

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The Effect of Perceived vs. Factual Knowledge on Exploration

Permalink

<https://escholarship.org/uc/item/5f88k3wq>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Fang, Yijin

Yang, Huiying

Christie, Stella

Publication Date

2024

Peer reviewed

The Effect of Perceived vs. Factual Knowledge on Exploration

Yijin Fang¹, Huiying Yang¹, Stella Christie^{1,2}

fangyj22@mails.tsinghua.edu.cn, hy-yang21@mails.tsinghua.edu.cn, christie@tsinghua.edu.cn

¹ Department of Psychological and Cognitive Sciences, Tsinghua University, Beijing, China, 100084

² Tsinghua Laboratory of Brain and Intelligence, Tsinghua University, Beijing, China, 100084

Abstract

Exploration and exploitation decision-making are crucial cognitive processes, often guided by an individual's knowledge state. The information gap theory posits that lesser knowledge enhances exploration, yet the differential impacts of *factual* versus *perceived* knowledge on exploration preferences are not thoroughly understood. This research aims to bridge this gap by independently manipulating factual and perceived knowledge to assess their separate effects on exploration behaviors. Through three studies, we discovered that individuals with less *factual* knowledge explored more intensely and for longer durations, but only when they were explicitly aware of their information gaps. Furthermore, our findings reveal that the *perception* of insufficient knowledge can trigger increased exploration, independent of the factual knowledge possessed. Our studies illuminate the significant impact of metacognitive states on exploration preferences, advancing our understanding of how people decide whether to explore or exploit.

Keywords: exploration; information gap theory; perception of knowledge state; metacognition; explore-exploit decisions.

Introduction

The decision between exploration and exploitation is a ubiquitous aspect of daily life, influencing choices ranging from mundane activities like selecting a meal to significant life decisions such as career planning. This dichotomy revolves around a critical question: should one venture into the unknown in search of potentially better outcomes (exploration), or adhere to the familiarity and safety of known options (exploitation)? Grounded in the information gap theory, research suggests that the drive to explore is often ignited by a need to bridge gaps in understanding (Loewenstein, 1994). Typically, exploration is the instinctive response in new or uncertain situations (Berlyne, 1966; Schulz & Bonawitz, 2007; Kidd & Hayden, 2015; Stahl & Feigenson, 2015). Through exploration, individuals accumulate information and reduce uncertainty enriching their knowledge of the situation. This process of exploration not only facilitates the making of more accurate causal inferences (Cook et al., 2011; Bonawitz et al., 2012; Sim & Xu, 2017) but also orients individuals toward potentially superior choices (Wilson et al., 2014; Blanco & Sloutsky, 2021; Meder et al., 2021). Importantly, many studies have demonstrated that initial exploration influences later decisions, often resulting in a transition from exploration to exploitation as the information gap is bridged (Cohen et al., 2007; Gershman & Niv, 2015; Wu et al., 2018; Schulz et al., 2019; Gopnik, 2020; Wilson et al., 2021).

According to this line of information gap theory, when an individual has more knowledge, they will explore less and are more likely to shift to exploitation earlier. Indeed, recent studies comparing experts and novices support this prediction. Evidence suggests that experts, such as doctors (Reyna & Lloyd, 2006) and policemen (García-Retamero & Dhami, 2009), explored less when addressing highly specialized issues compared to novices. One explanation is that these higher-knowledge groups possess a clear understanding of what is relevant. They rely on fewer dimensions of information, eliminating the need to explore all available information (Ericsson et al., 2007; Gigerenzer & Gaissmaier, 2011).

However, extensive knowledge and information do not always deter experts from further exploration. Experts, as opposed to novices, frequently make significant discoveries through comprehensive investigative efforts. When necessary, they exhibit a remarkably exploratory spirit despite having the smallest relative information gaps. To explain this variable preference in exploration, we propose that the interplay of factual and self-perceived information gaps exerts a complex influence on exploratory behavior. For instance, in contexts of scientific discovery and innovation, experts' exploration may be shaped not only by their actual information gaps, which set them apart from novices, but also by their self-assessment of their expertise in the specific area. Research indicates that self-perception biases can significantly influence self-assessment (Kruger & Dunning, 1999; Dunning, 2011), where less competent individuals may overestimate their capabilities, while the most skilled individuals tend to modest- or even under-estimate theirs. Consequently, if experts perceive remaining information gaps, they may be more inclined to explore further. Nevertheless, distinguishing between the effects of factual versus perceived information gaps remains a challenge in studying the exploration-exploitation dynamics.

