

The Influence of Search Behavior on Perception and Decision-Making

Zur Erlangung des akademischen Grades eines

Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

von der KIT-Fakultät für Wirtschaftswissenschaften

des Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

M.Sc. Thai Quoc Cao

Tag der mündlichen Prüfung: 18.11.2025

Referent: Prof. Dr. Benjamin Scheibehenne

Korreferent: Prof. Dr. Ulrich Ebner-Priemer

Karlsruhe, 2025

Abstract

Human decision-making is typically preceded by information search, a fundamental process that shapes belief and preference formation. Individual differences in search strategies lead to different experiences even within identical choice environments. This complexity is further compounded by noisy information environments, which expose decision-makers to inconsistent or contradictory evidence when evaluating the same options. Moreover, cognitive heterogeneity between individuals or situations may distort perception of identical information, influencing decisions. Therefore, understanding the intricate relationship between search behavior, choice environments, and perception is essential for comprehending human decision-making and preference formation. My thesis investigates the relationship between search behavior and decision-making through three empirical studies. The first study ($N = 130$) re-examines the choice overload phenomenon by examining search patterns across choice sets of varying sizes. The findings demonstrate that search efficiency improves when individuals employ a search breadth strategy, quickly sampling options to develop an initial understanding of the choice set. However, this benefit diminishes as choice set size increases beyond a certain threshold, revealing important boundary conditions for effective search strategies. The second project ($N = 220$) establishes that the sequential order of search significantly enhances the prediction of decision quality compared to non-sequential measures. Predictive accuracy increases by including more temporal data, whereas removing sequential information severely diminishes these models' performance. These findings underscore the necessity of incorporating dynamic search patterns into decision prediction frameworks. The third project ($N = 145$) demonstrates a bidirectional relationship between search behavior and the environmental structure. We found that in skewed distributions, individuals systematically oversample rare values, distorting judgments despite partial awareness of this bias. Collectively, these studies illuminate the dynamic bilateral relationship between search behavior and choice environments. The insights from this research contribute to the development of human-centered adaptive technologies that enhance decision quality while respecting the user's cognitive processes.

Contents

Chapter 1: General Introduction.....	8
Decision Making and Choice Environment	9
The Role of Search Behavior in Exploring the Choice Environment and Making Decisions ..	10
Research Goals and Scientific Method	12
Chapter 2: Re-examining Choice Overload: A Search Behavior Perspective.....	15
Abstract	16
Introduction.....	17
Information Search, Choice Overload, and the Decision-from-Experience Paradigm	18
Information Search Behavior and Set Size	19
Perception of Options' Value in Different Choice Environments	22
Method	24
Study Design	24
Stimuli Creation	25
Procedure and Incentives	25
Results.....	27
The Impact of Search Strategy and Choice Environment on Effort-Adjusted Decision Quality	27
Perception of Values in Large Set Sizes	29
Conclusion and Discussion	30
Theoretical and Practical Implications.....	31

Limitations and Future Directions	32
Chapter 3: The Role of Sequential Search Order in Predicting Decision Quality	34
Abstract	35
Introduction.....	36
Different ways to measure search behaviors.....	37
Predicting decision outcomes based on choice environment and search strategies.	40
Disentangling the Effects of Search Behavior and Information Integration on Probability Perception.....	41
Method	42
Study Design	42
Stimuli Creation	42
Procedure and Incentives	43
Result	43
Effect of Samples on Search Strategies.....	43
Interaction effect of Choice Environments and Strategies on Decision Quality	45
Estimation Accuracy as A Function of Number of Samples.....	47
The Role of Sequential Order of Search Behaviors	49
Conclusion and Discussion	51
Implications and Future Research	52
Limitations	53
Chapter 4: The Effect of Skewed Distributions on Sampling Behavior.....	54

Abstract	55
Introduction.....	56
Biased Sampling Behavior Toward Rare Events	58
Perception of Biased Experienced Means.....	59
Method	61
Study Design	61
Stimuli Creation	62
Procedure and Incentives	63
Results.....	64
Sampling Bias	64
Experienced Mean and Estimation Bias	66
A Potential Explanation for Biased Sampling Behavior.....	68
Conclusion and Discussion	70
Future Research and Implications.....	72
Limitations	74
Chapter 5: General Discussion	76
Future Research and Implications.....	77
Limitations	82
Acknowledgments	84
Reference.....	85
Appendix.....	98

Supplementary Material Chapter 2	98
Supplementary Material Chapter 3	105
Supplementary Material Chapter 4	109

Chapter 1: General Introduction

"Decision makers are not just passive recipients of information, but active seekers who construct their choice environments through their search patterns."

(Simon, 1955)

In an era where information is abundant yet attention is scarce, how individuals search for and process information plays a pivotal role in shaping their perceptions and decisions. Search behaviors—ranging from deliberate, systematic inquiries to quick, heuristic-based scans—fundamentally influence what information people encounter, how they interpret it, and ultimately, the choices they make. This dynamic is particularly relevant in classical decision-making paradigms, where biases in search strategies can lead to skewed judgments or suboptimal outcomes.

This thesis investigates the influence of search behavior on perception and decision-making across three distinct yet interconnected projects. In the first project, we reexamine the choice overload effect through the lens of search behavior, allowing participants to actively explore information and make decisions under varying conditions of set sizes. By recording their search patterns, final choices, and post-decision evaluations, this study explores how different search strategies interact with choice complexity to shape preferences and estimations of value. The second project shifts focus to the predictive power of the sequential order of search behavior in determining decision quality. Using a set of sequential based and non-sequential based metrics, this study examines whether utilizing sequential order of search patterns can improve forecast of decision quality across different environments. Finally, the third project investigates how prior knowledge about the choice environment—such as the distribution of options—shapes search strategies. By examining the bilateral relationship between choice environment and search behavior, we contribute an alternative explanation for well-documented biases of over- and underestimation of value. All three projects employ a naturalistic search paradigm, allowing participants to actively sample and learn about options before making a choice and estimations. This approach captures realistic search behaviors across diverse decision environments while providing insights into how search dynamics interact with classical decision-making phenomena. The overarching goal of this thesis is to demonstrate that search behavior is not merely a peripheral aspect of decision-making but a central driver of perception and choice. By integrating insights from behavioral economics, cognitive

psychology, and information science, this research contributes to traditional models by emphasizing the active role of information search in shaping judgments. The following sections will review key literature on decision-making, assess whether existing methodologies sufficiently account for search behavior, and argue that the interplay between search strategies and choice environments fundamentally influences various stage of judgment and decision-making. Through empirical investigation, my thesis aims to advance our understanding of how search behaviors shape—and sometimes distort—human cognition, offering implications for both theory and real-world decision support systems.

Decision Making and Choice Environment

Decision making is a crucial research topic in psychology and economics, where researchers aim to observe, predict, and explain decisions across various choice environments. Studies often manipulate different parameters of the choice environment and then observe aspects such as decision time (McShane & Böckenholt, 2018; Park & Jang, 2013), choice revision (Polman, 2012), and the final outcomes of the decision (Malone & Lusk, 2019; Reed et al., 2011). The most prominent line of research in behavioral science stems from the work of Amos Tversky and Daniel Kahneman (Tversky & Kahneman, 1979). Using binary gambles and simple description of decision scenarios, they systematically manipulated choice environment's parameters such as probabilities and outcomes to demonstrate consistent violations of rational choice theory, which had been the dominant model in economics. Their explanation—prospect theory—posits that these deviations from rational choice are driven by cognitive biases and psychological preferences. For instance, individuals often prefer a certain but smaller gain over a riskier option with a higher expected value, reflecting risk aversion (Schmidt & Zank, 2008). Likewise, people tend to avoid options that involve a potential loss, even if the expected value of that option is higher, illustrating what is known as loss aversion (Elabed & Carter, 2015; Jing & Cheo, 2013). The decision paradigm that Tversky and Kahneman used, often involving binary or multiple choices, is popular not only in behavioral economics and psychology but also in consumer psychology. Brands often research consumer preferences in different choice environments—for example, under varying set sizes or task complexities—to examine how these factors influence decisions, decision time, and contemplation (Buturak & Evren, 2017; Scheibehenne et al., 2010).

However, we argue that both traditional explanations and research paradigms have overlooked the crucial role of search behavior in decision-making (Hertwig & Pleskac, 2010; Olschewski et al., 2022). Although search behavior is a fundamental component of how people

form judgments and make choices, its role remains insufficiently understood. This thesis aims to highlight three critical aspects of search behavior that have been largely neglected in the existing literature. First, prior research often neglects how search behavior influences perception and decision-making across different choice environments (Le Mens & Denrell, 2011). The way individuals search for information can shape their understanding of the same decision context, leading to different conclusions. Yet, studies that focus solely on choices—without examining how people gather information—miss important insights into this process (Walters et al., 2023). Second, we emphasize that search is an adaptive process, where current search patterns are shaped by insights gained from earlier searches (Hutchinson et al., 2008; Spektor & Wulff, 2023; Wulff et al., 2015). Ignoring this dynamic nature could limit our understanding of how individuals come to different conclusions in the same choice environment. Third, search behavior is not only a tool for learning about the environment; it is also shaped by the structure of the environment itself (Bella-Fernández et al., 2022; Suzuki & Yamamoto, 2021; von Helversen et al., 2018). Different environmental settings can encourage or constrain certain search strategies or introduce various biases, influencing subsequent decisions. By exploring these three dimensions, this thesis aims to generate both practical and theoretical insights. The following sections examine the role of search behavior in greater depth, outline methodological approaches for studying it, and present the overarching goals of this research.

The Role of Search Behavior in Exploring the Choice Environment and Making Decisions

Search behavior plays a crucial role at various stages of decision-making, including information acquisition, information processing, and judgement. One of its primary functions is to reflect the effort individuals invest in gathering information before making a decision. While traditional behavioral research has largely focused on the final choice itself, real-life decisions often hinge on the trade-off between search effort and decision quality—an equally important aspect (Hertwig & Pleskac, 2010; Hills et al., 2010; Mehlhorn et al., 2015). Studies in consumer research have shown that more extensive search can lead to better decisions but may also reduce user satisfaction (Hills et al., 2013; Long et al., 2021; Nobel, 2021). Therefore, it is essential to examine search behavior alongside decision outcomes in behavioral psychology. Moreover, search behavior actively shapes individuals' perceptions of the choice environment by providing new information, whether actively sought or passively encountered. This information influences how people construct their understanding of the available options and, consequently, how they make decisions. Even when the objective characteristics of a choice environment—such as noise and signal—remain

constant, individuals may experience markedly different outcomes depending on the specific information they acquire. This area of research is captured by the decisions from experience paradigm (Hertwig & Pleskac, 2010; Hertwig et al., 2004), which emphasizes how personal experience, rather than described probabilities, guides decision-making. Unlike classical models such as prospect theory, which rely on predefined probabilities and outcomes, the decisions-from-experience approach offers predictions grounded in learning and search behavior.

In this thesis, we adopt the decisions-from-experience framework to examine behavior by simultaneously analyzing the choice environment, how people sample information, and what they observe through their sampling behavior. This approach enables us to understand how individuals interact with and construct their understanding of the choice environment. To support this multi-level analysis, we employ three key concepts throughout this thesis: the *true environment*, the *experienced environment*, and individuals' *estimation of the true environment*. The true environment refers to the objective structure of the decision context—such as the actual probabilities and outcomes in binary gambles or the underlying statistical features of a distribution in a sampling task. However, in search tasks where information is noisy or incomplete, individuals form a representation of the environment based on what they encounter during search—this is the experienced environment. It can be observed and recorded through participants' search patterns and the information they view. Lastly, the estimated environment reflects individuals' cognitive interpretation of the true environment based on what they observed (experienced environment). This estimation involves psychological processes such as reasoning, memory, or heuristics applied to the sampled information in an effort to reconstruct the true environment. Only by studying search behavior and recording what people observe during the search process can we disentangle the contributions of the true, experienced, and estimated environments—making it possible to identify which layer of the decision process drives specific behavioral outcomes.

Another reason we consider search behavior a crucial aspect of decision-making is its dynamic, bidirectional relationship with the choice environment. Search strategies not only shape how individuals perceive their environment, but the structure and features of that environment—such as search costs (R. Wang & Sahin, 2017), number of options (Mejía et al., 2021; Oppewal & Koelemeijer, 2005), and perceived risk (Walters et al., 2023)—also influence how people search. This creates a feedback loop in which strategies and perceptions continuously inform and reinforce each other. Importantly, individuals appear to have some awareness of their own search strategies and can adjust their estimation based on these insights. Research suggests that people exhibit

metacognitive insight into their cognitive and behavioral tendencies; for example, they revise their judgments based on confidence levels or discount risky options when uncertain (Mechera-Ostrovsky et al., 2022; Olschewski & Scheibehenne, 2023; Rosenbaum et al., 2022). A simple illustration of this is when someone estimates the price of an item but adjusts their guess upon recognizing limitations in their memory. This indicates that people can reflect on internal cues and adapt their judgments accordingly. However, the metacognitive dimension of search behavior—whether people are aware of how they search and can reflect on it—remains underexplored. Neglecting this aspect may cause us to overlook important insights into the adaptive mechanisms of human decision-making.

