater than 3 stars; bluish color represents a rating less than or equal to 3 stars) depict selected books while gray points represent books available but not selected. The arrows connecting two points denote the sequential order of consecutive book choices. **C.** Probability density plot of the distribution of people's exploration distance (black curve), which is measured as the Euclidean distance between semantic embedding vectors of consecutive book choices. The red distribution denotes the null distribution of exploration distance, which assumes people randomly select books. **D.** The experimental book selection landscape. A total of 15x15 options were arranged in a grid and presented for participants to make book selections. Each point encodes a book option, and the color encodes the genre of the books. Book options were selected from the real-world dataset, and arranged in a way such that semantically similar books are placed close to each other. **E.** Book selection experimental procedure. Participants completed a total of 15 trials of a click-read-rate task, where they clicked one option from the grid, read the synopsis of the book, and then rated their reading enjoyment. **F.** Participant sequential book selection trajectories in the experiment. Colored points (redish color indicates a favorable reading experience; bluish color indicates negative reading experience) represent people's choices and gray points represent books available but not selected. The arrow connecting two points encodes the sequential order of two consecutive choices.

similar options were placed close to each other and grouped in patches. Using this experimental paradigm, we collected sequences of book choices and reading enjoyment ratings from 250 participants (Figure 1E) and conducted further analysis to explore people's sequential book selection patterns (Figure 1F).

2.2 Reward learning and generalization

We assessed whether or not people learn to make better book choices over time in both datasets. Book ratings in the real-world datasets range from 1 star (lowest rating) to 5 stars (highest rating), while ratings in the experiment were measured by a 9-point scale ranging from 1 (extremely disliked) to 9 (extremely liked). We found an increasing trend in people's reading enjoyment over time (Figure 2 A&D) where correlations between the number of past reads and book ratings were significantly positive in both real-world ($r = 0.022, p < 0.001, 99\% CI = [0.021, 0.024]$) and experimental ($r = 0.106, p < 0.001, 99\% CI = [0.064, 0.147]$) datasets. The increased enjoyment indicates that people become better at choosing more favorable books over time, suggesting a learning curve in people's reading selection behaviors (Supplemental Section 4 reports alternative explanation tests[22]).

In addition, we looked at the evolving patterns of people's book explorations as a function of semantic distance between consecutive choices. We found that people's book exploration stabilized over time for both the real-world ($r = -0.006, p < 0.001, 99\% CI = [-0.008, -0.004]$) and experiment datasets ($r = -0.177, p < 0.001, 99\% CI = [-0.219, -0.135]$). This reflects a tendency towards less distant (in embedding space) options over time (Figure 2 B&E). Combined, people make more favorable and similar choices over time, indicating that they learn from previous book-reading experiences, gain better value estimations of book options, and choose high-value books accordingly. This mechanism is known as reward generalization[39], by which people use feature similarities to update their reward estimations as of function learning[38].

Importantly, reward generalization formalizes that, after a high-reward book reading experience, people estimate similar books to have higher values and do the contrary after a low-reward experience. Therefore, people tend to choose books that are similar to previously experienced high-reward books while avoiding books similar to low-reward books. Consistent with this prediction, we found that people tend to choose semantically similar books after a high-reward relative to a low-reward reading experience (Figure 2 C&F) among both the real-world ($r = -0.062, p < 0.001, 99\% CI = [-0.064, -0.060]$) and experimental datasets ($r = -0.315, p < 0.001, 99\% CI = [-0.353, -0.275]$). Thus, we found evidence indicating that readers employ a reward generalization strategy in book selection behaviors, which helps them quickly learn which option generates high rewards, discover their favorite book types, and improve their overall reading experiences (for sensitivity analyses, see Supplemental Section 5).

2.3 Directed and random exploration

People do not always choose books from their favorite genres. Sometimes, people select unfamiliar books and explore novel genres. This behavioral pattern suggests another critical exploration mechanism that might govern book selection—directed exploration—which involves deliberately seeking books with high novelty and uncertainty to gain knowledge while forgoing the immediate rewards of reading familiar books. Contrary to random exploration strategies, which specify that people's exploration is passively driven by the stochasticity of the decision-making process, directed