Although there are studies that compare perceived and factual knowledge by studying self-assessment and external assessment (Kruger & Dunning, 1999; Davis et al., 2006; Miller & Geraci, 2011), these have typically focused on domains reliant on prior knowledge or experience, such as humor, grammar, and school exams. These domains vary significantly among individuals and are not easily manipulable in an experimental setting. To fully understand the mixed effect of factual and perceived knowledge on exploration and exploitation behaviors, it becomes essential to devise a novel method that allows for the manipulation of both these variables.

In our study, we employed an adapted version of a clicking game designed for searching and collecting rewards (Wu et al., 2017). This game presents a range of options in a search space, where participants can earn rewards based on their choices within a set number of opportunities. They have the discretion to either explore a new option or exploit an already-discovered one. This task framework enables us to observe and analyze participants' exploration preferences, particularly focusing on their exploration frequency, as well as their exploration duration showing in the dynamics of transition from exploration to exploitation.

We hypothesize that *perception* of knowledge may have a more significant impact on exploration preferences than the *factual* knowledge itself. To explore this hypothesis, our initial study (Study 1) assessed differences in exploration behavior among individuals with varying actual knowledge. This was achieved by providing participants with differing amounts of information about the novel game, thereby simulating various knowledge bases. We hypothesized that exploration tendencies might not strictly conform to the information gap theory—those who know less explore more—due to the ambiguity in self-perception about the acquired knowledge. Subsequently, building upon Study 1, in Study 2 we provided individuals with perceptions that matched their factual knowledge by disclosing which level of hints they received. The goal of Study 2 was to investigate whether, after acquiring relatively accurate and matched perceptions, individuals' exploration would better conform to the information gap theory. Study 3 asked whether perception of or factual knowledge has a stronger influence on the decision to explore vs. to exploit. To test this, we employed a novel experimental design, maintaining constant factual knowledge while varying participants' perceived knowledge. This method enabled us to examine if the mere perception of knowledge could independently drive exploration preferences. Through this approach, our studies aim to shed light on how factual and perceived knowledge jointly influence exploratory behaviors.

Study 1

In light of the information gap theory, Study 1 sought to determine if different exploration tendencies are evident among individuals with varied levels of factual knowledge but in completely new situations. Unlike prior experts vs. novices studies where participants are aware of their expertise level in that domain, here we manipulated expertise in situ. Participants were exposed to either minimal (less knowledge) or relatively complete information (more knowledge) about the rules of the novel clicking game. Crucially, participants were not made aware of their relative knowledge levels, effectively creating groups of 'self-unaware experts' and 'self-unaware novices'. We assessed participants' total exploration frequency and the onsets of explore-to-exploit shifts in both conditions. The goal of this design was to assess whether the exploratory behaviors of

these unwittingly knowledgeable or unknowledgeable individuals would still align with the information gap theory.

Method

Participants. 36 university students from across China (23 females, aged 18-25). All participants confirmed they had no prior experience with similar tasks.

Design. Participants were randomly divided into two conditions. The Less-Knowledge group ($n = 17$, 12 females) received no additional contextual information other than the most basic rules and goal of the game. The More-Knowledge group ($n = 19$, 11 females) was provided with details about the task's patterns and rules.

Task. The task involved a clicking game for searching and collecting rewards (Wu et al., 2017). Participants were presented with a blank 8*8 grid search space with a total of 64 tiles. Participants could select any tile to reveal corresponding rewards, limited to a finite number of opportunities. They could decide whether to click a novel tile (exploration choice) or re-click an already-revealed one (exploitation choice). The rewards were represented both in numbers (1-32) and in varying shades of red, with darker shades indicating higher values, to facilitate an intuitive understanding and potentially reduce cognitive load. The search space (the 64 tiles grid) was a 'smooth environment'—where adjacent tiles had similar values. However, knowing this spatial correlation of values required knowledge about various aspects of the search environment: reward range, average value, and spatial distribution (see Figure 1). The Less-Knowledge groups had none of these knowledge, while the More-Knowledge groups were told of these variables. Participants engaged in three rounds of the game, aiming to maximize their total rewards within 25 clicks per round.