Despite the intuitive appeal of studying search behavior, capturing its complexity remains a major challenge. Traditional metrics—such as the total number of searches, searches per option, or number of switches—provide only coarse measures of effort and its distribution. These metrics often fail to account for the sequential structure of search, which can offer deeper insights into cognitive strategies and environmental constraints. In contrast, fields such as computer science have developed formal algorithmic models—like epsilon-greedy, Upper Confidence Bound (UCB), and uncertainty-driven exploration strategies—that specify how agents might balance exploration and exploitation (Audibert et al., 2009; Castellano, 2019; Kuleshov & Precup, 2014). However, the extent to which these models accurately reflect human behavior remains unclear (Mehlhorn et al., 2015; Vul et al., 2014). Foundational research is still needed to develop theoretical frameworks that capture the nuanced, adaptive, and sequential nature of human search behavior at an algorithmic level. The next section outlines the research goals of this thesis, introduces each specific study, and describes the scientific methods used to investigate them.

Research Goals and Scientific Method

This dissertation has three primary goals, each addressed through a dedicated empirical study. Collectively, the aim is to deepen our understanding of human search behavior in decision-making contexts, particularly how it interacts with and is shaped by features of the choice environment. Central to this investigation is the view of human behavior as part of an adaptive system—one that dynamically adjusts to environmental structure, feedback, and task demands over time. Gaining insight into this adaptive nature not only advances psychological theory but also informs the design of adaptive systems that respond to and support human behavior, enabling more effective human-computer interaction in real-world settings. To achieve this, we employed the scientific method of controlled experimentation, recruiting a large and diverse sample of

participants through the online platform Prolific for each study. This approach ensured a high level of data reliability and generalizability. All three studies were preregistered, including clear hypotheses and predicted outcomes, thus reinforcing methodological transparency and rigor.

Study 1: Revisiting Choice Overload Through the Lens of Search Behavior

The first study revisits the well-known phenomenon of choice overload by focusing on the role of search behavior—an aspect often discussed but rarely integrated systematically into the literature (Scheibehenne et al., 2010; Sethuraman et al., 2022). While larger choice sets may increase the likelihood of finding a desirable option, they also impose greater cognitive and physical search costs (R. Wang & Sahin, 2017). We argue that the effort-reward trade-off central to the choice overload debate can only be fully understood by examining how individuals navigate these trade-offs through their search strategies. To this end, we manipulated the choice environment by varying both the set size and the level of environmental noise. Participants were free to search among options, and we recorded several key variables: the number of searches, the distribution of searches across the choice set, the resulting decision quality, and participants' estimations of the value of their chosen option. We hypothesized that the distribution of search effort would significantly impact search efficiency—participants using more adaptive strategies would maintain high efficiency even in larger sets, challenging conventional findings on choice overload. Conversely, suboptimal strategies would reduce decision quality. Furthermore, we expected participants to demonstrate metacognitive awareness of potential search biases, adjusting their value estimations accordingly.

Study 2: Measuring Sequential Search and Its Role in Decision Quality

Building on the first study, the second project focuses on the measurement of search behavior, particularly the importance of its sequential structure (Cohen & Teodorescu, 2022; Mehlhorn et al., 2015). In the initial study, we developed a basic metric to capture the temporal ordering of search actions and test it predictive of decision quality against multiple alternative metrics. This study extends the literature on search behaviors by systematically comparing different ways to quantify search sequences and examining their predictive power across various decision environments. Participants completed tasks where they could perform a fixed number of searches before making a decision. We analyzed their search strategies using multiple measurements, including those sensitive to sequence order and those that are not (e.g., frequency-based or randomized versions of actual sequences). We then examined how well these different measures

predicted decision performance and value estimation accuracy. We predicted that metrics capturing the sequential order of search behavior would outperform other measures, thereby providing stronger evidence for the cognitive importance of structured search in human decision-making.

Study 3: How Choice Environments Shape Search Behavior

The third study investigates the bidirectional relationship between search behavior and the structure of the choice environment. While prior research has often asked how different search strategies perform across environments (Bartumeus et al., 2005; Everett & Byrne, 2004; Spektor & Wulff, 2023), this study takes the reverse perspective: how does the environment influence the search strategies people adopt. To test this, we provided participants with overviews of environments that were systematically skewed in structure, mimicking real-world biases such as negatively skewed product reviews. Our aim was to understand how decision biases, often attributed to preferences, might actually emerge from the way people search. We hypothesized that skewed environments would elicit biased search patterns, particularly a disproportionate allocation of limited search efforts toward rare outcomes. This biased allocation may result in a distorted experience of the environment, leading individuals to form misrepresentative perceptions based on the atypical information they encounter. However, we also predicted that participants would demonstrate a degree of metacognitive awareness—recognizing the potential distortion caused by their own search behavior—and accordingly adjust their value estimations to more accurately reflect the true structure of the choice environment. By emphasizing search-driven perception, this study introduces a novel framework for understanding how individuals both adapt to and are shaped by the informational characteristics of their environment.

Chapter 2: Re-examining Choice Overload: A Search Behavior Perspective

Thai Quoc Cao¹

Benjamin Scheibehenne¹

¹Karlsruhe Institute of Technology, Kaiserstraße 89-93, 76133 Karlsruhe, Germany

CRedit statement

Thai Quoc Cao: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Software, Validation, Visualization, Writing – Original Draft, Writing – Review & Editing. **Benjamin Scheibehenne:** Conceptualization, Methodology, Supervision, Writing – Review & Editing.

Acknowledgments

This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – GRK2739/1 – Project Nr. 447089431 – Research Training Group: KD²School – Designing Adaptive Systems for Economic Decisions.

Abstract

In this study, we adopt a search behavior perspective to re-examine the choice overload phenomenon, where having too many options can lead to poorer decisions and negative psychological outcomes. A total of 130 participants took part in a search-then-decide experiment, in which they allocated their search effort across varying numbers of options to gather information and make the best possible choice. By analyzing the trade-off between search effort and decision quality using a measure called Effort-Adjusted Decision Quality (EADQ), we found that participants who employed a search breadth strategy (i.e., initially sampling widely before narrowing their focus) achieved higher EADQ scores, even in larger choice sets. However, the benefits of this strategy diminished when the set size reached 40 options. Additionally, while participants showed some awareness of sampling bias in noisy environments, they were largely unaware of and did not adjust for biases associated with larger set sizes. These findings underscore the cognitive challenges of searching in large choice sets and contribute to the literature on choice overload by offering insights into effective search strategies and their limitations. The results also have practical implications for improving consumer search behavior in complex decision environments.

Introduction

Many everyday decisions—such as shopping, investing, or finding a romantic partner online—involve choosing from a wide array of alternatives. While larger choice sets theoretically increase the likelihood of aligning with a decision maker’s preferences, the choice overload hypothesis posits that an abundance of options can produce detrimental effects. These may include increased decision time, reduced decision quality, heightened post-decision regret, and diminished satisfaction (Iyengar & Lepper, 2000; Scheibehenne et al., 2010; Schwartz, 2005). Despite decades of empirical investigation, meta-analysis findings remain inconclusive, with some studies confirming the overload effect and others finding little to no evidence for it (Scheibehenne et al., 2010; Sethuraman et al., 2022). This ongoing debate raises critical questions about the conditions under which choice overload occurs, and why.

Both sides of the debate present valid points. On one hand, expanding choice sets in consumer retails offers clear advantages for firms by attracting a broader customer base and accommodating heterogeneous preferences (Kinjo & Ebina, 2015; Scheibehenne et al., 2010). On the other hand, increasing the number of options also amplifies the cognitive and physical demands of decision-making, as individuals must invest greater effort in searching for and evaluating alternatives. This can lead to lower purchase rates (Long et al., 2021) and increased information-seeking in contexts marked by more uncertainty (Olschewski et al., 2021). The heart of the issue lies in a fundamental trade-off: while larger choice sets offer a greater likelihood of better matches, they also demand more cognitive effort to search and evaluate (Long et al., 2021; R. Wang & Sahin, 2017; Zwick et al., 2003). Some argue that the added search effort is a rational investment in improved decision quality (Saltsman et al., 2021), while others contend that the cognitive costs often outweigh the benefits (Schwartz et al., 2002), particularly when task design or contextual factors make comprehensive search infeasible (F. Wang et al., 2021; Yun & Duff, 2017).

A critical yet often overlooked dimension of this cost-benefit analysis is the role of search strategies in shaping the experience of choice overload. This matters because individuals may adopt different search patterns even within the same decision environment, resulting in divergent subjective experiences and outcomes. Research on both animal foraging and algorithmic search behavior shows that exploration strategies fundamentally shape how agents perceive task complexity and represent their environment (Bartumeus et al., 2005). Furthermore, the effectiveness of any given search strategy depends on how well it matches environmental constraints such as the structure and size of the choice set (Kuleshov & Precup, 2014).

The present study aims to investigate the interplay between search strategies and choice set size, thereby advancing understanding of the mechanisms underlying choice overload. Drawing on insights from the literature on information search, we examine how the same search strategy may lead to different outcomes depending on the size of the choice set, particularly in terms of search effort and decision quality. By observing participants' search behavior and choices, we aim to uncover how strategies influence their mental representations of the decision environment and how these representations guide subsequent adaptations. In doing so, we propose that search behavior in large choice sets may yield patterns and predictions that diverge from those proposed by traditional choice overload theory.

Information Search, Choice Overload, and the Decision-from-Experience Paradigm

A powerful framework for examining information search behavior and its influence on decision-making is the decision-from-experience (DfE) paradigm. Unlike the traditional paradigms in behavioral sciences, in which options are fully described and choices are recorded without observing the search process, the DfE approach allows individuals to learn about options by drawing samples from them before making a decision (Hertwig & Erev, 2009). This setup captures the dynamic and interactive nature of real-life decision-making, such as consumer search behavior, where individuals actively sample and evaluate options based on incomplete information.

Crucially, Decision-from-Experience (DfE) paradigms not only allow researchers to examine final choices but also enable a systematic investigation into how decisions unfold over time. By capturing search processes, decision strategies, and the evolving mental representations individuals form of the choice environment, DfE provides a dynamic window into decision-making. While much of the literature on choice overload has emphasized decision outcomes—such as satisfaction and choice quality (D'Angelo & Toma, 2017; Diehl & Poynor, 2010; Haynes, 2009)—search effort itself is a central component of the decision experience. It contributes to cognitive load, shapes post-decision satisfaction (Bollen et al., 2010), and constitutes one of the most direct behavioral indicators of overload. When search becomes lengthy, complex, or mentally taxing, the overall utility of the decision process—often referred to as the “consumer experience”—may diminish (Bearden et al., 2006).

However, most prior research tends to treat either search effort or decision quality as independent outcome variables, rarely investigating both in tandem with respect to choice set size. To address this gap, we propose a composite measure— Effort-Adjusted Decision Quality

(EADQ)—to serve as the focal point of this study. EADQ is operationalized as the ratio between standardized decision quality and the standardized total number of searches, capturing the average quality gained per unit of search effort. This metric is not only easy to compute, but it also aligns closely with the central trade-off posited by choice overload theory: as choice sets expand, the likelihood of encountering higher-quality options increases, but so too does the cognitive and physical cost of identifying them.

Information Search Behaviors and Set Size

Most studies on choice overload focus on how set size influences decision-making, often in relation to features of the choice environment such as information presentation formats or attribute variability (Besedeš et al., 2015; Dellaert et al., 2017; Townsend & Kahn, 2014). However, relatively few have investigated the compatibility between sampling strategies and different choice environments. Different sampling strategies may perform differently depending on the environment, and conversely, environmental characteristics can shape the strategies individuals adopt (Spektor & Wulff, 2023). Notably, Hills et al. (2013) found that individuals adjust their sampling behavior in response to set size. As the number of options increases, people tend to increase their total number of samples while reducing the number of samples per individual option. This reflects a strategic trade-off between search breadth and search depth. Another, less frequently observed but still notable sampling strategy in large choice sets involves devoting substantial effort to evaluating one option at a time before moving on to the next. Levav et al. (2012) documented both this pattern and search breadth, which has been associated with maximizing and satisficing tendencies in consumer decision-making, respectively.