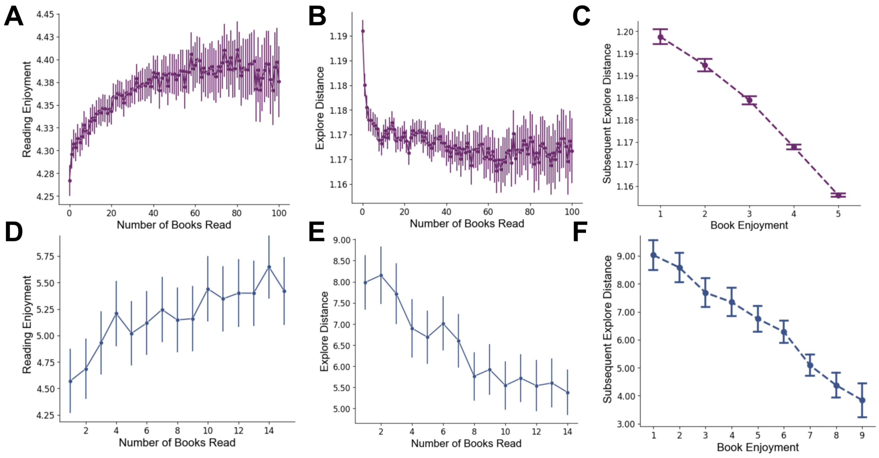

Fig. 2 Signatures of learning and reward generalization in the real-world data and the experimental data. **A.** Real-world data; **D.** Experiment data: People's reading enjoyment by the number of books that have been read. **B.** Real-world data; **E.** Experiment data: The exploration distance, measured as the Euclidean distance between semantic embedding vectors of consecutive book choices, by the number of books that have been read. **C.** Real-world data; **F.** Experiment data: The exploration distance by its immediate preceding reading rating. The points in the line plot indicate the mean estimates, and the vertical lines indicate the 99% confidence intervals.

exploration hypothesizes that people actively add an information bonus to the reward estimation of high-uncertainty options in order to encourage choices toward uncertain options.

We operationalized book uncertainty by evaluating the number of ratings and the variance of ratings in the GoodReads metadata. Compared to books that have rarely been rated or have heterogeneous ratings, books with a large number of homogeneous ratings should give readers more confidence in estimating their reading rewards and, hence, have lower uncertainty. Indeed, people's book exploration distance is associated with the number of ratings and rating variance (Figure 3 A&B). We found a significant negative correlation between the logarithm of rating counts and exploration distance ($r = -0.042, p < 0.001, 99\% CI = [-0.044, -0.040]$), as well as a significant positive correlation between rating variance and the exploration distance ($r = 0.076, p < 0.001, 99\% CI = [0.074, 0.078]$) (for sensitivity analyses, see Supplemental Section 5).

2.4 Computationally modeling sequential selection

We further probed the signatures of reward generalization and directed exploration by evaluating and comparing their corresponding computational models: the Guassian Regression (GP) model and the Upper Confidence Bound (UCB) model. The GP model only specifies a generalization mechanism to estimate the mean reward for books using a Gaussian regression function of previous reading enjoyment and the semantic similarities between previously read books and estimated-but-unread books. By comparison, the UCB model, representing the directed exploration mechanism, calculates the upper confidence bound of value estimates by adding the uncertainty of

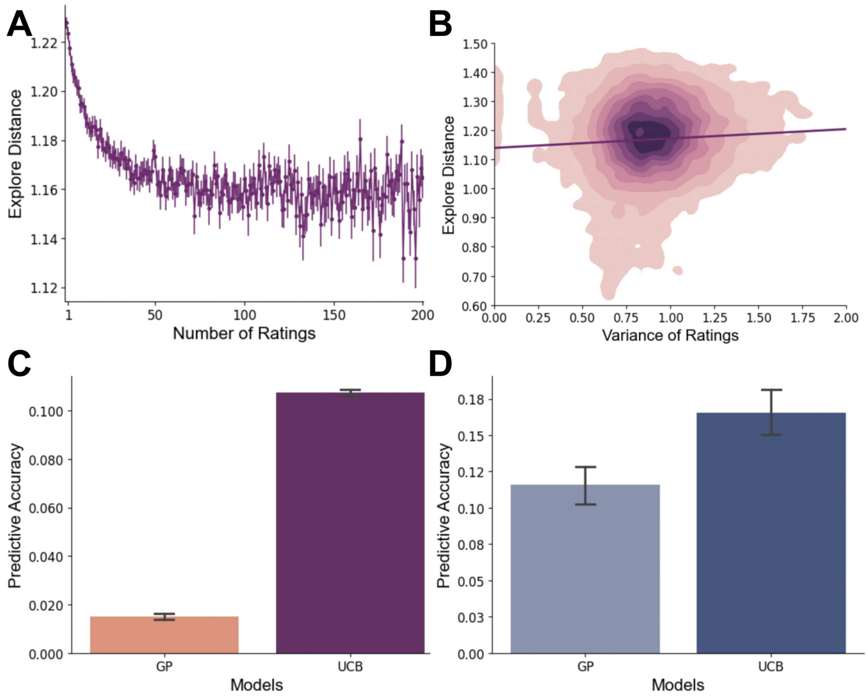

Fig. 3 Signatures of directed exploration in the real-world data, and model prediction accuracy in both real-world and experiment data. **A.** Exploration distance by the number of ratings for subsequent book choices. Mean estimates were plotted as points, and 99% confidence intervals were plotted as vertical lines. **B.** Bivariate distributions of exploration distance and the variance of ratings for subsequent book choices. Darker colors indicate higher probability density. The solid line is plotted as the regression line of variance of ratings regressing on exploration distance. The predictive accuracies for the Gaussian Process (GP) and the Upper Confidence Bound (UCB) models for each subject are plotted for real world dataset (**C**) and experimental data (**D**). Bar height indicates a mean estimate for the predictive accuracy for all subjects in the dataset, and the vertical line indicates the 99% confidence interval of the predictive accuracy.

reward estimation as an information bonus to the value estimate and then utilizes the upper confidence bound to probabilistically determine book selections (Supplemental Section 6).