Figure 1: The explore-exploit task: a clicking game for searching and collecting rewards. There are 3 knowledge types of the search environment: reward range (Range), average value (Mean), and spatial distribution (Space rule)

Procedure. All participants commenced with the same foundational instructions, explaining that they were expected to search and collect the different rewards hidden behind the 64 tiles of the 8*8 grid world. The initial instruction emphasized two key rules: first, that darker tile colors indicated higher rewards; and second, that any tile could be re-clicked multiple times. After this introductory briefing, the procedure diverged for the two groups. Participants in the

Less-Knowledge condition were not given any further information about the search environment. Participants in the More-Knowledge condition continued to receive a third piece of information about the searching environment, including the range, the mean, and the space rule. The instruction they got was: “(3) thirdly, the values of rewards range from 1 to 32. They exhibit a smoothly gentle fluctuation around 16 and 17. And the distribution of the rewards in such environments is regular and smooth so that closer tiles would likely have more similar rewards.” Finally, all participants were told their goal of the search task: “There are 3 rounds of searching games, with each round offering 25 clicking opportunities. Your goal is to accumulate as many total rewards as possible.”

Results and Discussion

In our analysis of participants’ exploratory behavior, we focused on two key aspects: the total exploration ratio (clicking new versus already-revealed tiles) and the timing of the first explore-to-exploit shift. Inconsistent with the information gap theory, we found no significant disparity in exploration choices between participants with less factual knowledge (41.73% exploration choices, $SD = 0.24$) and those with more knowledge (39.72% exploration choices, $SD = 0.22$). A linear mixed model analysis (exploration ratio ~ factual knowledge base + round + (1|subject), $R^2 = 0.21$, $t(103) = 8.70$, $p < .001$) revealed no significant effect of knowledge on the frequency of exploration, $\beta = 0.01$, 95% CI [-0.05, 0.03], $t(103) = -0.62$, $p = 0.538$. To corroborate exploration preferences, we further looked at the first explore-to-exploit shift point (defined as the first repetitive choice). Participants with less factual knowledge made this shift on the 9.06th trial on average ($SD = 5.66$), while those who possessed more factual knowledge did so on the 9.23th trial on average ($SD = 5.22$). Similar to the total exploration ratio, the effect of factual knowledge bases was also not significant on the explore-to-exploit shift, $\beta = -0.02$, 95% CI [-0.96, 0.92], $t(103) = -0.04$, $p = 0.968$.

In conclusion, our Study 1 findings challenge the conventional assumption of the information gap theory. We observed no significant differences in exploration preferences and performances between participants with less or more factual knowledge. Since the study design creates self-unaware experts and novices, the role of clear self-evaluation and perception regarding one’s knowledge may be critical in explaining the difference between these results and the prediction of information gap theory. Specifically, individuals with more extensive factual knowledge in our study might not have a congruent perception of their knowledge level. To further investigate this possibility, Study 2 was designed to align participants’ perceptions more closely with their factual knowledge, thereby testing if a more accurate self-assessment affects exploration preferences.

Study 2

In Study 1, the absence of notable differences in exploration preferences among participants with varying levels of factual

knowledge led to intriguing questions about the role of knowledge perception. Although Study 1 provided insights into behaviors under different knowledge conditions, it didn’t explicitly address participants’ self-awareness of their knowledge levels. In Study 2, our objective is to elucidate the role of the perception of knowledge in exploratory behavior. Moving beyond prior studies that used self-reported knowledge perception, here we actively manipulated participants’ awareness of their knowledge state before playing the clicking game. Similar to Study 1, participants were provided with varying degrees of information about the search environment. But crucially, they were informed about their current knowledge state. This was done in a non-comparative manner, focusing solely on each participant’s understanding of the completeness of their own information about the task. We aim to examine whether a decrease in exploration occurs as people accurately perceive their diminishing information gap, aligning with the prediction of the information gap theory.