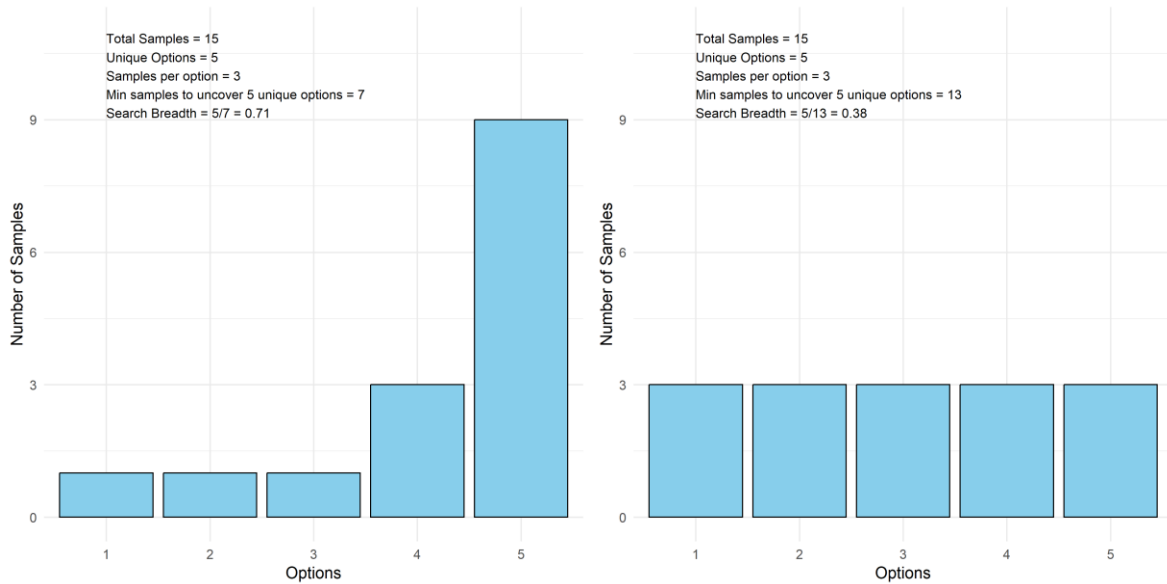
While set size appears to shape the type of sampling strategy individuals adopt, it remains unclear how well these strategies perform when used in the same choice environment. Addressing this question requires comparing both the effort required to reach a decision and the decision outcomes produced by different sampling strategies. To quantify how individuals allocate their sampling effort, Hills et al. (2013) proposed a metric—the ratio of total samples to the number of unique options sampled. This measure captures the average distribution of sampling effort across the choice set and reveals a general shift toward broader, shallower search patterns as set size increases. While this metric captures overall sampling patterns, it overlooks a critical temporal dimension—the sequence in which samples are collected. Sequential patterns can shape early impressions of the option set and influence how options are mentally represented. For instance, consider two individuals who each allocate 15 samples across 5 options. One might sample each

option once before concentrating the remaining samples on two promising candidates—exhibiting a search breadth strategy. Another might allocate 3 samples to each of the 5 options in turn—illustrating a search depth strategy. Although both yield the same average (3 samples per option), the underlying sampling experiences and mental representations of the choice environment may differ significantly (Figure 2.1). To capture this sequential aspect, we introduce a novel metric that reflects the temporal dynamics of sampling strategies. Specifically, we define two distinct sampling styles. The first, search breadth, describes an early emphasis on exploring as many unique options as possible before narrowing focus. In contrast, search depth reflects an early emphasis on thoroughly sampling one option before moving on to the next.

We operationalize these styles using a different sampling ratio: the number of unique options sampled divided by the minimum number of samples required to observe all those unique options. A ratio of one indicates that each sample revealed a new option; lower ratios indicate repeated sampling of the same option(s) early in the search. Figure 2.1 illustrates how these sampling strategies lead to distinct distributions, even when the average number of samples per option remains the same.

Figure 2.1

A demonstration of how different search strategies lead to different ways to distribute the samples even when the number of samples per option is the same.



Note. Left panel: a prioritizing search breadth participant samples each option once then focuses search on options 4 and 5. Right panel: a prioritizing search depth participant samples each option evenly.

Previous research suggests that individuals adaptively prioritize search breadth in large choice sets, and some have considered this behavior as a rational adaptation to increased complexity, allowing individuals to quickly scan the environment and focus on promising alternatives. However, this shift in sampling style may have different implications depending on the structure of the choice environment (Biella & Hütter, 2024; Descamps et al., 2016; Vul et al., 2014). More specifically, the effectiveness of this strategy critically depends on the structure of the choice environment—particularly the level of noise in option values. In environments where option values are noisy, meaning that information about the same options across different samples varied a lot, for example in a uniform distribution, a broader search might initially seem efficient. But because the signal is weak and noisy, individuals may actually need more samples per option (i.e., more depth) to make reliable decisions—leading to higher effort and longer deliberation time. In contrast, in environments with less noise or more central tendency, prioritizing search breadth may be advantageous: individuals can quickly identify good options, verify them with a few samples, and stop searching—resulting in lower effort and good decision quality.

At the same time, an alternative line of reasoning could suggest that search breadth in high-noise environments may actually amplify the perceived differences between options, helping individuals mentally sharpen distinctions and efficiently identify promising alternatives—even with minimal sampling. While this may initially seem counterintuitive, consider a high-noise environment in which options fall into two latent groups: one of generally high-value options and one of generally low-value options. If an individual samples each group, the high-value option might produce an extreme (e.g., near-maximum) outcome, while the low-value option yields a poor (e.g., near-minimum) outcome—purely due to noise. This creates the illusion of a stark difference between the two, making it easier to distinguish the better options early on. Although such lucky contrasts might seem rare, as the number of high-value and low-value noisy options increases with larger set size, so does the probability of encountering extreme observations, giving the search breadth strategy a unique advantage in large, high-noise environments. This idea stands in contrast to traditional choice overload theories, which argue that greater environmental variability complicates information integration and increases cognitive load (Lee & Lee, 2004; Lurie, 2004). By taking into account the interaction between choice environment and search strategies, we want to investigate whether search breadth can actually simplify decision-making—particularly in high-noise environments where random variation can exaggerate meaningful differences across a wide set of options. In such contexts, search breadth strategy may allow individuals to identify promising alternatives more efficiently

H1: Individuals who prioritize search breadth will exhibit higher EADQ in larger set sizes compare to smaller set sizes, but only in high noise environment and not in low noise environment.

Perception of Options' Value in Different Choice Environments

Another crucial aspect of decision-making is how individuals perceive the value of their chosen option. As individuals sample each option, the outcomes they observe accumulate into an *experienced mean*, which they use to estimate the option's *true mean*. Consequently, the way people allocate their sampling effort affects the number of samples collected for each option, shaping their mental representation of those options. This implies that different sampling strategies lead to different levels of accuracy in value representation. For example, a search breadth strategy—where individuals sample broadly across many options but only take a few samples per option—may be sufficient for rough understanding of the choice set, making relative comparisons between options easier. However, it limits accurate evaluation of any individual prospect due to the small number of

samples per option. In contrast, a search depth strategy—involving deeper sampling of fewer options—provides more reliable estimates but requires greater effort to cover the full choice set.

One notable insight from the search breadth strategy is that it can result in positively biased experienced means. Because people tend to interpret positive noise as evidence of high value (and vice versa), limited sampling increases the risk of mistaking noise for signal. Consider a scenario in which an individual samples each of three options once. If all options have the same true mean, the person is likely to choose the one that happened to produce the highest observed value due to random variation and pay no attention to the one that happen to have the lowest observed value. Consequently, this bias might lead them to overestimate the chosen option's true value. Furthermore, as both the number of options (set size) and environmental noise increase, the bias in experienced mean becomes more pronounced—a phenomenon known as search-amplified risk (Hills et al., 2013).

The positive bias described in the three-option sampling example stems from three factors. First, the use of a search breadth strategy that reduces the number of samples per option, making individuals more prone to interpreting noise as signal. Second, the task's objective—choosing the best option—leads individuals to discard options that happen to return a low value after one sample, even if that value was due to noise. Third, the level of noise in the environment itself plays a critical role. In high-noise environments, where information about an option is inconsistent, small sample bias is substantially larger than in low-noise environment. Understanding these factors plays a critical role in how individuals adjust their expectations. A useful method for investigating this is to ask participants to estimate the true mean of an option after sampling. Estimation tasks focus on the cognitive aspect of judgment by instructing participants to provide the most accurate estimate of the option's true value. This contrasts with valuation tasks, where individuals state how much they are willing to pay—responses that are influenced by subjective preferences such as risk aversion. By using estimation rather than valuation, we can more precisely assess the degree to which participants' judgments deviate systematically from their experienced means.

Prior research suggests that asking individuals to internalize statistical principles—like small-sample bias or mean estimation—is too cognitively demanding, making such adjustments unlikely in practice (Juslin et al., 2007; Zhu et al., 2020). Meanwhile, others argued that although people seemingly behave irrationally, they are sometimes aware of the risks involved in their decisions and make corresponding adjustments (Le Mens & Denrell, 2011; Olschewski & Scheibehenne, 2023). Therefore, it is particularly important to examine how people perceive their

experienced mean in a task that feels natural or intuitive, even if it is not framed in explicitly statistical terms. Drawing on research in cognitive adjustment and decisions from experience, we hypothesize that individuals are at least partially aware of how their own sampling strategy interacts with environmental structure. Specifically, we expect that individuals adjust their estimation to be lower than the experienced mean, and the deviation from estimation and experienced mean will be larger in high-noise environment and when the number of samples on the chosen option is lower

H2: Higher environmental noise, fewer samples on the chosen option, and smaller set size will each independently predict greater underestimation of the experienced mean.

Method

Study Design

Participants completed an online search task designed to examine how people search for information and make decisions. In this task, participants were asked to identify the busiest restaurant from a set of options, where “busiest” was defined as the restaurant with the highest average daily customer count. On each trial, multiple restaurants were displayed on the screen, each corresponding to a hidden distribution of daily customer numbers. Participants could sample the number of customers visiting a restaurant on a given day by clicking the button associated with that restaurant. Each click revealed a single value drawn from the restaurant’s distribution. Participants were free to sample as many times as they wished before selecting the restaurant they believed had the highest average. After making their choice, they were asked to estimate the average number of customers for the selected restaurant. The purpose of developing a cover story about choosing the best restaurant is to make the task easier to understand for participants, instead of directly presenting abstract concepts about distributional properties, which many may not be familiar with.

To test H1 and H2, we used a mixed design with two noise conditions and four set size conditions. The task included two between-subject noise conditions: a high-noise condition and a low-noise condition. In the high-noise condition, the underlying distributions of customer counts had greater variability, whereas in the low-noise condition, the distributions were more consistent and predictable. Additionally, the task featured four within-subject set size conditions: 5, 10, 20, and 40 restaurants. The order of these conditions was not randomized; all participants began with the largest set size and progressed to smaller ones. This design was chosen based on prior literature, which shows that descending set size order can promote search breadth strategy (Levav et al., 2012;

Hills et al., 2013). This was implemented to enhance the salience of search breadth strategies—our main focus—by encouraging participants to prioritize exploring more options, especially early in the task.

Stimuli Creation

To generate the underlying distribution for each restaurant, we first determined the mean value of the distributions. We draw the mean values from two distributions so that one group of restaurants has higher average customer counts (the high-value group) and one with lower averages (the low-value group). The means of the low-value group were drawn from a uniform distribution $U(99, 101)$, and those of the high-value group from $U(109, 111)$. Across all set size conditions, 20% of the restaurants were drawn from the high-value group. This meant there was one good option out of every five options across all set size conditions. This fixed proportion ensured that finding a good restaurant is equally challenging across set sizes, controlling for task difficulty of the within-subject conditions.

Once the mean for each option was generated, its distribution shape was determined by the noise condition. In the high-noise condition, values were drawn from a uniform distribution with a range of ± 10 around the option's mean. In the low-noise condition, values followed a normal distribution with a standard deviation of 2.5, truncated at ± 1 SD from the mean to limit the range. This manipulation ensured that the signals of high- and low-value options overlapped more in the high-noise condition—making them harder to distinguish. Meanwhile, the signals in the low-noise condition were more consistent, making high and low-value options easier to differentiate. Lastly, because the mean values were generated from either the high-value or low-value group, there was a risk that participants could learn the typical value ranges and apply this knowledge across trials. To prevent this, each trial's values were scaled by a random multiplier. This discouraged reliance on prior knowledge or numerical anchors from earlier trials and reduced the potential for cross-trial learning.


Procedure and Incentives

The experiment was implemented using oTree (Chen et al., 2016). We recruited 130 participants ($M_{age} = 42$; $SD = 14.15$; 60% female) via Prolific (Prolific, 2024). Each participant received a base payment of €0.80 and a performance-based bonus ranging from €0.20 to €1.00 ($M = €0.73$, $SD = €0.14$). Participants first read instructions and completed a practice trial to familiarize themselves with the task. They then answered three comprehension check questions; only those

who correctly answered at least two were included in the analysis. After completing the main task, participants were asked for demographic information. Performance bonuses were based on the true underlying mean of the selected restaurant, incentivizing accurate identification of the busiest option. After making their selection, participants estimated the average customer count of their chosen restaurant. Accurate estimates (within 10% of the true mean) were rewarded with a fixed additional bonus, further encouraging careful engagement with the task.

Figure 2.2

A snapshot of a trial in the experiment.



City 4: There are 5 restaurants
 20% of restaurants in city 4 are busier than the rest.

Click on any restaurant button to observe how many customers visited that restaurant in one random day. You can do this until you feel confident enough to invest.

Restaurant 1
Restaurant 2
Restaurant 3
Restaurant 4
Restaurant 5

423 customers visited restaurant 1

Choose a restaurant to invest in:

..... ▼

On average, how many customers do you expect will visit your restaurant?

Click **Next** to submit your investment. You cannot change your choice after this.

Next

Note. Participants engaged in sampling a restaurant by clicking on the corresponding button, subsequently making a single decision and providing an estimation of the average number of customers expected to visit their selected restaurant.