The GP model consists of two decision parameters. The first is a generalization parameter (λ), which controls the length scale of the radical kernel function in a Gaussian regression model. In principle, $\lambda \to 0$ leads to zero generalization and independent value estimation among options, whereas $\lambda \to \infty$ leads to maximum generalization, such that the dependency of value estimation is linear to feature distances. The second parameter is a temperature parameter (τ) of the softmax function that determines the randomness in the probabilistic mechanism of exploration choices in a way that $\tau \to 0$ leads to zero randomness such that the highest valued option is always chosen, whereas $\tau \to \infty$ leads to maximum randomness with a uniform probability of selecting any option. In addition to the generalization and temperature parameters, the UCB model contains a third decision parameter, exploration bonus (β), which controls the extent of directed exploration, with higher β leading to a stronger bias towards options with high uncertainty.

We evaluated the one-step-ahead prediction accuracy for the book selection sequences for each subject in both the real-world and experimental datasets (Figure 3C&D). This prediction accuracy is calculated as the empirical model fit that is normalized by a null random model, which assumes an equal probability of selection. Since the number of options in real-world book selection is massive and unobservable because people can choose any published book in principle, following[17], we calculated the model's estimated likelihood of the observed choice relative to an artificial null book option with averaged semantic features of all possible book options (see Supplemental Section 7 for a discussion). We found that the GP model with the generalization and random exploration mechanisms performed better than chance prediction in both real-world and experimental datasets (real-world: $R2 = 0.015, Z = 30.011, p < 0.001, 99\%CI = [0.014, 0.016]$; experimental: $R2 = 0.116, Z = 17.024, p < 0.001, 99\%CI = [0.098, 0.133]$). Moreover, the UCB model with an additional directed exploration term (information bonus) had a significantly higher prediction accuracy compared to the GP model (real-world: $R2 = 0.107, Z = 512.134, p < 0.001, 99\%CI = [0.106, 0.109]$; experimental: $R2 = 0.165, Z = 12.984, p < 0.001, 99\%CI = [0.146, 0.185]$). Thus, our results support the existence of both a generalization mechanism, by which readers learn from previous reading experiences to make better book selections based on semantic similarities between books, as well as a directed exploration information bonus mechanism, by which readers are actively biased to read books that have a high uncertainty in order to obtain knowledge that aids future book selection.

2.5 Curiosity modulates exploration in book selections

Epistemic curiosity is an intrinsic motivator for uncertainty reduction and knowledge acquisition[9, 24, 40, 41]. Empirically, curiosity promotes information seeking via media use[42], and explains individual differences in media usage patterns[10, 31**?**]. Therefore, we asked if people's trait curiosity characterizes individual variability in book explorations. To answer this question, we identified computational phenotypes[43] of exploration decisions by estimating the decision parameters (i.e., τ, β) for each participant independently (Figure 4). Then, we applied multiple regression with the log-transformed decision parameters for exploration (i.e., τ, β) as dependent variables and the five-dimensional curiosity measure (i.e., joyous exploration, thrill-seeking, stress tolerance, derivative sensitivity, and social curiosity;[31] as independent variables while controlling for participants' age (an important covariate of human exploration patterns[18] and media selection [44]).

Regressing (Supplemental Section 8) the directed exploration bonus parameter (β) on joyous exploration revealed a positive association (Figure 4B; $b = 0.605, SE = 0.214, t(238) = 2.824, p = 0.005, 95\%CI = [0.183, 1.027]$). In addition, regressing random exploration temperature parameter (τ) on thrill-seeking shows a positive association (Figure 4C; $b = 0.225, SE = 0.092, t(238) = 2.440, p = 0.015, 95\%CI = [0.043, 0.406]$). By comparison, there is a negative association between the random exploration temperature parameter (τ) and joyous exploration (Figure 4D;

$b = -0.201, SE = 0.101, t(238) = -1.983, p = 0.049, 95\%CI = [-0.401, -0.001]$). Together, participants high in joyous exploration (the desire to seek out and take joy in new knowledge) practice less random but more directed exploration, suggesting a strong preference for books with high uncertainty by systematically assigning an information bonus to the uncertain books when evaluating their subjective value. On the other hand, participants high in thrill-seeking (the willingness to take risks in order to achieve high-variance experiences) have a stronger tendency to adopt a random exploration book selection strategy instead of following the directed exploration strategy.

Joyous exploration modulates people's directed exploration behaviors by attaching intrinsic incentives to high-uncertainty options. Following this idea, we looked at people's reward experiences in book exploration as related to their joyous exploration. We found that people with high joyous exploration (above a median split) generally experienced higher reading enjoyment compared to people with low joyous exploration (below a median split; Figure 4E). Furthermore, we identified that this difference in reading enjoyment mainly arises from reading semantically different books rather than semantically similar books (Figure 4F). Importantly, mixed-effect regression model with reading enjoyment as a dependent variable revealed a positive main effect of trait joyous exploration ($b = 0.183, SE = 0.084, Z = 2.191, p = 0.028, 95\%CI = [0.019, 0.347]$), and an interaction effect between the semantic similarities between consecutive books and people's joyous exploration trait ($b = 0.016, SE = 0.007, Z = 2.215, p = 0.032, 95\%CI = [0.001, 0.030]$; see Supplemental Table 4 for full results). Thus, compared to people with low joyous exploration, people with high joyous exploration gain higher rewards from their exploratory book choices. This result echoes previous studies showing that curiosity positively boosts experienced enjoyment during information-seeking[30, 45].