Method

Participants. 38 university students (20 females, aged between 18 and 25 years). Similar to Study 1, all participants confirmed they had no prior experience with the task.

Design. Participants were randomly assigned to one of two conditions: Explicitly-Less-Knowledge ($n = 19$, 16 females), and Explicitly-More-Knowledge condition ($n = 19$, 4 females). For each condition, we emphasized the respective availability or unavailability of additional information, enabling participants to form perceptions aligned with their factual knowledge state.

Task. The task was identical to that used in Study 1.

Procedure. The procedure in Study 2 was largely identical to Study 1, with specific modifications to align participants’ perceptions with their factual knowledge levels. In the Explicitly-Less-Knowledge condition, participants first received the same basic task rules and goals as those in the Less-Knowledge condition of Study 1, without additional information about the search environment. A key distinction was a display of an opaque block masking some hidden information on the screen. Participants were told ‘There is another piece of information related to the search environment that cannot be disclosed at this moment.’ This approach was designed to make participants aware of the existence of further task-related information that they had not obtained, thereby aligning their perception with their actual limited knowledge.

In the Explicitly-More-Knowledge condition, the screen first displayed an opaque block masking the same hidden information. Participants were informed, ‘There is another piece of information related to the search environment additionally shared with you at this moment.’ Then the masking is removed and participants acquired comprehensive information about the search environment, paralleling the More-Knowledge condition in Study 1. This approach ensured that participants developed a perception of sufficiency regarding their current knowledge state.

Results and Discussion

Unlike in Study 1, in Study 2, we observed a difference in exploration behavior. Participants in the Explicitly-Less-Knowledge condition engaged in exploration 53.96% of the time ($SD = 0.25$), significantly higher compared to participants in the Explicitly-More-Knowledge condition, (39.37%, $SD = 0.23$). A linear mixed model analysis (exploration ratio \sim matching perception + round + (1 | subject), $R^2 = 0.42$, $t(109) = 9.72$, $p < .001$) showed a significant increase in exploration frequency in the Explicitly-Less-Knowledge group ($\beta = -0.15$, 95% CI [-0.22, -0.07], $t(109) = -4.05$, $p < .001$). This pattern was also observed in the duration of exploration as indicated by the first explore-to-exploit shift point: the Explicitly-Less-Knowledge group made the first shift from exploration to exploitation on the 11.93th trial on average ($SD = 6.10$) while those in the Explicitly-More-Knowledge condition shifted earlier on the 10.23th trial ($SD = 5.43$). A linear mixed model (shift point \sim matching perception + round + (1 | subject), $R^2 = 0.40$, $t(109) = 8.35$, $p < .001$) confirmed that the 2 groups differed on when they shifted to exploitation ($\beta = -1.70$, 95% CI [-3.39, -0.01], $t(109) = -2.00$, $p = .048$).




Figure 2: Exploratory behavior of participants with more or less knowledge across Studies 1 and 2. Those in Study 1 had no explicit perception of knowledge, whereas those in Study 2 were aware of how much knowledge they knew.

Combining the data from both Study 1 and Study 2, we observed the main effects of explicit knowledge states on both the exploration frequency ($\beta = 0.11$, 95% CI [0.03, 0.19], $t(215) = 2.85$, $p = .005$) and duration ($\beta = 2.58$, 95% CI [0.78, 4.38], $t(215) = 2.83$, $p = .005$). This supported the hypothesis that individuals' perceptions of their factual knowledge bases significantly influence their exploration. Moreover, an interaction effect between matching perception and factual knowledge on exploration frequency was identified ($\beta = -0.12$, 95% CI [-0.22, -7.94e-03], $t(215) = -2.12$, $p = .035$). Specifically, less knowledgeable individuals significantly increased exploration when made aware of their limited knowledge, whereas those with comprehensive knowledge did not alter their exploration behavior based on perception clarity. However, this interaction was not observed in the duration of the exploration.