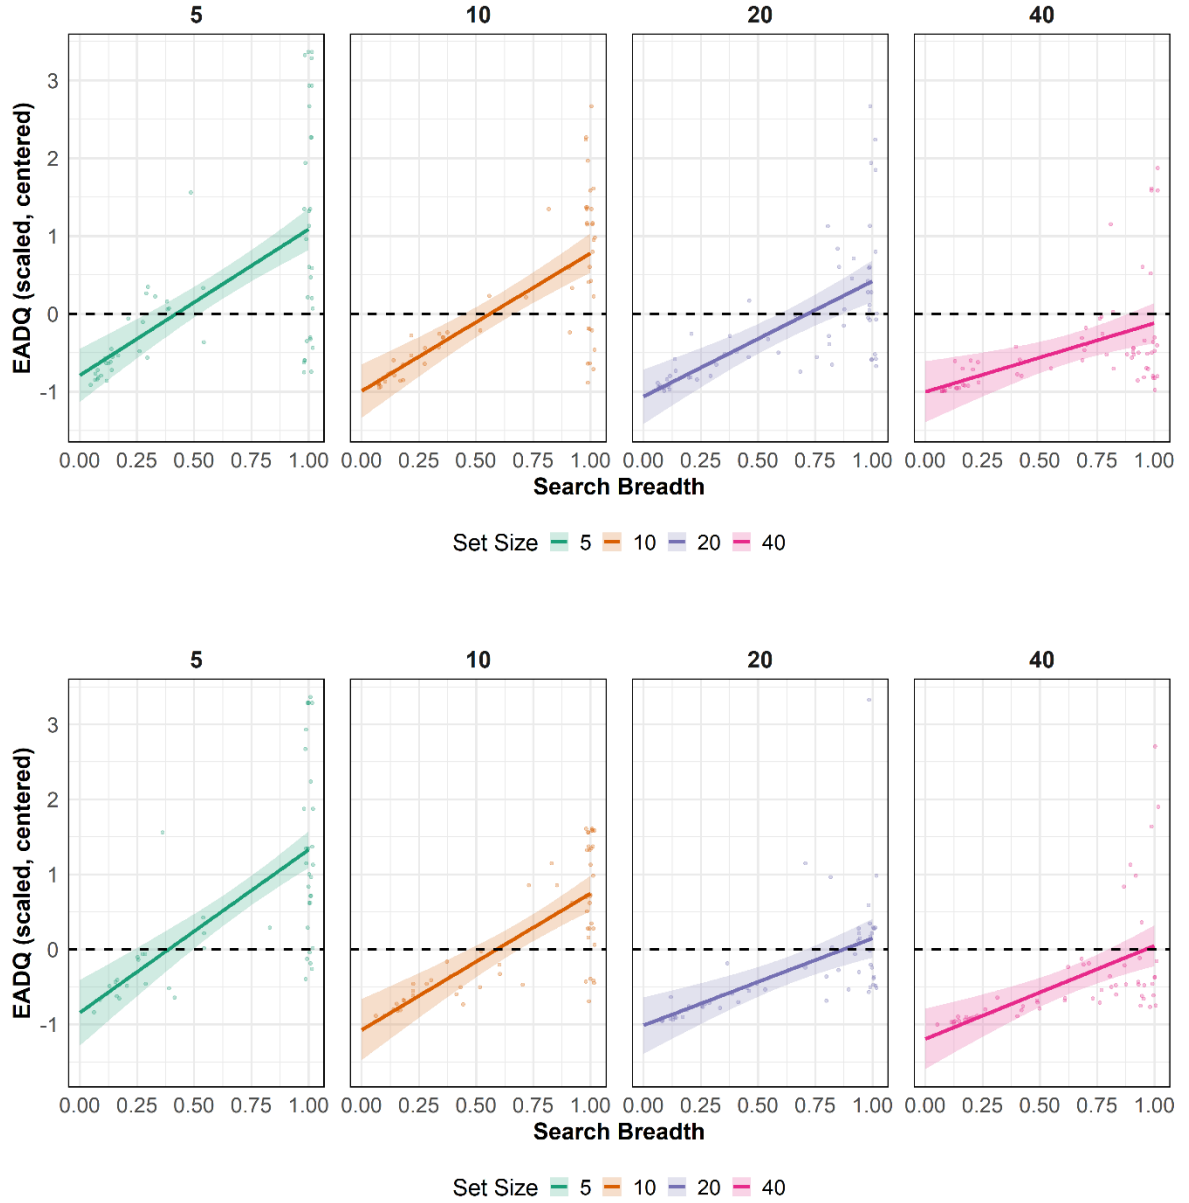
Results

The Impact of Search Strategy and Choice Environment on Effort-Adjusted Decision Quality

Our analyses were performed using the "brms" function (Bürkner, 2021) in R (R Core Team, 2016) to construct Bayesian regression models for confirmatory analyses. The focus of the investigation on this section was on EADQ, defined as the scaled ratio between the underlying mean of the chosen option (decision quality) and the total number of samples participants spent on that trial. To assess the first hypothesis, we modeled EADQ as a function of the interaction between search breadth and set size. Given our expectation that this interaction effect would be substantial only under high-noise conditions, we conducted two separate analyses: one for the high-noise data subset and another for the low-noise subset. The model from the high-noise condition showed an intercept estimate of -0.66 ($SE = 0.17$, $95\% CI [-1.00, -0.33]$), which is below zero. This suggests that at the reference levels of the predictors—that is, when people use search depth in the five options condition—participants' EADQ was substantially below the overall mean. We found a strong main effect of search strategy on EADQ, such that greater search breadth was associated with higher EADQ on average, $b = 1.88$, $SE = 0.24$, $95\% CI [1.40, 2.36]$. This indicates that participants who distributed their attention more broadly during search tended to make more efficient decisions relative to their effort in high-noise environment. Meanwhile, the number of available options did not exert any evident main effect on EADQ at any level (10 options: $b = -0.20$, $SE = 0.25$, $95\% CI [-0.70, 0.28]$; 20 options: $b = -0.27$, $SE = 0.25$, $95\% CI [-0.77, 0.22]$; 40 options: $b = -0.21$, $SE = 0.27$, $95\% CI [-0.73, 0.30]$). Interestingly, a significant interaction emerged between search strategy and the largest set size (40 options), $b = -1.00$, $SE = 0.36$, $95\% CI [-1.71, -0.29]$. This interaction contradicted our prediction that larger set size will emphasize the positive effect of search breadth strategy. However, it seems that the beneficial effect of a broader search strategy on EADQ was attenuated in contexts where participants faced 40 options. In other words, although broader search generally enhanced decision efficiency, this advantage diminished when participants were confronted with a high number of options, possibly reflecting cognitive overload or diminishing returns in complex environments.

Figure 2.3

Effect of search breadth and set size on EADQ in high-noise (Upper) and low-noise (Lower) environments



We conducted a similar analysis in the low-noise environment to assess the effect of search strategy and set size in a different environment. The intercept was significantly negative, $b = -0.85$, $SE = 0.22$, 95% $CI [-1.28, -0.41]$, suggesting that when search depth is used at five options condition, participants' EADQ was at below average level. There was a substantial positive main

effect of search breadth on decision efficiency, $b = 2.17$, $SE = 0.27$, $95\% CI [1.63, 2.72]$, indicating that participants who engaged in search breadth strategy tended to make better EADQ. Although we did not find any main effect of set sizes on EADQ, the interaction effects between search breadth and set size revealed a more nuanced pattern. At set size 20, the positive effect of search breadth on EADQ was substantially reduced, $b = -1.01$, $SE = 0.38$, $95\% CI [-1.76, -0.27]$. A similar attenuation was observed at set size 40, $b = -0.93$, $SE = 0.40$, $95\% CI [-1.71, -0.15]$. These findings suggest that the efficiency benefits of search breadth diminish in larger set size, even when the level of noise in the environment was low.

Taken together, these results indicate that employing a search breadth strategy generally improved EAQD, consistent with our initial predictions. However, in the high-noise condition, this benefit was not amplified as set size increased; rather, EADQ was reduced when participants applied a search breadth strategy. One plausible explanation for this unexpected pattern is that in large set sizes, participants must distribute their limited search effort across too many alternatives, which may increase the likelihood of selecting suboptimal options. Although the relative proportion of good to bad options remained constant, the absolute number of poor options increased, possibly diluting the effectiveness of search breadth strategies. Notably, a similar pattern was observed in the low-noise condition, where the attenuation of search breadth benefits emerged already at the 20-option level. These findings suggest that the advantage of exploratory behavior on decision efficiency diminishes as the set size increases, even while the level of noise was greatly reduced.

Perception of Values in Large Set Sizes

We hypothesize that three factors independently predict the degree of underestimation of the experienced mean: environmental noise, number of samples on the chosen option, and set size. Specifically, we expect that underestimation will be greater in high-noise environments compared to low-noise ones, when the number of samples on the chosen option is lower, and when the set size is smaller. We test this hypothesis using a Bayesian linear regression model with three independent predictors (noise level, number of samples on the chosen option, and set size), each compared to the model intercept. The dependent variable is estimation adjustment, which was defined as the difference between the estimated mean and the experienced mean. A negative estimate of this difference both indicates underestimation of the experienced mean and that the estimation is closer to the true mean than the experienced mean.

We expected participants to adjust their estimates more when the number of samples for the chosen option was low, the set size was larger, and the environment was noisy. The intercept was significantly negative ($b = -2.31$, $SE = 0.56$, $95\% CI [-3.40, -1.21]$), indicating that, when all predictors were at their reference levels (average number of samples on the chosen option, a set size of five options, and a high-noise environment), participants' adjusted estimates were, on average, 2.31 units closer to the true mean compared to the experienced mean. Main effect analysis revealed a significant positive effect of the number of samples on adjusted estimates ($b = 0.04$, $SE = 0.01$, $95\% CI [0.02, 0.06]$), suggesting that increasing the number of samples on the chosen option reduced the need for estimation adjustment. As the number of samples on the chosen option increased, the experienced mean approached the true mean (Appendix Chapter 2), rendering further adjustment unnecessary, thus supporting our hypothesis.

Additionally, the effect of noise was significant. Specifically, adjusted estimates were substantially higher in low-noise environments compared to high-noise environments ($b = 1.17$, $SE = 0.50$, $95\% CI [0.18, 2.15]$), indicating that participants accounted for noise when adjusting their estimates. The effects of set size (10, 20, and 40 items compared to the reference of 5 options) were not evident, as the 95% credible intervals for these predictors included zero. This suggests that participants did not adjust for the bias introduced by larger set sizes in the experienced mean, possibly due to a lack of awareness of this bias. Overall, the results indicate that participants made more accurate adjustments to their estimates when fewer samples were chosen, in noisy environments, and with smaller set sizes. However, they did not appear to adjust for the bias introduced by larger set sizes in the experienced mean, likely due to an inability to recognize this bias.

Conclusion and Discussion

The present study investigated how search strategies interact with choice set size and environmental noise to influence decision efficiency, operationalized through the Effort-Adjusted Decision Quality (EADQ) metric, and the perception of option values. Our findings provide nuanced insights into the dynamics of search and selection, challenging traditional accounts of choice overload and offering a more context-sensitive view of adaptive decision-making. We found that to some extent, larger set size does not harm search efficiency under the search breadth strategy, suggesting the role of search strategy in minimizing effort and maximizing decision quality, and that people's search behavior is influenced by strict search cost constraints (Hertwig & Pleskac, 2010; Vul et al., 2014). However, the benefit of search breadth diminished in the largest set

sizes, contrary to our initial hypothesis (H1). Specifically, participants' EADQ scores plateaued or declined when the number of options reached 40 in both high and low-noise environments, suggesting that searching in large choice environments require too much effort for the given potential benefit in our task regardless of the noise levels in the environment.

Supporting our second hypothesis (H2), participants systematically underestimated the experienced mean of the chosen option when the number of samples from that option was low, and this underestimation was even stronger under high-noise conditions. This pattern suggests a partial, albeit limited, awareness of small-sample bias, echoing prior research on statistical intuitions in experience-based decisions (Hertwig & Erev, 2009). These findings also support previous work on people's metacognitive abilities (Desender et al., 2018; Le Mens & Denrell, 2011; Olschewski & Scheibehenne, 2023). However, relatively little research has explored human metacognition regarding their own search strategies and the consequences of those strategies. Participants adjusted their estimates most when searching in a high-noise environment, suggesting a clear understanding of how their search behavior and the environment interact. While prior research has emphasized the role of opportunity costs and counterfactual comparisons in post-choice dissatisfaction (Schwartz, 2004), our findings suggest an additional mechanism: individuals may develop overly high expectations for an option in large set sizes due to a failure to appropriately adjust their beliefs. This, in turn, increases the likelihood of dissatisfaction when those expectations are not met. These results open new avenues for research linking perceived decision quality to satisfaction and regret, potentially integrating theories of phantom alternatives and decision justification (Diehl & Poynor, 2010).

Theoretical and Practical Implications

Our results contribute to the broader literature by integrating ecological perspectives on search—such as foraging theory (Bartumeus et al., 2005)—into the study of the choice overload phenomenon. Specifically, we introduce the EADQ metric, a more integrative measure of decision quality that incorporates both outcomes and effort, addressing calls for more dynamic and ecologically valid assessments of decision-making performance (Appelhoff et al., 2022; Song et al., 2019; Yeung & Summerfield, 2012). This measurement highlights the trade-offs between search costs and decision outcomes, which have been a central focus in recent research (R. Wang & Sahin, 2017). Furthermore, our findings are consistent with previous research on search behaviors, showing that search breadth can simplify decision-making at various set sizes (Baumann et al., 2022; Everett & Byrne, 2004; Spektor & Wulff, 2023). Our measure of search breadth aligns

conceptually with the satisficing-maximizing tendencies often studied in personality research on choice overload. However, our metric captures actual behavior rather than self-reported traits. While personality may influence search strategies, prior research also highlights the influence of contextual factors—such as initially small set sizes—on encouraging deeper search in subsequent trials (Hills et al., 2013; Levav et al., 2012). Our finding that search breadth alone does not improve decision quality in the largest set sizes aligns partially with the choice overload hypothesis. In such conditions, the cognitive effort required to identify a satisfactory option becomes excessive, and the potential benefits of more options fail to outweigh this cost. As a result, larger sets may hinder decision-making by increasing search difficulty without enhancing value (Reutskaja et al., 2022). These insights have practical implications. Designers of choice environments—such as retailers, policymakers, and interface designers—should consider reducing unnecessary options, increasing the informational value of available alternatives, and minimizing noise. Enhancing the accessibility and clarity of information can lessen the cognitive burden of estimation and adjustment, thereby supporting better-informed decisions.