2.6 Robustness and additional analysis

We conducted linear regressions on directed and random exploration after removing non-significant covariates, and the results remain robust (Supplemental Section 9). We note that a considerable portion of the directed exploration parameter estimates are at a boundary close to zero. We consider this a feature rather than a bug because this near-zero parameter estimate captures meaningful exploration characteristics for non-exploratory participants. To verify the robustness of the regression results, we fit a censored regression model specifying a left boundary near zero, and the results remain consistent (Supplemental Section 10). Additional robustness checks were carried out to evaluate the stability of the parameter estimation, given that the global optimization method is non-deterministic. We repeated the optimization method 100 times independently, and parameter estimates for each decision parameter and each participant were reliably stable (Supplemental Section 11).

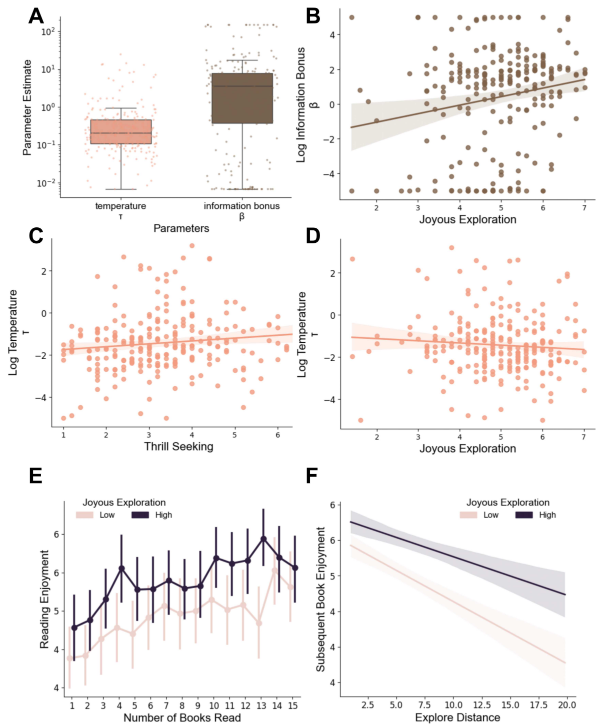

Fig. 4 Curiosity modulates exploration patterns in people's book selection. **A.** Boxplots of cross-validated parameter estimates of the temperature parameter τ (left box plot) and the information bonus parameter β (right box plot) for each participant in the experimental dataset. Parameter estimates for each participant are shown as points displayed for each box plot, and are horizontally jittered for improved visual interpretation. The box plots show the median (middle horizontal line), interquartile range (box), and 1.5-times interquartile range (whiskers). **B.** Logarithm of information bonus parameter β regressed on joyous exploration. **C.** Logarithm of temperature parameter τ regressed on thrill seeking. **D.** Logarithm of temperature parameter τ regressed on joyous exploration. The solid line indicates the estimated regression line, and the shaded area around the regression line represents the 95% confidence interval of regression coefficient estimates. **E.** Line plot of reading enjoyment by the number of books that have been previously read in the experimental dataset, for participants high (darker colored line) and low (lighter colored line) in joyous exploration. Mean estimates are plotted as points, and 95% confidence intervals are plotted as the vertical lines. **F.** The regression models of subsequent book enjoyment regressed on exploration distance for participants high (darker colored line) and low (lighter colored line) in joyous exploration.

3 Discussion

In real-world book selection scenarios, readers encountering the exploration-exploitation dilemma optimize reading enjoyment by identifying and selecting high-rewarding books while effectively reducing uncertainty about unknown books by exploring uncertain options. Heuristic decision mechanisms, including reward generalization and directed exploration, provide plausible solutions to address this computationally intractable problem. However, the unique domain-specific characteristics of book selection challenge the applicability of these domain-general mechanisms to this complex real-world decision problem. Here, we approached this challenge by integrating behavioral analysis and computational modeling of book selections on large-scale real-world digital trace data and then confirmed and replicated the findings with a behavioral experiment.

We identified that people's sequential book choices are represented and constrained in a patchy semantic embedding space, consistent with foraging patterns in semantic search[34]. Through a sequence of book explorations, readers rapidly improve their reading enjoyment by learning to select more favorable and similar books over time. This learning process is implemented through a reward generalization mechanism, such that readers expect a similar reading reward from semantically similar books and, therefore, select books similar to high-rewarding books and avoid books similar to low-rewarding ones. Moreover, a directed exploration mechanism biases book selection towards options with high estimated uncertainty, signaled by a small number or a high variance of ratings. Our computational modeling results revealed that the UCB model (a directed exploration model accompanied by a reward generalization mechanism), achieved the highest predictability for people's future book selections. Together, people's sequential book selections are governed by reward generalization and directed explorations in a semantic space via reward generalization and uncertainty bias, thus helping readers to address the book value estimation challenge and navigate the exploration-exploitation dilemma.