In summary, the results showed that when individuals' perceptions are manipulated to align better with what they

truly know, their exploration preferences conform more closely to the information gap theory. These findings emphasized the significant impact of perceptions about knowledge on exploration, potentially outweighing the effect of factual knowledge.

Study 3

Across Studies 1 and 2, individuals with lower factual knowledge engaged in more intensive exploration, but only when they were explicitly exposed to their knowledge status. This suggests that the perception of information gaps may have a more substantial impact than the factual information gap itself on exploratory behaviors. To test this hypothesis, Study 3 was designed to isolate the effect of perception. While maintaining a constant level of factual knowledge among all participants, we manipulated their perceptions of knowledge. The primary aim of this study is to test if the mere perception of knowledge changes exploration behavior.

Method

Participants. 105 university students from across China (76 females, aged between 18 and 25 years). All participants confirmed they had no prior experience with the task.

Design. As in Study 2, we manipulated the perception of knowledge by presenting less (masked) or more (eventually unmasked) information. But in contrast to Study 2 where the perception of and factual knowledge are aligned, here, everyone received the *same factual* information, either about the Range, Mean, or Space-rule. We created these 2 (Perception: Perceive More vs. Perceive Less) by 3 (Type of Knowledge: Range vs. Mean vs. Space-rule) to probe whether the effect of knowledge perception is generalizable across different types of knowledge. The detailed participant assignment for each condition is outlined in Table 1.

Task. The task was identical to that used in Studies 1 and 2.

Post Check. After the task, to further confirm the effectiveness of the manipulation of perception, we asked participants to evaluate the knowledge that they received at the beginning of the task on a scale from 1 (not at all sufficient) to 5 (completely sufficient).

Procedure. The procedure was largely identical to Study 2. All participants were assigned randomly to either the Perceived-More or the Perceive-Less condition. All participants received the same factual knowledge (just one variable), but of different types—either Range, Mean, or Space-rule.

Participants in the Perceive-Less condition first received basic instruction on the task rules and goal, along with information on one of the specific knowledge types. This information, however, was followed by some additional information that was obscured by an opaque block, suggesting the presence of undisclosed details. This setup was designed to create a perception of incomplete knowledge. Similar to Study 2, participants were informed, 'There is another piece of information related to the search environment that cannot be disclosed at this moment.'

Participants in the Perceive-More condition also received instruction on the task rules and goal, accompanied by one of the specific knowledge types. Initially, the specific knowledge details were concealed behind an opaque block. Participants were then told: “There is another piece of information related to the search environment additionally shared with you at this moment.” Subsequently, the block was removed to reveal the information. No further information is displayed on the screen, masked or not. This approach framed the same revealed specific knowledge as full and extra, fostering a perception of having sufficient information.

Finally, all participants were asked to finish the post-check scale after the task.

Results and Discussion

To examine the impact of perception and knowledge types on exploration preferences, Linear mixed-effect models were constructed (exploration ratio/first shift point ~ perception * types + round + (1 | subject)). The analysis revealed a significant main effect of perception on exploration frequency ($\beta = -0.19$, 95% CI [-0.34, -0.03], $t(308) = -2.42$, $p = .016$)—Perceived-Less groups explored more. Perceived-Less groups also explored for a longer duration, indicated by differential first explore-to-exploit shift point compared to the Perceived-More group ($\beta = -5.97$, 95% CI [-9.71, -2.23], $t(308) = -3.14$, $p = .002$).

This pattern of exploring more and longer when thinking that one knows less is generalizable across types of knowledge, except for the Space-rule. There is an interaction effect between perceptions and knowledge types also the first explore-to-exploit shift point ($\beta = 1.74$, 95% CI [4.80e-03, 3.48], $t(308) = 1.97$, $p = .049$)—participants in the Space-rule knowledge type tend to explore a lot regardless of their perceived knowledge. Further research is required to fully understand the distinct effects of various knowledge types on exploration behavior. It’s possible that participants who got the space-rule information think of this dimension as difficult or complex, and hence perceive themselves to know less about the game, resulting in high exploration. The detailed results of participants’ exploration preferences in each condition are shown in Table 1 and visualized in Figure 3.