Limitations and Future Directions

This study has several limitations. As noted earlier, a more effective investigation of the trade-off between search effort and set size would require direct manipulation of the quality of options within different set sizes. A key limitation of our design is that the added value of larger sets was relatively small, as the options were too homogeneous. Future research should explore scenarios where larger sets include options that offer substantially better outcomes than smaller sets, akin to rare target detection tasks in the information and attentional search literature (Mazyar et al., 2013; Pleskac et al., 2021; Zilker, 2022). Another important factor that warrants better control is search cost. In our study, participants were allowed to search freely and among up to 40 options until they felt confident in their decision. However, individuals likely differ in patience levels and tolerance for effort, which could lead to variations in their perceived "budget" for search effort. These individual differences may influence both the extent and the strategy of their search behavior (van Opheusden et al., 2023; R. Wang & Sahin, 2017). Future studies should explicitly examine how search is distributed under varying effort constraints to better understand adaptive search regulation. Overall, this study contributes to the search behavior perspective on choice overload by demonstrating that search efficiency hinges on the alignment between user strategy and choice environment. While search breadth tends to improve efficiency, its advantages diminish as the number of available options grows too large. The finding highlights the need for more adaptive

models of decision-making that consider both the structural characteristics of the environment and the cognitive limitations of decision-makers. Future research should continue to explore how these elements jointly influence not only decisions but also perceptions, satisfaction, and longer-term outcomes.

Chapter 3: The Role of Sequential Search Order in Predicting Decision Quality

Thai Quoc Cao¹

Benjamin Scheibehenne¹

¹Karlsruhe Institute of Technology, Kaiserstraße 89-93, 76133 Karlsruhe, Germany

CRedit statement

Thai Quoc Cao: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Software, Validation, Visualization, Writing – Original Draft, Writing – Review & Editing. **Benjamin Scheibehenne:** Conceptualization, Methodology, Supervision, Writing – Review & Editing.

Acknowledgments

This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – GRK2739/1 – Project Nr. 447089431 – Research Training Group: KD²School – Designing Adaptive Systems for Economic Decisions.

Abstract

The goal of this study is to highlight the role of the sequential order of search behavior in predicting decision outcomes. Participants ($N = 220$) took part in a search task, where they were instructed to identify the binary gamble with the highest expected value by sampling outcomes of different gambles. Controlling for the total number of searches, we found that when the number of searches were limited, participants tended to sample broadly across all available options rather than concentrating their efforts on a few, suggesting that people have a preferred strategy under resource constraints. Furthermore, by comparing various metrics designed to capture search strategies, we found that those incorporating the sequential structure of search behavior predicted decision outcomes more effectively than non-sequential measures. Our findings reveal the complexity of search processes and demonstrate that accounting for the sequential dynamics of search can yield deeper insights into how individuals navigate decision environments.

Introduction

Making informed choices in today's information-rich world is rarely straightforward, as individuals must navigate not only their own preferences but also shifting and sometimes contradictory features of the decision environment. Furthermore, both the choice environment and the individuals themselves are dynamic and continually influence one another. As people have different approaches to search, they encounter different pieces of information at different times, and the same environment may sometimes provide inconsistent or even conflicting information about a given option. As a result, a seemingly straightforward process such as gathering information to form an opinion or make a choice can become unexpectedly complex.

One way to study the intricacies of information search behavior and its consequences for decision-making is by simplifying the components—either by controlling the structure of the choice environment or by constraining how individuals search through it. Prospect theory, one of the most influential theories in decision science, adopts the former approach: participants are typically presented with two binary options in a descriptive format (e.g., a gamble with a certain probability of winning a specified amount), effectively eliminating the need for information search and focusing instead on how people respond to fixed parameters of the choice environment. In contrast, research in consumer science often emphasizes more naturalistic settings, allowing participants to search freely and observing how different search patterns shape their perception and evaluation of the same information. Other studies examine the effects of varying information presentation formats—such as verbal versus visual displays (Townsend & Kahn, 2014)—on judgment and decision outcomes.

While these simplified and controlled approaches have yielded critical insights into how people search and make decisions, they often do so at the expense of overlooking important aspects of search behavior itself—particularly how search unfolds over time. One key factor we highlight in this study is the sequential order of information search. This dimension is crucial because in real-world contexts, information is often dynamic, incomplete, or inconsistent. Two individuals seeking answers to the same question may arrive at completely different conclusions simply due to the sequence in which they encountered relevant information. Moreover, individuals also differ in how they approach the search process: their strategies, priorities, and resources can all significantly impact the quality and success of their search process and decisions. Folk wisdom and psychological research alike suggest that both first impressions (Biella & Hütter, 2024; Kadwe et al., 2022; Ma et al., 2020; Soll et al., 2019) and recency effects (Appelhoff et al., 2022; Mason et

al., 2024) play a significant role in shaping judgment, implying that search is a temporally dependent process. It is not only where individuals search, but also in what order they do so that reveals crucial information about how their decisions are formed.

Therefore, the goal of this research is to systematically incorporate the sequential order of search into a measurement of search behavior, and to use these insights to predict decision quality across choice environments that may either support or challenge different search strategies. We argue that understanding the temporal dynamics of search is essential to fully capturing how people make decisions in a world characterized by information overload, inconsistency, and constant change. In the following sections, we review prior research on how search behaviors have been measured and present our hypotheses concerning the role of sequential search in shaping decision outcomes.

Different ways to measure search behaviors

The literature on information search and foraging behavior has proposed multiple ways to measure and describe how humans and animals allocate limited resources or effort across multiple available options. In the domain of information search, relatively simple metrics—such as the number of searches per option (i.e., how often people search a given option on average) or the percentage of switch versus stay behavior (i.e., whether individuals continue to explore the same option or switch to another)—have been used to assess how search effort is distributed (Mehlhorn et al., 2015). These measures offer insight into whether individuals prefer to stay with a single option to learn more about it or instead explore multiple alternatives. In foraging behavior research, more complex models such as the Lévy flight pattern—where probabilistic models describe the decision to remain in or leave a food patch—have been used to capture how animals and humans search under resource constraints (Bella-Fernández et al., 2022).

However, we argue that both approaches tend to overlook the importance of the sequential order in search behavior. In the first class of measures, although we can compute how often people switch between or stay with options, these summaries ignore the temporal dynamics of decision making. For instance, individuals may begin their search by switching frequently among many alternatives and later focus their attention on a few selected ones. Conversely, if the cost of switching is high, people may prefer to remain with a single option for extended periods. In either case, aggregating behaviors into average statistics—such as number of searches per option or percentage of switch and stay—fails to capture the evolving motivations that underlie search

behaviors at different stages. These dynamics can only be revealed by analyzing the sequence in which search actions occur.

For the second class of descriptive and algorithmic models in foraging and computational research such as the Lévy flight model or uncertainty-based sampling algorithms (e.g., Upper Confidence Bound, UCB), we argue that they provide crucial normative frameworks for interpreting behavior. These models help benchmark human behavior against theoretically optimal strategies (Kuleshov & Precup, 2014). Yet at their core, such models often assume that the next decision is determined by probabilistic rules—for example, with an 80% probability to remain in the current option and 20% to explore a new one. Although such probabilistic strategies may implicitly encode past information, we are still not sure whether people utilize the same signals as these frameworks to conduct their search (Audibert et al., 2009; Loecher, 2021; Speekenbrink, 2022).

Previous research has shown that search is systematic and goal-directed (Hills, 2006; Mehlhorn et al., 2015), serving both to evaluate individual options and to understand the broader structure of the choice environment. For example, repeated sampling of an option near the end of a search process may signal it as the eventual choice (Bhatnagar & Orquin, 2022; Orquin et al., 2021), while early impressions can disproportionately shape subsequent attention allocation (Biella & Hütter, 2024). To understand such strategies, our initial approach was to examine how individuals distribute search effort across available options. Studies by Levav et al. (2012) and Hills et al. (2013) show that when faced with more alternatives, people often choose to sample more options at the cost of reducing the depth of exploration per option. This search breadth first approach appears to be a robust behavioral pattern in both humans and animals even when effort is limited (Bartumeus et al., 2005; Bella-Fernández et al., 2022). In light of this, we reintroduce a metric developed in our first project that captures both the breadth of search and its sequential order of search. This metric measures how quickly an individual samples across all available unique options and is operationalized as the ratio of the number of unique options sampled to the minimum number of samples required to observe those unique options. A ratio of 1 indicates maximal breadth—each sample introduces a new option—while lower ratios reflect earlier repetition and more depth-oriented search patterns.

To understand this measurement, consider, for example, two individuals shopping online for a high-end PC. Both have the opportunity to read a total of 15 short articles about five different models. Person A chooses to begin with one article from each model to get a broad overview, then

returns to the two most interesting options for deeper reading. This represents a breadth-first search strategy. Person B, in contrast, reads three articles about the first model before moving to the next, reflecting a depth-first approach. Although both individuals read the same total number of articles and average the same number of articles per model, the sequence in which they gathered information differs markedly. Our measure of search breadth captures this distinction: Person A would have a search breadth of one, indicating faster exposure to the unique options, while Person B's lower score would indicate slower exposure to new options.

Several important distinctions separate this measure from others. While it may appear similar to metrics such as the average number of searches per option, our measure of search breadth captures how many samples it takes on average to uncover a new option, until all unique options have been discovered. This includes information about the temporal order of sampling that averaging-based metrics lack. Another important note about search breadth metric is that it captures sequential search behavior only up to the point at which all unique options have been observed. Consequently, it does not reflect the sequencing of searches beyond that point. Therefore, in this research, a comparison was made between this metric—which captures sequential order up to a critical threshold—with other approaches in the literature that either incorporate the full sequence of searches or ignore sequential order altogether.

Capturing the temporal dynamics to highlight the difference between search breadth and search depth strategies is not trivial. Although neither strategy is inherently superior, they may create different perceptions of the available options and ultimately lead to different decisions. This contrast becomes especially salient when the number of total searches is fixed, as it allows for a clearer observation of participants' prioritization patterns. By manipulating the total number of searches in our research, we can examine the prevalence of search breadth strategies and examine how they influence subsequent choices. Based on previous research on search behavior, which shows that search breadth is the preferred strategy when the cost of searching between options is low or comparable to the cost of searching within an option—and that this strategy is employed even under limited search effort (Vul et al., 2014)—we expect that reducing the total number of allowed searches will lead to greater reliance on search breadth. This suggests that when more searches are available, people may adopt more diverse strategies, but when the number of searches is limited, they will default to search breadth.

H1: When the total number of allowed searches is reduced, participants will rely more on search breadth as a dominant strategy.

Predicting decision outcomes based on choice environment and search strategies.

An effective way to highlight the impact of search strategies and choice environments on decision quality is to examine how a given strategy performs in different choice environments (Hutchinson et al., 2008; Spektor & Wulff, 2023; Wulff et al., 2015). By observing the choices individuals make and comparing the expected values of their selected options across these different settings, and among people using different search strategies, we can assess the interaction effect between search strategies and choice environments on decision quality.

Building on the insight that search breadth is a strategy that enables individuals to quickly gain an overview of the available options but may come at the cost of a deeper understanding of single option, we can design environments that reward different kinds of information acquisition. Notably, Edwards (1953) developed two types of environments that require learning about either the probabilities or outcomes of different gambles to find the one with the highest expected value. In the first environment, termed "frequent-low" (or "long shots"), options with higher expected value yield frequent but small rewards, while those with lower expected value offer rare but large rewards. Accurately evaluating expected value in this setting requires a deep understanding of the probabilities associated with each outcome, favoring search depth. In contrast, the rare-high (or "short shots") environment features options where the most rewarding outcome occurs infrequently but is of high magnitude, while options with lower expected value yield smaller but more frequent outcomes. This environment tends to favor search breadth because individuals who scan a wide range of options are more likely to encounter one with a rare but highly attractive outcome. When such an outcome is observed, the option may be ranked more favorably due to its high payoff, even if the probability of that outcome is low. This is because individuals with a search breadth strategy are less likely to develop a deep understanding of the probabilities associated with each option compared to those with a search depth strategy. As a result, they tend to prioritize outcomes over likelihood—an approach that aligns well with the demands of the rare-high environment. We hypothesize that the fit between a person's search strategy and the choice environment will predict decision quality: when strategy and environment align, decision quality improves; when they misalign, decision quality declines.

H2: The relationship between search breadth and decision quality will depend on the choice environment. In the frequent-low environment, individuals with higher search breadth will exhibit lower decision quality (negative relationship). In contrast, in the rare-high environment, higher search breadth will be associated with higher decision quality (positive relationship).