We probed whether curiosity characterizes people's book selection. We found that people's joyous exploration and thrill-seeking curiosity modulate book exploration behaviors by promoting directed and random exploration tendencies, thus supporting previous theorizing of curiosity as an intrinsic motivational drive that triggers directed exploration[46]. Additionally, we found curiosity fosters enjoyment for selected books, especially for semantically distant explorations. These inflated rewards may reflect the added information bonus specified in the directed exploration mechanism, thus demonstrating that curiosity may encompass both motivational (i.e., wanting) and affective (i.e., feeling) mechanisms. This result is consistent with recent arguments that people's motivational incentives operate independently from the predicted value outcomes[47].

Strikingly, domain-general decision mechanisms applied well to book selection, a domain-specific type of information-seeking behavior in a semantic space. Remarkably, and despite the distinct nature of decision contexts, the mechanisms that govern people's book selection behavior are generally consistent with real-world food-ordering behaviors[17]. These generic mechanisms widely exist for a range of decisions that involve the exploration-exploitation dilemma, such as animal and human foraging behaviors[48, 49] and might be a result of evolutionary adaptations that optimize the exploration-exploitation tradeoff[50]. Our findings showcase additional evidence that these mechanisms extend to abstract semantic space[10, 34, 37] for efficiently searching conceptual items (e.g., books, videos, news, etc.) to optimize information gain and intrinsic reward.

However, challenges and limitations exist in our attempt to apply computational models to real-world book selection data. These challenges and limitations reflect a systematic discrepancy between complex real-world decision contexts and a idealized theoretical decision environment. For instance, decisions are assumed to be made by comparing the estimated value of different options, thus requiring knowledge of all options available, which is usually inaccessible in many real-world decision problems, including book selection. In addition, computational decision models specify a learning process based on experienced rewards that are objective and explicit to decision-makers. However, the experienced reward in real-world decisions is usually implicit, subjective, and easily modulated by the intrinsic value of options, such as interests and curiosity, thus complicating the reward learning and generalization process. In fact, the assumptions in decision theories usually can not be fully satisfied in complex real-world decision situations[51, 52], thus slowing down the application of decision theories on real-world use cases.

The opportunistic applicability of decision theories and models points out promising directions for future research. Empirical studies may explore how to apply domain-general decision theories and models to domain-specific real-world decision problems. It is evident that distinctive human behaviors, such as social interactions, media usage, food foraging, and purchasing behaviors, may share common decision mechanisms. These potential application studies may benefit from increased explainability and predictability[53]. For instance, recommendation systems with a Gaussian Process regression model have shown improved recommendation performance for increased user clicks[54]. Additionally, theoretical works stand to benefit from considering the applicability of decision theory to real-world decision problems. Doing so will help verify the generality of different decision mechanisms to real-world decision environments, increase practical applicability, and generate novel insights for next-generation decision theories.

4 Methods

4.1 The Amazon Dataset

This dataset consists of a representative subset of readers' book selections and ratings on Amazon[32]. It contains 35,478 readers leaving 2,083,630 reading and rating records for 416,797 books. Each reader left 59 ($SD = 40$) book reading and rating records on average. We arrived at this dataset by filtering out readers who left less than 30 records or more than 300 records. Doing so helps us maintain a reasonable horizon length that is long enough to probe learning and exploration dynamics while not too long to demand unaffordable computational expense. Additionally, we filtered out readers for whom more than ten percent of records were placed at the same timestamp, because the true temporal ordering of book selection was missing for these readers. We scrapped the book metadata, including synopsis, rating distributions, and genres from GoodReads, and reading records without corresponding metadata were excluded ($N = 28,303$; 1.3% of all records). All preprocessed data and the code necessary to reproduce the results reported in this manuscript is available online (https://anonymous.4open.science/r/sequential_book_selection-EC8A).

4.1.1 Book Semantic Embedding

We created a latent semantic space for book embeddings. We encoded the preprocessed book synopsis into 384-dimensional embedding vectors for each book, using the state-of-the-art sentence-transformer model `all-MiniLM-L6-v2` (https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2). Thus, the pairwise semantic distance between books was calculated as the Euclidean distance between their semantic embedding vectors.

Distinct from previous studies[17, 22], which measure the frequency of non-repetitive choices as indices for exploration, book choices are non-repetitive in nature. Thus, we measured the extent of exploration, captured by the semantic distance between consecutive choices, as an valid way to quantify people's book exploratory selections.

4.1.2 Embedding Validation

To verify the validity of the semantic distance measures, we asked 248 participants (131 female; M $\pm$ SD age: 40$\pm$13 years), on Prolific to rate the perceived pairwise similarities among 22 randomly sampled books. These 22 books give a total of 231 combinations of book pairs to be evaluated. Each participant was paid $4.47 (equivalent to $12/hour) to evaluate similarities of 15 randomly sampled book pairs after reading the book synopses for both books. Similarity was evaluated on a scale ranging from 1 (extremely dissimilar) to 9 (extremely similar). On average, each pair of books received 16 ratings. Finally, we took the averaged similarity ratings for each book pair to test the validity of the semantic distance measures from the embedding method. We constructed two distance matrices, one with the Euclidean distance metrics in semantic space and the other with participant evaluated similarities. Then, a comparison between these two distance matrices was conducted using the Mantel test[33], which evaluates the association between distance matrices while accounting for the inflated number of observations of pairwise distances.