Additionally, in our post-check, the model (self-reported sufficiency ~ perception * types + (1 | subject)) showed a significant difference in individuals’ self-perceived knowledge sufficiency across two kinds of manipulated perceptions ($\beta = 1.48$, 95% CI [0.51, 2.44], $t(102) = 3.04$, $p = .003$), confirming the effectiveness of our perceptual manipulation. Particularly, a significant interaction effect was observed between manipulated perceptions and types of knowledge ($\beta = -0.65$, 95% CI [-1.10, -0.20], $t(102) = -2.88$, $p = .005$). This suggests that participants’ judgments of knowledge sufficiency more closely aligned with our intended manipulations under the Range and Mean knowledge conditions, in contrast to the Space-rule knowledge conditions. The potential explanation of these findings is discussed in the general discussion section.

Table 1: Participants’ assignment and their exploration preferences in each condition of Study 3

Type	Conditions	n _{males} / n _{females}	Exploration	
			frequency M (SD)	duration M (SD)
Range	Perceive-less	5/15	0.57 (0.29)	13.90 (7.37)
	Perceive-more	3/17	0.43 (0.27)	9.92 (6.17)
Mean	Perceive-less	4/13	0.60 (0.27)	13.66 (6.34)
	Perceive-more	2/9	0.40 (0.27)	9.73 (6.32)
Space -rule	Perceive-less	7/11	0.66 (0.29)	15.02 (7.69)
	Perceive-more	8/11	0.59 (0.30)	14.56 (7.22)

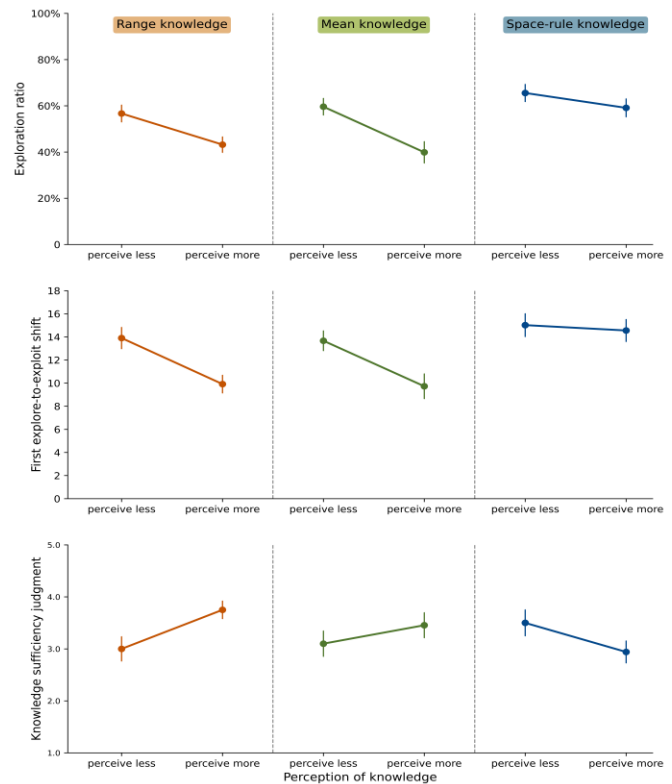


Figure 3: Participants explored more and longer when they perceived themselves as knowing less (Perceive-Less) than when they thought they knew more (Perceive-More) even though everyone had the same factual knowledge.

In addition to exploration behavior, we also examined participants’ total reward performance in three studies. In Study 1, no significant difference was observed in total reward garnered by participants with less or more factual knowledge ($M_{less} = 1952.71$ points, $SD = 182.79$; $M_{more} = 2012.00$, $SD = 147.09$), $t(30.74) = -1.06$, $p = 0.296$, $d = -0.36$. However, a different pattern emerged in Study 2. Participants in the Explicitly-Less-Knowledge condition gathered fewer

rewards ($M = 1805.21$, $SD = 175.74$) than those in the Explicitly-More-Knowledge condition ($M = 2086.11$, $SD = 166.06$), $t(35.89) = -5.06$, $p < .001$, $d = -1.64$). Furthermore, the negative relationship between the trends of exploration and rewards performances was observed again in Study 3: Participants in the Perceive-Less conditions generally gathered slightly but not significantly lower rewards ($M = 1831.27$, $SD = 215.76$) than those in the Perceive-More conditions ($M = 1911.00$, $SD = 218.83$), $F(1, 101) = 3.44$, $p = .066$, $\eta^2 = .03$. Integrating these results, those who thought they knew less increased exploration, but as a result garnered fewer rewards—there is a cost to exploration.