Disentangling the Effects of Search Behavior and Information Integration on Probability Perception

Consistent with the first project, the final goal of this study is to examine how individuals perceive their chosen option under varying search strategies and choice environments. In binary gambles, where outcome values are typically visible, the critical unknown that individuals must infer is the probability of each outcome. However, to understand how people perceive the probability of a gamble, several distinctions must be made about three variables: the *true probability*, which reflects the actual likelihood of an outcome; the *experienced probability*, which is based on the frequency of outcomes observed during sampling; and the *estimated probability*, which represents the individual's best estimate of the true probability. Discrepancies between these levels can illuminate why people sometimes make suboptimal judgement. For instance, when individuals sample rationally and develop an unbiased experienced probability but still misestimate the true probability, decision quality likely suffers due to perceptual or cognitive distortion. In contrast, when a search strategy leads to biased sampling and individuals base their probability estimates on this biased experienced probability, poor decisions can be traced back to limitations in the sampling process itself. By distinguishing these sources of error, we can better understand whether suboptimal decisions arise from flawed inference, biased sampling, or a mismatch between the decision-maker and the structure of the environment (Lee & Lee, 2004; Nobel, 2021).

Therefore, it is crucial to recognize that different search strategies can produce fundamentally different experiences and perceptions of the choice environment. Prior research has shown that relying on a small number of samples across many options can result in a positively biased representation of an option's true mean, as people are often drawn to higher outcomes while neglecting lower ones (Erev et al., 2022; Hertwig & Pleskac, 2010; Hills et al., 2013; Noguchi & Hills, 2016). In environments involving multiple binary gambles, a search breadth strategy may reduce the average number of samples per option (Hertwig & Pleskac, 2010), making probability estimates more extreme due to the small-sample effect. As a result, the probability that participants "experience" may be substantially higher than the true probability of the gamble. Such biased experiences can lead to poor decisions if people do not correct for this distortion—i.e., if they treat their biased experience as objective truth. However, based on previous research on metacognitive adjustment (Olschewski & Scheibehenne, 2021), we anticipate that when the number of samples from the chosen option is small, decision-makers will be aware of their own sampling bias and adjust their probability estimates accordingly. Therefore, we hypothesize that individuals will show

stronger corrective adjustments when they have fewer samples from the chosen option, indicating accurate awareness of both the structure of the environment and the consequences of their search strategies.

H3: When the number of samples on the chosen option is smaller, people underestimate the experienced probability more.

Method

Study Design

The hypotheses and experimental design were preregistered on AsPredicted ([link](#)). In this online experiment, all participants were shown ten slot machines, each with a unique expected value and two potential outcomes: winning an amount $[x]$ with probability $[p]$, or receiving nothing with probability $[1-p]$. Participants could click on a slot machine to observe a past outcome from that machine in order to learn about $[x]$, $[p]$, and ultimately the expected value of that option. In each trial, they had a fixed number of samples to explore any of the available slot machines in any order they chose.

We employed a mixed design with four between-subject conditions and two within-subject conditions. Participants were randomly assigned to one of four between-subject conditions, each differing in the number of total samples provided (10, 20, 30, or 40 samples). By controlling the total number of samples, we could assess decision quality under equivalent sampling constraints and examine which strategy is preferred under limited resources (H1), as well as the effect of sample count on search strategies and perception (H3). The two within-subject conditions varied in the choice environments participants encountered: rare-high or frequent-low. Each participant completed four trials in each within-subject condition, and the order of condition blocks was randomized. Comparing search strategies across the same environments allowed us to test the second hypothesis (H2).

Stimuli Creation

To create the two within-subject conditions—frequent-low and rare-high—we first generated an ordered sequence of ten expected values (EVs) ranging from 2 to 8 for the slot machines. In the frequent-low condition, the probability vector $[p]$ consisted of ten ordered values ranging from 0.1 to 0.9. Outcome values $[x]$ were then calculated by dividing each EV by its

corresponding probability ($x = EV/p$). This setup ensured that the option with the highest EV also had the highest winning probability, making its outcome frequent but low in value. In the rare-high condition, we reversed the probability vector used in the frequent-low condition, so that the option with the highest EV was paired with the lowest winning probability. As before, outcome values were computed using $x = EV/p$, resulting in a best option with a rare but high-magnitude outcome.

Procedure and Incentives

The experiment was programmed using Otree (Chen et al., 2016) and distributed on Prolific (Prolific, 2024) to 220 participants. The median completion time was 10 minutes ($SD = 1.13$), and the average participant fee was £7.50/hour. After applying the preregistered exclusion criteria, 179 participants remained, of whom 56% were female and the average age was 40 years ($SD = 13$).

Participants first provided informed consent, then read detailed instructions explaining the task. They were told their goal was to identify the slot machine with the highest expected value by observing past outcomes. An introduction to the concept of expected value was included in the instructions. The incentive rules were explained: one random participant would be rewarded based on the EV of their chosen option in a randomly selected trial. Participants then completed an example trial and answered four comprehension questions about the instructions. Only participants who passed the comprehension check were included in the analysis. Each participant completed four trials per choice environment (frequent-low and rare-high). In each trial, participants were also asked to estimate the $[p]$ of their chosen option by answering the question: "How many times do you expect to win if you get 100 draws from this slot machine?" Comparisons of their estimates with the sampled and true $[p]$ values were used to test H3. After completing the main task, participants answered demographic questions and provided open-ended responses about their decision strategies.

Result

Effect of Samples on Search Strategies

Data analyses were conducted in R studio (R Core Team, 2021) with the brms package for Bayesian mixed-effect model (Bürkner, 2021). All models were tested using random intercepts and fixed slopes. Categorical variables were dummy coded, and variables with expected skewed distribution, such as the number of samples on chosen options, were log-transformed. For our first hypothesis, we conducted a regression analysis using search strategy as the dependent variable and

sample size as the predictor. Our hypothesis posited that search breadth would diminish as the number of samples increased. The descriptive and confirmatory analyses presented in Table 3.1 lent robust support to this hypothesis: in the 10-sample condition, participants predominantly prioritized search breadth, indicating a preference for examining all available options rather than focusing intently on a smaller subset. As the sample size increased, the mean estimate for search breadth decreased, accompanied by an increase in the standard deviation. This pattern suggests a greater variability in strategies employed with larger sample sizes. We conducted an additional model including the choice environment and its interaction with the number of samples, which revealed no significant effects. This suggests that search strategy was primarily influenced by the number of samples, rather than by contextual factors such as the probabilities and outcomes associated with the gambles.

Table 3.1

The effect of the number of samples on search breadth strategies.

Between-subjects Condition	M	SD	Diff	Lower bound	Upper bound
10 samples	0.96	0.11	—	-	-
20 samples	0.85	0.24	-0.12	-0.22	-0.01
30 samples	0.80	0.28	-0.16	-0.25	-0.06
40 samples	0.72	0.32	-0.24	-0.33	-0.14

Note. The "Diff" column shows the mean difference compared to the 10-sample condition, which serves as the reference. The lower and upper bounds represent the 95% credible interval of the difference. Results indicate that higher sample counts are associated with reduced use of search breadth strategies.

Interaction effect of Choice Environments and Strategies on Decision Quality

To test our main hypothesis (H2), we included choice environment conditions, search breadth, and their interaction terms as predictors for the expected value (EV) of the chosen option. We anticipated a positive slope between chosen EV and search breadth in rare-high conditions and a negative slope in frequent-low conditions. Models 3 and 4 of Table 3.2 tested this hypothesis, both incorporating interaction effects between choice environment and strategy. The distinction between these models is that Model 4 included the number of samples as a covariate. If the interaction effect of choice environment and strategy were confounded by the number of samples, statistically controlling for the sample variable could eliminate the interaction effect.

However, interaction effects between strategy and environment were observed in both Model 3 and Model 4, indicating that these effects were not primarily influenced by the number of samples. A negative interaction effect reveals that prioritizing search breadth in a frequent-low condition significantly degrades decision quality (negative slope in frequent-low condition in Figure 3.1). Conversely, a positive, albeit smaller, main effect of search breadth suggests that using this strategy in rare-high conditions enhances decision quality. This effect, not apparent in Model 3, emerged when considering the number of samples in Model 4, indicating that prioritizing search breadth in rare-high conditions is less effective when sample numbers are limited. Thus, our findings supported H2: the interaction effects between choice environment and search strategy significantly predict decision quality.

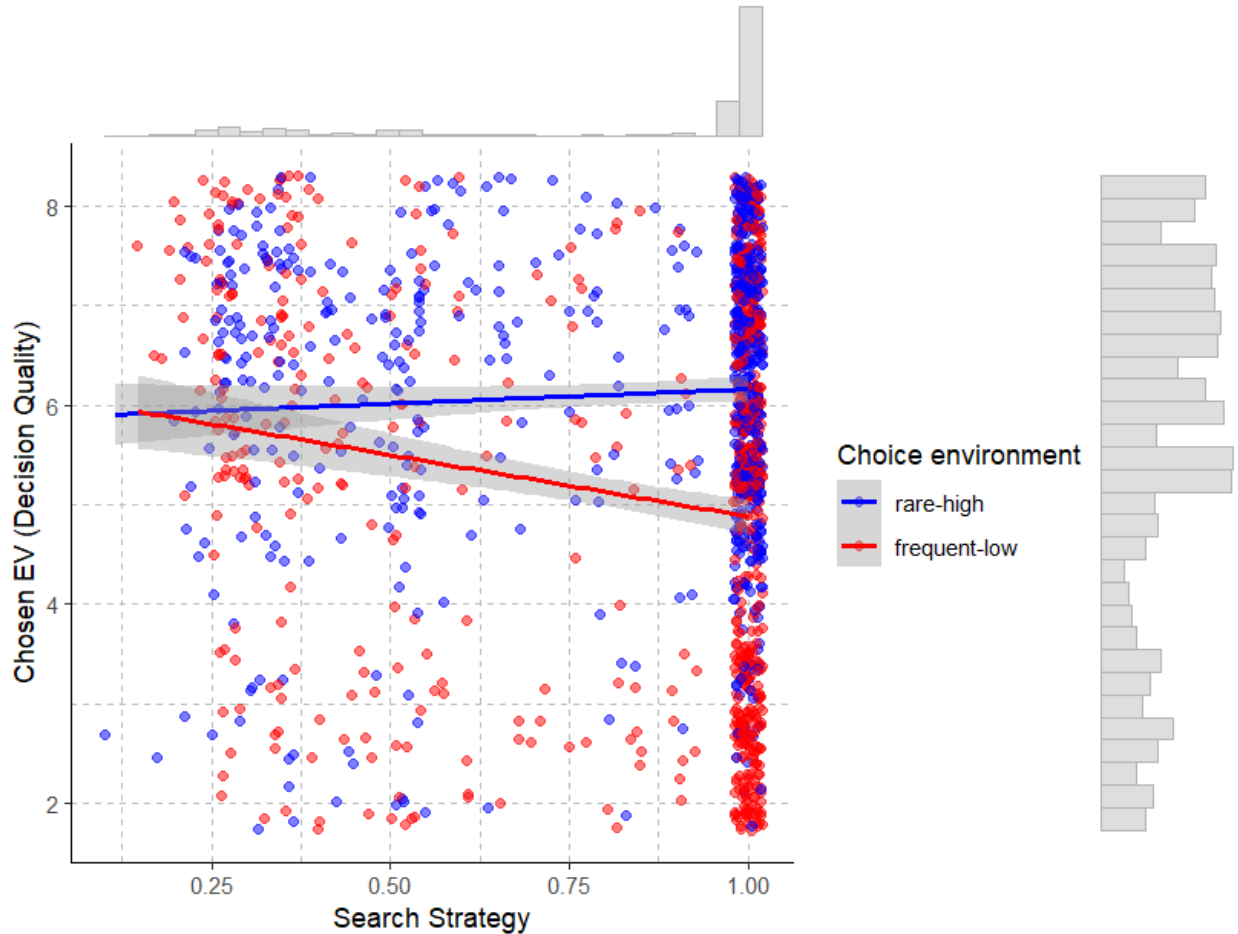
Table 3.2

Regression table. And LOO model comparisons model with the baseline model.

	Model 1	Model 2	Model 3	Model 4
Constant	6.11 [5.97; 6.25]	6.52 [6.19, 6.84]	5.86 [5.44, 6.28]	5.22 [4.72, 5.71]
CO (frequent-low)	-1.01 [-1.18; -0.83]	-1.01 [-1.19, -0.83]	0.25 [-0.30, 0.81]	0.22 [-0.34, 0.78]
Strategy (breadth)	-	-0.50 [-0.84, -0.14]	0.30 [-0.18, 0.78]	0.57 [0.09, 1.06]
Samples 20	-	-	-	0.21 [-0.09, 0.51]
Samples 30	-	-	-	0.48 [0.19, 0.77]
Samples 40	-	-	-	0.76 [0.46, 1.05]
CO x Strategy	-	-	-1.53 [-2.18, -0.90]	-1.49 [-2.13, -0.84]
LOOIC	5640.1	5636.1	5615.8	5597.5
Δ LOO		4	24.2	42.5

Figure 3.1

The effects of search strategy and choice environment (CE) on decision quality.



Note. The data points were jittered to improve visualization.