4.2 The Experimental Dataset

4.2.1 Participants

Participants (n = 250) were recruited from Prolific and paid $6 (equivalent to $12/hour) for their time (31.6 $\pm$ 15.9 minutes) in the experiment. Participants (n = 5) who failed the attention check were excluded from the analysis, thus resulting in a final sample size of 245 (129 female; $M \pm SD$ age: 40 $\pm$ 13 years). The Institutional Review Board at [REDACTED] provided ethical approval of the experimental protocol and the methods were carried out in accordance with the relevant guidelines and regulations. All participants provided informed consent before participating in the experiment.

4.2.2 Stimulus Preparation

Book selection resembled a multi-armed bandit task that simulated the real-world book selection environment. We first selected a subset of 225 books (Supplemental Section

2), which include the 22 books used for semantic embedding validation, from the real-world dataset. Then, we applied a multidimensional scaling technique[55] to project each book's 384-dimensional semantic embedding vectors down into two-dimensional vectors. This dimensionality reduction method maximally preserves the pairwise distances between books from high to low-dimensional space. Next, we arranged these book options into a 15x15 grid based on their two-dimensional embedding vector in a way such that the Euclidean distance on the grid represents the semantic distance between books.

4.2.3 Measures

We measured participant curiosity using the Five-Dimensional Curiosity Scale[31]. This scale consists of 25 survey questions that evaluate five curiosity dimensions: deprivation sensitivity, joyous exploration, stress tolerance, social curiosity, and thrill-seeking. For each dimension, participants are asked to rate five statements on a 0 ("Does not describe me at all") to 6 ("Completely describes me") scale. These subscales' reliability (Cronbach's α) was good ($\alpha > 0.75$) for all five dimensions (Supplemental Table 9). Thus, we averaged the responses for each curiosity dimension and used them as our curiosity measures for the analysis.

4.2.4 Experimental procedure

Once the study began, participants sat at a computer and gave informed consent using a digital form. Next, after a brief training session, participants made a total of 15 selections for their preferred books by clicking one cell on the 15x15 decision grid. After each selection, the corresponding book synopsis was displayed, and participants were asked to evaluate how much they enjoyed the story on a 9-point Likert scale ranging from 1 (extremely dislike) to 9 (extremely like) After the book selection task, participants were redirected to the Qualtics platform to answer questions to measure their trait curiosity[56] and demographics, including age, gender, and race.

4.3 Computational model fitting and evaluation

For the real-world dataset, following[17], we constructed the GP and UCB models with a default parameter setting ($\lambda = 1, \tau = 1, \beta = 1$). For the experimental data, following[12], we used the cross-validated maximum likelihood estimation method to estimate a set of parameters (λ, τ, β) for each subject independently. We used a `scipy`[57] implementation of the global optimization differential evolution method to optimize the likelihood objective function, defined as the sum of the log likelihood for all leave-one-out predictions. Since the differential evolution method is non-deterministic, we repeated the parameter estimation 100 times for each subject, and took the average as the parameter estimates.

Finally, model performance was evaluated based on the predictive accuracy of each model's leave-one-out predictions. We computed a pseudo-R^2 measure, which normalizes the log loss prediction error of model M with that of a random model

M_{rand}, which assumes an uniform distribution of option selection:

$$R^2 = 1 - \frac{log\mathcal{L}(M)}{log\mathcal{L}(M_{rand})} \quad (1)$$

where $R^2 > 0$ indicates a prediction accuracy better than the null model, since $log\mathcal{L}(M) < log\mathcal{L}(M_{rand})$, while $R^2 \leq 0$ indicates a poor predictive accuracy worse than or equal to chance prediction.

References

[1] Fischer, S. R. *A History of Reading* (Reaktion Books, 2004).

[2] Perrin, M. F. & Andrew. A third of americans say they read both print and digital books in the past year (2022). URL https://www.pewresearch.org/short-reads/2022/01/06/three-in-ten-americans-now-read-e-books/ft_22-01-06_bookreading_03-png/.

[3] Cunningham, A. E. & Stanovich, K. E. Early reading acquisition and its relation to reading experience and ability 10 years later. *Developmental Psychology* **33**, 934–945 (1997).

[4] Nicholas, D., Rowlands, I. & Jamali, H. R. E-textbook use, information seeking behaviour and its impact: Case study business and management. *Journal of Information Science* **36**, 263–280 (2010).

[5] Rentfrow, P. J., Goldberg, L. R. & Zilca, R. Listening, watching, and reading: The structure and correlates of entertainment preferences. *Journal of Personality* **79**, 223–258 (2011).

[6] West, R. F., Stanovich, K. E. & Mitchell, H. R. Reading in the real world and its correlates. *Reading Research Quarterly* **28**, 35–50 (1993).