General Discussion

How do individuals decide whether to explore or to exploit? The information gap theory posits that an information gap fuels their drive to explore: people who know less tend to explore more. While prior studies comparing experts and novices corroborated this prediction, real-life phenomena seem to challenge this notion: experts do explore when needed and make novel discoveries, even as they accumulate knowledge. Here we investigate these varying exploratory behaviors by manipulating perceived vs. factual knowledge. Our findings across three studies highlight the pivotal roles of *perceived* information gaps in exploration behavior. Specifically, we observed that higher exploration is initiated only when individuals are explicitly made aware of the existence of information gaps. In scenarios where individuals remained ambiguous about their knowledge state—regardless of its factual sufficiency—there was no significant difference in exploration behavior. Moreover, our studies reveal that even with identical factual knowledge, perception of its sufficiency could significantly influence exploration. Individuals engaged in more exploration when knowledge was presented in a manner that implied the existence of further information gaps. Our three studies help reconcile the information gap theory on exploration-exploitation dynamics, highlighting the significant role of metacognition in the exploration process.

Our method of perception manipulation provides a novel pathway to explore how the metacognitive state, or self-perceived knowledge state, influences exploration preferences. Past research in this domain has yielded inconsistent results regarding the effect of perceived knowledge on the inclination to explore. For example, asking people “Do you think your guess is correct” and “Do you know the correct answer” elicited distinctive trends of exploration (Wade & Kidd, 2019). These inconsistencies may, in part, be attributed to the complexities and biases inherent in self-reporting methods, which do not always directly reflect an individual’s perception of knowledge state. By manipulating participants’ perceptions concerning a novel task, we were able to isolate and more clearly define the influence of perception of knowledge state on exploration behavior.

One potential limitation of our study was the absence of a direct confirmation of the validity of our perception

manipulation in Study 2. In Study 3, we validated our manipulation strategy through a post-check scale assessing participants’ self-judged knowledge sufficiency. The results supported the reliability of our manipulation, showing that participants’ perceptions aligned with our intended control. The interaction between manipulated perceptions and types of knowledge highlights the distinctive role of Space-rule knowledge. Its informativeness allows participants to strategically search, while its complexity introduces multiple uncertainties, such as judging the position of the peak and the direction of reward distribution. These complexities, combined with different prior expectations, can contribute to cognitive dissonance (Festinger, 1962; Harmon-Jones & Mills, 2019) in different directions. Participants in the Perceived-less conditions, starting with lower expectations, might perceive an ‘information bonus’, leading them to adjust their judgments of sufficiency upward when recalling. Conversely, those in the Perceived-more conditions begin with higher expectations, but the unresolved complexities and resulting questions cause a decrease in their perceived knowledge sufficiency.

Future research could examine how various types of knowledge could affect exploratory behaviors. Our findings suggest the type of knowledge presented to participants may play a crucial role in their exploration patterns. Specifically, in conditions where Range and Mean knowledge were presented, individuals consistently exhibited increased exploration when the environment suggested a lack of knowledge. Conversely, in the Space-rule knowledge condition, exploration levels were uniformly high, irrespective of any explicit indication of information gaps (see Figure 3). This could be attributed to the inherently complex and indirect nature of Space-rule knowledge, which likely necessitates more extensive exploration and strategic reasoning compared to the more directly comprehensible Range and Mean knowledge.

Our study also contributes to the discussion of explore-exploit trade-offs, particularly the notion that exploration yields an ‘information bonus’ conducive to better understanding the global environment and future predictions (Wilson et al., 2014). The observed negative correlation between exploration and immediate reward performance in our results aligns with this perspective, indicating that exploration might serve broader objectives beyond immediate reward maximization, such as information gathering and knowledge acquisition. This exploration tendency, despite not correlating with higher rewards in the short term, could prove beneficial in longer-term tasks, as evidenced by the highest rewards being collected under conditions where individuals had the most comprehensive knowledge.