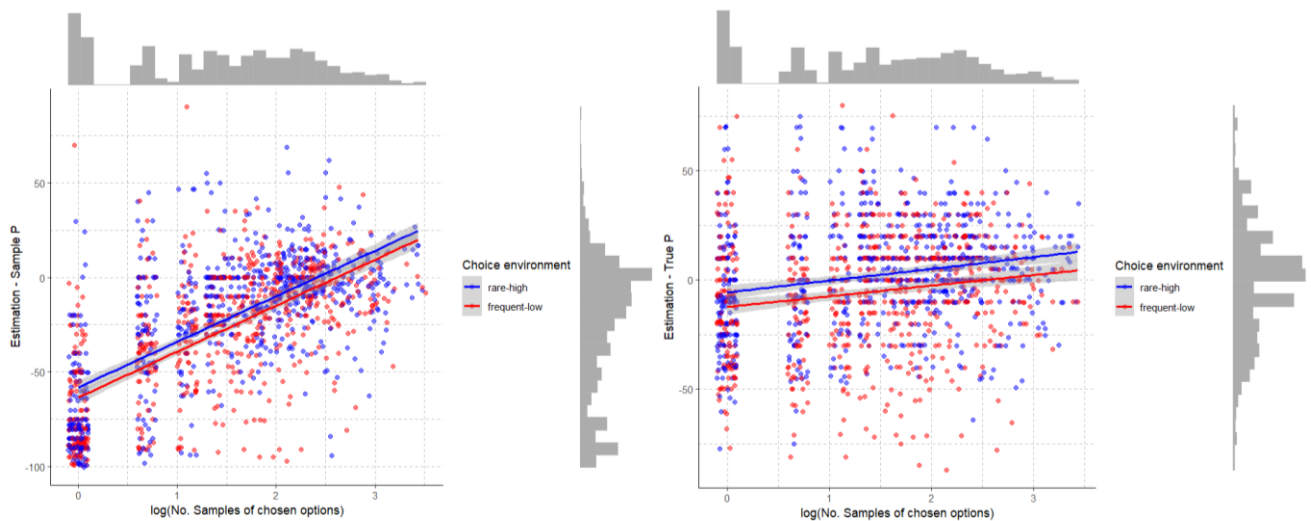
Estimation Accuracy as A Function of Number of Samples

To test H3, we conducted two mixed-effects Bayesian regression models comparing participants' probability estimates to two benchmarks: the experienced probability and the true probability of winning for the chosen gamble. The key predictor variable was the number of samples drawn from the chosen option. We hypothesized that participants would underestimate the experienced probability when the number of samples was low. We further expected this underestimation to occur similarly across both choice environments, suggesting that estimation adjustments are primarily driven by sample size rather than by the environment type.

Figure 3.2 presents the results of these analyses, including effects of choice environment. In the left panel, the dependent variable is the difference between participants' estimated probability and the experienced probability, plotted against the log-transformed number of samples on the chosen option. Most of the slopes fall below the zero point on the y-axis, indicating a general tendency to underestimate the experienced probability. Moreover, the degree of underestimation was greater when the number of samples was small, supporting our hypothesis: participants tended to underestimate the probability more strongly when they had fewer samples. As the number of samples increased, participants' estimates converge more with the experienced probability and eventually slightly overestimated it. Additional analyses revealed no interaction between choice environment and number of samples on the chosen option, ruling out the possibility that the environment conditions had a moderating effect on estimation adjustment.

Figure 3.2

The effect of the number of samples and the choice environment on estimation adjustment and estimation accuracy



Note. Left: The effect of the number of samples and the choice environment on the difference between participants' estimates and the experienced mean (estimation adjustment). Right: The effect of the number of samples on the chosen option and the choice environment on the difference between participants' estimates and the true mean (estimation accuracy). The x-axes in both plots are log-transformed.

The right panel of Figure 3.2 illustrates participants' estimation accuracy, measured as the difference between their estimates and the true probability of the chosen option. Overall, estimates did not deviate substantially from the true probability, even when participants had a low number of samples, suggesting generally accurate estimation and effective adjustment from the sample probability. However, we observed a moderating effect of the choice environment: participants in the frequent-low condition tended to underestimate the true probability more often, whereas those in the rare-high condition tended to overestimate it.

The Role of Sequential Order of Search Behaviors

To examine the role of sequential order in search behavior as a predictor of decision quality, we conducted model comparisons using various sequential and non-sequential search strategy metrics. Specifically, we replaced the search breadth metric in Model 4 of Table 3.2 with alternative strategy measures and compared model fits to evaluate which search measurement best accounted for the data. Among the sequential measures, we first included search inertia, which captures the exploitation-exploration trade-off in sampling behavior (Ashby & Teodorescu, 2019; Mehlhorn et al., 2015). This metric indicates whether participants tend to continue sampling the same option or switch to others. It was computed as the proportion of samples where the option sampled at time $t+1$ differed from the option at time t . A value of one represents constant switching between options, while a value of zero indicates that all samples were drawn consecutively from a single option without any switching. The second sequential metric was the average length of search sequences, adapted from run-length encoding methods in computer science (Rahman, 2025). The simple form of this metric captures the average number of consecutive samples drawn from the same option. For example, the sequence [1, 2, 3, 4, 4] yields run lengths of [1, 1, 1, 2], with an average length of 1.25. The maximum possible value equals the total number of samples, indicating complete exploitation of a single option. The minimum possible value is 1, indicating that no option was sampled twice in a row.

To evaluate the importance of sequential order, we compared these sequential metrics to a commonly used non-sequential measure: entropy. Entropy assesses the unpredictability of a sequence (Qu et al., 2022; Scheibehenne et al., 2010). An entropy value of 1 reflects a perfectly predictable strategy (e.g., repeatedly sampling one option), while lower values indicate higher randomness. However, entropy does not capture the temporal order of samples—it treats each selection independently. To further isolate the role of sequential structure, we removed the sequential order from participants' search behavior data by randomly shuffling their sampling

sequences. We then recalculated two key metrics that are sensitive to order—search breadth and average run length—and refitted them into the same predictive model. A decline in model fit following this manipulation indicated that these metrics rely on sequential information and lose predictive power when that order is disrupted.

Table 3.3

Model comparison of prioritizing search breadth strategy with other measurements of search strategies

Models	LOO	Δ LOOIC	p_loo
Search breadth	5597	0	30.3
Average search length	5590	7	32.1
Search Inertia	5605	-8	31.8
Entropy	5615	-18.2	28.9
Search breadth shuffled	5606	-9	29.6
Average search length shuffled	5618	-20	27.8

Table 3.3 presents the performance of six models evaluated using Leave-One-Out Cross-Validation (LOO-CV), a standard method for comparing Bayesian models. The LOO Information Criterion (LOOIC) provides an estimate of out-of-sample predictive accuracy, with lower values indicating better performance. The Δ LOOIC values show differences from the baseline model—defined here as the model using the search breadth strategy—with higher Δ LOOIC indicating worse predictive performance than baseline model. Model complexity was assessed using the effective number of parameters (p_{loo}), where higher values suggest a more complex model. Among all models, the average search length model achieved the lowest LOOIC (5590), indicating the best predictive accuracy. This metric incorporates the most amount of sequential order data in search behavior. However, it is also the most complex model, with the highest p_{loo} (32.1). The prioritizing search breadth strategy, our main proposed metric, had the second-best predictive accuracy (LOOIC = 5597) and was the second least complex model, ranking just after entropy. Notably, both metrics that incorporated sequential order information—average search length and prioritizing search breadth—outperformed models using search inertia and entropy, which did not account for the order of sampling. In contrast, when the order of search behavior was removed by shuffling participants' sampling sequences, the performance of the average search length and prioritizing search breadth models deteriorated, falling below that of entropy and search inertia. This finding underscores the critical role of sequential order in modeling search strategies and their influence on decision quality.

Conclusion and Discussion

In the complex landscape of decision-making, individuals frequently operate in noisy informational environments, striving to make choices that align with their preferences. This study aimed to better understand how search strategies—particularly those incorporating the sequential order of search—interact with the structure of the choice environment to influence decision outcomes and perceptions. Our findings reveal a discrepancy between how the search process is theorized and how it is typically measured. While prior literature has characterized search as goal-oriented and systematic (Hills et al., 2006; Mehlhorn et al., 2015), many existing measures neglect the sequential order in which search unfolds. This omission can obscure important insights into the nature of human decision-making. By incorporating sequential search order into our analysis, we uncovered that participants often adopted a systematic approach: they initially distributed their search broadly before narrowing in on a subset of options. This behavior was particularly

pronounced when the number of samples approached the number of available options, suggesting that search breadth was prioritized under limited search effort condition. These results reinforce previous findings that highlight the structured nature of human search and support the idea that people prioritize learning available options first instead of searching deeply when navigating complex environments (Levav et al., 2012).

Crucially, we showed that integrating the sequential structure of search into our metrics—particularly through measures like search breadth and average sequence length—yielded better predictions of decision outcomes. These effects were especially clear when comparing two types of environments: rare-high vs. frequent-low. In rare-high environments, broader search led to better outcomes; in frequent-low environments, the same strategy underperformed. The predictive power of these measures diminished and underperformed non-sequential measurements when search sequences were shuffled, suggesting that the observed benefits arise from the actual temporal order of search.

By examining participants' estimations of the probability of the chosen option, we found that participants tended to underestimate the bias observed probabilities, effectively correcting toward the true probability. This suggests a metacognitive awareness of their own strategies: individuals recognized that sparse sampling could yield misleading impressions, and adjusted accordingly. Rather than showing a correct perception of a biased observed probability, participants demonstrated adaptive calibration of beliefs, supporting recent claims about strategic adjustment to their own meta cognitive (Olschewski & Scheibehenne, 2021).

Implications and Future Research

Our study contributes to the decision sciences by emphasizing the importance of sequential order in search behavior. By introducing sequence-sensitive measurements, we move toward a more ecologically valid and cognitively grounded model of human search. Moreover, our findings show that search strategies are contingent, not universally optimal: what works well in one environment may backfire in another. This challenges static notions of rational search and supports a contingency-based approach to modeling behavior, where strategy effectiveness emerges from the interaction of environmental properties and cognitive constraints.

Practically, our results offer several directions for designing decision support systems, recommendation engines, and choice architectures. For instance, interface designers and online retailers might use behavioral indicators of search sequence to infer when users are over-searching

or under-searching, and provide adaptive cues accordingly. Educational tools could also train individuals to recognize when to switch from breadth-focused to depth-focused strategies depending on the distribution of valuable options in an environment. Several promising avenues remain. Future studies should explore how search order interacts with individual traits, such as cognitive reflection, need for cognition, or maximization tendency, to affect performance. Investigating search behavior in real-world high-stakes decisions (e.g., job selection, financial investments) could also help validate and extend the findings. Additionally, efforts to formalize a model that incorporates sequential dependencies—either through machine learning, reinforcement learning, or algorithmic modeling—may help explain how strategies emerge and evolve across domains.

Limitations

Despite these contributions, several limitations should be acknowledged. First, our prioritizing search breadth metric captures only a narrow slice of search behavior. While it reflects how broadly participants search across options, it does not fully represent goal-directedness or the recursive structure of decision-making. Future work should aim to develop more comprehensive models that capture both search structure and intention. Second, our study does not directly incorporate search costs—a key factor in ecological and economic models of decision-making (Hertwig & Pleskac, 2010; Spektor & Wulff, 2023). In natural environments, the costs of sampling between options versus within options can vary dramatically. For example, in foraging contexts, switching between patches may be more costly than exploiting a single resource. These structural differences can lead to very different search patterns, such as Levy flight or stay-switch heuristics, which our paradigm does not capture. Finally, while our results apply to the choice environments we tested, they may not generalize to all forms of sequential search. Prior research has shown that different paradigms—such as foraging, menu selection, or hypothesis testing—can elicit fundamentally different cognitive behaviors (von Helversen et al., 2018). Thus, caution is warranted in applying these findings across all decision domains.

Chapter 4: The Effect of Skewed Distributions on Sampling Behavior

Thai Quoc Cao¹

Benjamin Scheibehenne¹

¹Karlsruhe Institute of Technology, Kaiserstraße 89-93, 76133 Karlsruhe, Germany

CRediT statement

Thai Quoc Cao: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Software, Validation, Visualization, Writing – Original Draft, Writing – Review & Editing. **Benjamin Scheibehenne:** Conceptualization, Methodology, Supervision, Writing – Review & Editing.

Acknowledgments

This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – GRK2739/1 – Project Nr. 447089431 – Research Training Group: KD²School – Designing Adaptive Systems for Economic Decisions.

Abstract

This study examines how biases in information search behavior, particularly the tendency to over-sample rare outcomes, influence the perception and evaluation of numerical distributions. Participants in an online experiment ($N = 145$) disproportionately sampled rare events more in skewed distributions, leading to overestimations of positively skewed distributions and underestimations of negatively skewed ones. While increasing the total number of samples reduced estimation errors, participants also made cognitive adjustments by partially compensating for their oversampling of rare events when estimating means. These findings contribute to a better understanding of the interplay between information search biases and cognitive processes in decision-making, insights that are increasingly relevant as autonomous systems shape how humans gather and interpret information. Such systems, by influencing both the information and how we sample, may affect agency and judgment in various contexts such as e-commerce, social media usage, and even morally complex scenarios such as politics.

Introduction

When making decisions, people often possess prior knowledge of an option's quality or possible outcomes from past experiences. For instance, e-commerce platforms commonly display distributions of past customer ratings as an indicator of a product's quality. Likewise, in many countries, lottery providers are legally required to disclose the probability of winning for a given prize tier. These disclosures are expected to support more informed and rational decisions. However, despite widespread awareness that winning a lottery jackpot is statistically improbable, global spending on lottery tickets amounts to approximately \$250 billion annually (Kim & Oswald, 2021). Similarly, marketing research suggests that rare but highly negative reviews can disproportionately damage a product's reputation, even when positive feedback is abundant (Wu, 2013). These examples suggest that rare outcomes may be subjectively overweighted in these contexts.