[7] Xuan, Q. *et al.* Modern food foraging patterns: Geography and cuisine choices of restaurant patrons on yelp. *IEEE Transactions on Computational Social Systems* **5**, 508–517 (2018).

[8] Hills, T. T. *et al.* Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences* **19**, 46–54 (2015).

[9] 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).

[10] Lydon-Staley, D. M., Zhou, D., Blevins, A. S., Zurn, P. & Bassett, D. S. Hunters, busybodies and the knowledge network building associated with deprivation curiosity. *Nature Human Behaviour* 1–10 (2020).

[11] Whittle, P. Multi-armed bandits and the gittins index. *Journal of the Royal Statistical Society. Series B (Methodological)* **42**, 143–149 (1980).

[12] Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D. & Meder, B. Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour* **2**, 915–924 (2018).

[13] 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).

[14] Gershman, S. J. Uncertainty and exploration. *Decision* **6**, 277–286 (2019).

[15] Krebs, J. R., Erichsen, J. T., Webber, M. I. & Charnov, E. L. Optimal prey selection in the great tit (parus major). *Animal Behaviour* **25**, 30–38 (1977).

[16] Gong, X. & Huskey, R. Consider the time dimension: theorizing and formalizing sequential media selection. *Human Communication Research* hqad051 (2023).

[17] Schulz, E. *et al.* Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences* **116**, 13903–13908 (2019).

[18] Giron, A. P. *et al.* Developmental changes in exploration resemble stochastic optimization. *Nature Human Behaviour* **7**, 1955–1967 (2023).

[19] Dubois, M. & Hauser, T. U. Value-free random exploration is linked to impulsivity. *Nature Communications* **13**, 4542 (2022).

[20] Blanchard, T. C. & Gershman, S. J. Pure correlates of exploration and exploitation in the human brain. *Cognitive, Affective & Behavioral Neuroscience* **18**, 117–126 (2018).

[21] Hogeveen, J. *et al.* The neurocomputational bases of explore-exploit decision-making. *Neuron* **110**, 1869–1879.e5 (2022).

[22] Riefer, P. S., Prior, R., Blair, N., Pavey, G. & Love, B. C. Coherency-maximizing exploration in the supermarket. *Nature Human Behaviour* **1**, 1–4 (2017).

[23] Sharot, T. & Sunstein, C. R. How people decide what they want to know. *Nature Human Behaviour* **4**, 14–19 (2020).

[24] Murayama, K. A reward-learning framework of knowledge acquisition: An integrated account of curiosity, interest, and intrinsic-extrinsic rewards. *Psychological Review* **129**, 175–198 (2022).

[25] Kidd, C. & Hayden, B. Y. The psychology and neuroscience of curiosity. *Neuron* **88**, 449–460 (2015).

[26] Dubourg, E., Thouzeau, V., de Dampierre, C., Mogoutov, A. & Baumard, N. Exploratory preferences explain the human fascination for imaginary worlds in fictional stories. *Scientific Reports* **13**, 8657 (2023).

[27] Baranes, A., Oudeyer, P.-Y. & Gottlieb, J. Eye movements reveal epistemic curiosity in human observers. *Vision Research* **117**, 81–90 (2015).

[28] Jach, H., DeYoung, C. G. & Smillie, L. Why do people seek information? the role of personality traits and situation perception. *Journal of Experimental Psychology: General* (2021). URL https://psyarxiv.com/tcqe8/.

[29] Risko, E. F., Anderson, N. C., Lanthier, S. & Kingstone, A. Curious eyes: Individual differences in personality predict eye movement behavior in scene-viewing. *Cognition* **122**, 86–90 (2012).

[30] Hsiung, A., Poh, J.-H., Huettel, S. A. & Adcock, R. A. Curiosity evolves as information unfolds. *Proceedings of the National Academy of Sciences* **120**, e2301974120 (2023).

[31] Kashdan, T. B. *et al.* The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *Journal of Research in Personality* **73**, 130–149 (2018).

[32] Ni, J., Li, J. & McAuley, J. Inui, K., Jiang, J., Ng, V. & Wan, X. (eds) *Justifying recommendations using distantly-labeled reviews and fine-grained aspects.* (eds Inui, K., Jiang, J., Ng, V. & Wan, X.) *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 188–197 (Association for Computational Linguistics, Hong Kong, China, 2019). URL https://aclanthology.org/D19-1018.

[33] Mantel, N. The detection of disease clustering and a generalized regression approach. *Cancer Research* **27**, 209–220 (1967).

[34] Hills, T. T., Jones, M. N. & Todd, P. M. Optimal foraging in semantic memory. *Psychological Review* **119**, 431–440 (2012).

[35] Wilke, A. *et al.* A game of hide and seek: Expectations of clumpy resources influence hiding and searching patterns. *PLOS ONE* **10**, e0130976 (2015).

[36] Lawson, R. G. & Jurs, P. C. New index for clustering tendency and its application to chemical problems. *Journal of Chemical Information and Computer Sciences* **30**, 36–41 (1990).

[37] Hills, T. T., Todd, P. M. & Jones, M. N. Foraging in semantic fields: How we search through memory. *Topics in Cognitive Science* **7**, 513–534 (2015).