Overall, our research highlights the intricate interplay between perception, prior knowledge, and exploration behavior, opening new avenues for understanding cognitive processes in explore-exploit decision-making contexts.

Acknowledgments

We thank Zi Yin for supporting the construction of the online game platform and Likai Tang for insightful discussions. Both are PhD students from the Tsinghua Laboratory of Brain and Intelligence (THBI).

References

- Berlyne, D. E. (1966). Curiosity and Exploration: Animals spend much of their time seeking stimuli whose significance raises problems for psychology. *Science*, *153*(3731), 25-33.
- Blanco, N. J., & Sloutsky, V. M. (2021). Systematic exploration and uncertainty dominate young children's choices. *Developmental Science*, *24*(2), e13026.
- Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive psychology*, *64*(4), 215-234.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 933-942.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*(3), 341-349.
- Davis, D. A., Mazmanian, P. E., Fordis, M., Van Harrison, R. T. K. E., Thorpe, K. E., & Perrier, L. (2006). Accuracy of physician self-assessment compared with observed measures of competence: a systematic review. *Jama*, *296*(9), 1094-1102.
- Dunning, D. (2011). The Dunning-Kruger effect: On being ignorant of one's own ignorance. In *Advances in experimental social psychology* (Vol. 44, pp. 247-296). Academic Press.
- Ericsson, K. A., Prietula, M. J., & Cokely, E. T. (2007). The making of an expert. *Harvard business review*, *85*(7/8), 114.
- Festinger, L. (1962). Cognitive dissonance. *Scientific American*, *207*(4), 93-106.
- Garcia-Retamero, R., & Dhami, M. K. (2009). Take-the-best in expert-novice decision strategies for residential burglary. *Psychonomic Bulletin & Review*, *16*, 163-169.
- Gershman, S. J., & Niv, Y. (2015). Novelty and inductive generalization in human reinforcement learning. *Topics in cognitive science*, *7*(3), 391-415.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, *62*, 451-482.
- Gopnik, A. (2020). Childhood as a solution to explore-exploit tensions. *Philosophical Transactions of the Royal Society B*, *375*(1803), 20190502.
- Harmon-Jones, E., & Mills, J. (2019). An introduction to cognitive dissonance theory and an overview of current perspectives on the theory. In E. Harmon-Jones (Ed.), *Cognitive dissonance: Reexamining a pivotal theory in psychology*. American Psychological Association.
- Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, *88*(3), 449-460.
- Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of personality and social psychology*, *77*(6), 1121.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, *116*(1), 75.
- Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021). Development of directed and random exploration in children. *Developmental science*, *24*(4), e13095.
- Miller, T. M., & Geraci, L. (2011). Unskilled but aware: reinterpreting overconfidence in low-performing students. *Journal of experimental psychology: learning, memory, and cognition*, *37*(2), 502.
- Reyna, V. F., & Lloyd, F. J. (2006). Physician decision making and cardiac risk: effects of knowledge, risk perception, risk tolerance, and fuzzy processing. *Journal of Experimental Psychology: Applied*, *12*(3), 179.
- Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., & Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, *116*(28), 13903-13908.
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental psychology*, *43*(4), 1045.
- Sim, Z. L., & Xu, F. (2017). Learning higher-order generalizations through free play: Evidence from 2- and 3-year-old children. *Developmental psychology*, *53*(4), 642.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91-94.
- Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic bulletin & review*, *26*, 1377-1387.
- Wilson, R. C., Bonawitz, E., Costa, V. D., & Ebitz, R. B. (2021). Balancing exploration and exploitation with information and randomization. *Current opinion in behavioral sciences*, *38*, 49-56.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore-exploit dilemma. *Journal of Experimental Psychology: General*, *143*(6), 2074.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature human behaviour*, *2*(12), 915-924.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2017). Mapping the unknown: The spatially correlated multi-armed bandit. *bioRxiv*, 106286.