[38] Reverdy, P., Srivastava, V. & Leonard, N. E. Modeling human decision-making in generalized gaussian multi-armed bandits. *Proceedings of the IEEE* (2019). URL http://arxiv.org/abs/1307.6134. ArXiv:1307.6134 [cs, math, stat].

[39] Wimmer, G. E., Daw, N. D. & Shohamy, D. Generalization of value in reinforcement learning by humans. *The European Journal of Neuroscience* **35**, 1092–1104 (2012).

[40] Cogliati Dezza, I., Cleeremans, A. & Alexander, W. H. Independent and interacting value systems for reward and information in the human brain. *eLife* **11**, e66358 (2022).

[41] Poli, F., O'Reilly, J. X., Mars, R. B. & Hunnius, S. Curiosity and the dynamics of optimal exploration. *Trends in Cognitive Sciences* **28**, 441–453 (2024).

[42] Qiu, J. & Golman, R. Curiosity in news consumption. *Applied Cognitive Psychology* **38**, e4195 (2024).

[43] Schurr, R., Reznik, D., Hillman, H., Bhui, R. & Gershman, S. J. Dynamic computational phenotyping of human cognition. *Nature Human Behaviour* 1–15 (2024).

[44] Gong, X., Huskey, R., Eden, A. & Ulusoy, E. Computationally modeling mood management theory: a drift-diffusion model of people's preferential choice for valence and arousal in media. *Journal of Communication* **73**, 476–493 (2023). URL https://doi.org/10.1093/joc/jqad020.

[45] Noordewier, M. K. & van Dijk, E. Deprivation and discovery motives determine how it feels to be curious. *Current Opinion in Behavioral Sciences* **35**, 71–76 (2020).

[46] Dubey, R. & Griffiths, T. L. Understanding exploration in humans and machines by formalizing the function of curiosity. *Current Opinion in Behavioral Sciences* **35**, 118–124 (2020).

[47] Berridge, K. C. Separating desire from prediction of outcome value. *Trends in Cognitive Sciences* **27**, 932–946 (2023).

[48] Garrett, N. & Daw, N. D. Biased belief updating and suboptimal choice in foraging decisions. *Nature Communications* **11**, 3417 (2020).

[49] Humphries, N. E., Weimerskirch, H., Queiroz, N., Southall, E. J. & Sims, D. W. Foraging success of biological lévy flights recorded in situ. *Proceedings of the National Academy of Sciences* **109**, 7169–7174 (2012).

[50] Pyke, G. H. Optimal foraging theory: A critical review. *Annual Review of Ecology and Systematics* **15**, 523–575 (1984).

[51] Anvari, F., Billinger, S., Analytis, P. P., Franco, V. R. & Marchiori, D. Testing the convergent validity, domain generality, and temporal stability of selected measures of people's tendency to explore. *Nature Communications* **15**, 7721 (2024).

[52] Olschewski, S., Luckman, A., Mason, A., Ludvig, E. A. & Konstantinidis, E. The future of decisions from experience: Connecting real-world decision problems to cognitive processes. *Perspectives on Psychological Science* **19**, 82–102 (2024).

[53] Gong, X. & Huskey, R. Media selection is highly predictable, in principle. *Computational Communication Research* **5**, 1 (2023).

[54] Vanchinathan, H. P., Nikolic, I., De Bona, F. & Krause, A. *Explore-exploit in top-n recommender systems via gaussian processes*, RecSys '14, 225–232 (Association for Computing Machinery, New York, NY, USA, 2014).

[55] Borg, I. & Groenen, P. J. *Constructing MDS Representations*, 19–35 (Springer, New York, NY, 2005). URL https://doi.org/10.1007/0-387-28981-X_2.

[56] Kashdan, T. B., Disabato, D. J., Goodman, F. R. & McKnight, P. E. The five-dimensional curiosity scale revised (5dcr): Briefer subscales while separating overt and covert social curiosity. *Personality and Individual Differences* **157** (2020).

[57] Virtanen, P. *et al.* Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods* **17**, 261–272 (2020).

Declarations

Funding

This project was funded by [REDACTED]. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Competing interests

The authors declare no competing interests.

Ethics approval and consent to participate

The Institutional Review Board at [REDACTED] provided ethical approval of the experimental protocol and the methods were carried out in accordance with the relevant guidelines and regulations. All participants provided informed consent before participating in the experiment.

Consent for publication

All authors consent to publication.

Data availability

The data supporting this study's findings are publicly available. The Amazon book rating data are from the Amazon Review Data (https://cseweb.ucsd.edu/~jmcauley/datasets/amazon_v2/)[32]. The experimental behavioural data are available on GitHub (https://anonymous.4open.science/r/sequential_book_selection-EC8A).

Materials availability

The code necessary to reproduce the experimental paradigm is publicly available on Github (https://anonymous.4open.science/r/sequential_book_selection-EC8A/s2/behavioral_experiment_pavlovia).

Code availability

All custom code required to reproduce the results are available on GitHub (https://anonymous.4open.science/r/sequential_book_selection-EC8An).

Author contribution

Author contribution was redacted for blinded review.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementalinformationblinded.pdf