The influence of distribution features, such as the rareness of outcomes, on perception and judgments has garnered substantial research attention. Prospect Theory, a prominent explanatory framework for decision making under risk, posits that individuals disproportionately overweight low probabilities, a tendency attributed to the curvature of the probability weighting function (Kahneman & Tversky, 1977). However, Prospect Theory's explanation is largely derived from studies on choice behavior in binary gambles, with limited exploration of how individuals gather and process information—processes that may be shaped by distinct biases (Azzopardi, 2021; Schulz-Hardt et al., 2000). Recent research has increasingly focused on how the characteristics of a distribution influence its perception (Ludvig & Spetch, 2011; Ludvig et al., 2014; Mason et al., 2024; Stewart, 2009). However, further research is required to examine how these characteristics influence sampling behavior, a process that precedes perception (Hills et al., 2010). Without examining sampling behavior, the biases observed in lottery-like gambles—commonly attributed to cognitive factors—may instead reflect residual effects of biased sampling patterns. For example, Niese and Hütter (2022) demonstrated that the negative framing effect, traditionally explained by Prospect Theory through motivational biases such as loss aversion, can also be explained by sampling processes—where negative framing prompts individuals to retrieve more negative information about an option, leading to framing-dependent biases. This underscores the important role of sampling behavior in reexamining previously established psychological phenomena. In other words, sampling biases could shape cognitive processes such as perception (Johnson & Tversky, 1984; Olschewski et al., 2022; Walters et al., 2023), and preferences (Dohmen et al., 2018;

Mallpress et al., 2015; Weber, 2010) which then in turn influence decision outcomes. This suggests a potential causal pathway in which sampling behavior influences perception, ultimately leading to decision biases such as overweighting of rare events. By concurrently analyzing both sampling behavior and information perception, this study seeks to disentangle these components with a special focus on overweighting of rare events.

Evidence from multiple studies indicates that search behavior is sensitive to both external and internal influences. For instance, Biella and Hütter (2024) demonstrated that sampling strategies vary with motivational context: when individuals are driven by interest, they tend to truncate sampling after encountering early counter evidence, whereas in disinterested contexts, they sample more extensively and systematically. External factors, such as rare events, also attract disproportionate attention during information search—even when the shape of the distribution is known. For example, research on lottery buyers suggests that their decision to purchase risky lotteries is driven by a preference for skewness rather than risk, leading them to overweight the observations of a few jackpot winners while neglecting the vast majority of losers in a highly positively skewed environment (Åstebro et al., 2015; Garrett & Sobel, 1999). In contrast, in negatively skewed environments like online reviews, consumers may react negatively to skewness, giving disproportionate attention to a few negative reviews while downplaying many positive ones. (Jung et al., 2020; Wu, 2013). Despite the seemingly contradiction, both examples hint at an overweighting of rare outcomes in the decision-making process. While these studies did not experimentally test the influence of these distributional properties on sampling behavior, and the cause of the overweighting of rare values remains unclear, they suggest that sampling behavior inevitably distorts perception by amplifying rare outcomes.

Sometimes, humans also ignore the probability of an event regardless of its consequence, a phenomenon termed 'probability neglect' by (Sunstein, 2003). This can lead individuals to downplay dangerous risks, such as a lightning strike in a storm, while overestimating another risk with a similar probability, such as a terrorist attack. The key difference lies in how people gather and process information about these risks. In the case of terrorism, sensational media coverage or recommendation systems driven by attention economy often increases people's exposure to such news, making the risk feel disproportionately large. In contrast, while lightning strikes are also rare, they receive much less media coverage than terrorism due to the lack of sensationalism. Consequently, although humans are often biased toward rare events, they may also downplay similar extreme risks when their experience to that risk is limited (Sunstein, 2003). This distinction

underscores the importance of understanding the context and sampling behaviors driving different types of biases in risk perception and decision-making.

Here, we aim to investigate how the skewness of a distribution influences individuals' sampling behaviors and their perception of options. By hypothesizing that individuals pay more attention to rare events, as evidenced by the greater frequency of their sampling, we seek to understand seemingly irrational behaviors, such as excessive lottery spending or exaggerated reactions to rare negative reviews. Prior to presenting our study design, we will review relevant literature on how biased sampling and knowledge of distributional properties influence decision-making.

Biased Sampling Behavior Toward Rare Events

Research on human sampling behavior spans multiple domains, though findings often diverge across paradigms (von Helversen et al., 2018). For instance, in foraging studies, where search incurs energy and opportunity costs, individuals tend to explore locally before moving on to the next food patch (Charnov, 1976; Hills, 2006; Hills et al., 2015). The distinctive feature of this paradigm is the high cost of switching (e.g., energy expenditure to travel *between* food patches), which compels the agent to weigh the trade-offs between continuing with the current depleting option and exploring alternatives with a great cost (von Helversen et al., 2018).

A different paradigm is needed to capture the nature of online information search, where information is abundant, and switching costs between options are minimal. The decision-from-experience (DfE) paradigm offers a more suitable framework, focusing more on how people collect, perceive, and evaluate information from multiple options (Hertwig et al., 2004). In typical DfE experiments, participants sample between two options, each drawing a random outcome from a distribution or pre-generated number sequence, to assess which option is preferable or their willingness to pay (Johnson & Tversky, 1984; Olschewski et al., 2022; Walters et al., 2023). Using the DfE paradigm, Hills and Noguchi (2013) demonstrated that as the number of available options increases, individuals sample a broader range of options but gather fewer samples from each. This sampling pattern appears across various contexts, including consumer psychology (Levav et al., 2012), goal-directed search (Hills et al., 2010; Vul et al., 2014), and social perception (Biella and Hütter, 2024). Studies using eye-tracking in search behavior have shown that individuals spread their attention across numerous alternatives, focusing more closely on options with prominent or favorable features (Bella-Fernández et al., 2022; Rajsic et al., 2015).

One important drawback of the sampling paradigm is that in both, foraging and information search studies, the shape of the distribution is typically unknown. This lack of information complicates efforts to isolate the influence of distributional shapes such as skewness on sampling behavior. By contrast, in contexts such as online reviews or lotteries, individuals often have a general overview of the distribution's shape while searching for additional information (e.g., a histogram of star ratings). To mimic this scenario, we designed an environment where participants are shown the shape of different histogram distribution with five bins and the probability of each bin, but the range of values remains hidden. This setup necessitates information search to learn the objective mean (i.e., the true mean) of the distribution. In such situations, where individuals have access to the distribution's shape but not its range and aim to estimate an option's true value, they might ideally employ stratified sampling—an efficient and unbiased strategy that allocates samples based on the probability of each outcome (a more formal description of the context and supporting proof can be found in the Appendix). By contrast, an excessive allocation of samples to rare events, particularly in skewed distributions, may lead to distorted perceptions of the true mean.

To further examine how people prioritize information search, we consider an alternative hypothesis: that individuals focus more on rare *and* extreme outcomes, rather than just rarity alone. In an environment where all outcomes are equally probable, a search strategy driven purely by rarity would result in an even distribution of sampled outcomes. However, if individuals disproportionately sample values at the distribution's upper and lower bounds, despite their equal likelihood, this would suggest that information search is guided by the extremity of outcomes rather than rarity alone.

Overall, we hypothesize that individuals exhibit biased sampling behavior toward rare outcomes, even when they are aware of the distribution's shape. This study aims to examine how individuals allocate their search efforts across different outcome categories when provided with explicit knowledge of the distribution.

H1: When facing skewed distributions, participants will over-sample the rarest outcome relative to its actual probability, whereas in the uniform condition, their sampling distribution will resemble the true distribution.

Perception of Biased Experienced Means

Another important aspect is how people perceive the information they sample. Past research employed estimation tasks to assess how people perceive a sequence of outcomes. Unlike valuation

and choice tasks, estimation tasks are incentivized based on accuracy, thus not influenced by risk preferences (Olschewski et al., 2021, 2022). Studies have found that people often underestimate the mean of a presented number sequence, possibly due to a "compressed mental number line"—a cognitive bias that leads to estimates lower than the actual mean of the distribution (Oberholzer et al., 2021). However, it is important to note that in these tasks, the number sequences were presented to participants as a continuous stream on the screen rather than requiring an active sampling process, thereby removing the influence of any sampling biases. Consequently, it remains unclear how sampling biases affect the perception of a distribution's true mean.

When individuals draw information from a distribution, the limited number of observations they encounter and how they allocate their search may lead to two distinct concepts of the mean. The first is the *true mean*, representing the central tendency of the underlying distribution. The second is the *experienced mean*, derived by averaging the subset of observations they have encountered. In skewed distributions, biased sampling behavior is expected to produce a biased experienced mean. For instance, oversampling rare outcomes in a negatively skewed distribution may result in an experienced mean lower than the true mean, whereas oversampling rare outcomes in a positively skewed distribution may lead to an experienced mean higher than the true mean. To the extent that individuals infer the true mean based on their experienced mean, any bias in the experienced mean will systematically lead to over- or underestimation of the true mean. Therefore, we hypothesize that:

H2a: For negatively skewed distributions, the experienced mean will be smaller than the true mean (negative bias).

H2b: For positively skewed distributions, the experienced mean will be larger than the true mean (positive bias).

H2c: For uniform distributions, there will be no systematic difference between the experienced mean and the true mean (no bias).

While the experienced mean depends on sampling behavior, how people perceive and adjust for it is influenced by cognitive factors. Individuals adapt their valuations not only based on external factors such as social norms (Fleischhut et al., 2022) but also on internal factors, including confidence (Olschewski & Scheibehenne, 2023; Soll et al., 2019; Yeung & Summerfield, 2012). These types of adjustments require a basic level of causal reasoning about the environment or metacognitive awareness of the decision-maker's own limitations. (Olschewski & Scheibehenne,

2023). For example, consider a real-world scenario in which someone asks a friend about the price of an item purchased during a Black Friday sale. If the friend underestimates the typical price of the item, their reasoning may involve internal adjustments—such as accounting for the likelihood of discounts during sales events. Crucially, this type of adjustment relies on prior knowledge or familiarity with Black Friday, suggesting that better prior information facilitates more accurate adjustments.

By extension, we hypothesize that people will recognize and adjust for sampling biases when sampling from skewed conditions, shifting their estimates closer to the true mean, particularly as the sample size increases. This hypothesis is grounded in research on metacognition adjustment (Olschewski & Scheibehenne, 2023; Yeung & Summerfield, 2012), and rational learning (Le Mens & Denrell, 2011; Olschewski et al., 2024) which suggests that individuals adjust their judgement based on internal factors such as their confidence, and the awareness of their cognitive and behavioral limitations. While participants may over-sample from rare outcomes, we expect that mental adjustments away from the experienced mean will improve estimation accuracy when people have more experience with the skewed distributions due to observing more samples. We do not expect an increase in estimation accuracy with sample size in the uniform condition because here we do not expect the experienced mean to deviate from the true mean in the first place.

H3: Increasing the number of samples from skewed distributions will reduce the difference between participants' estimates and the true mean. However, this effect will not be observed in the uniform condition.

Method

Study Design

Participants took part in an online shopping simulation designed to examine how they sample information to estimate the average score of customer reviews. The task consisted of multiple trials. In each trial, participants were presented with a histogram of customer ratings ranging from 1 to 5 stars, where 5 represented the highest rating. They were informed that each star rating was associated with a text “review” that was converted into a numerical positivity score ranging from 0 (very negative) to 100 (very positive). The cover story was designed to make the instructions more intuitive compared to an abstract task involving statistical distributions and histograms. We also explicitly informed participants that both the ratings and positivity scores were

entirely computer-generated for the experiment and not derived from real customer reviews to eliminate potential inference bias.

There were three within-subject conditions, each consisting of different histogram shapes: negatively skewed, positively skewed, and uniform (four trials per condition). The uniform condition served as a non-skewed baseline to determine whether participants focused just on rare events or a combination of rare *and* extreme events (Mason et al., 2024). The skewed conditions aimed to assess how distribution shape influenced sampling behavior and mean estimations (H1 and H2). During each trial, participants could click on individual bins of the histogram to sample a positivity score corresponding to that star rating (e.g., clicking on the "1-star" bin would reveal the positivity score of a simulated customer who rated the product with 1 star). However, they had a limited number of samples they could collect during each trial.

To assess the impact of number of samples, participants were randomly assigned to one of two between-subject conditions: one group could collect 10 samples, while the other group could collect 20 samples from the distribution. This manipulation was designed to investigate whether participants adjusted their sampling strategies and estimations when given more samples (H3).

Participants' primary task was to estimate the product's average positivity score based on their samples. To incentivize accuracy, participants were offered a potential reward of up to £8, depending on how closely their final estimate matched the true average positivity score.

Stimuli Creation

The experimental design and hypothesis were preregistered on AsPredicted ([link](#)). To create the stimuli, we used three distinct beta distributions, each representing a specific shape condition: negatively skewed $\beta(3,1)$, positively skewed $\beta(1,3)$, and uniform $\beta(1,1)$. These distributions were then scaled by a random multiplier between 50 and 100 for each trial to mask the true range of positivity scores and to prevent participants from inferring the underlying distribution's exact range based on previous trials.

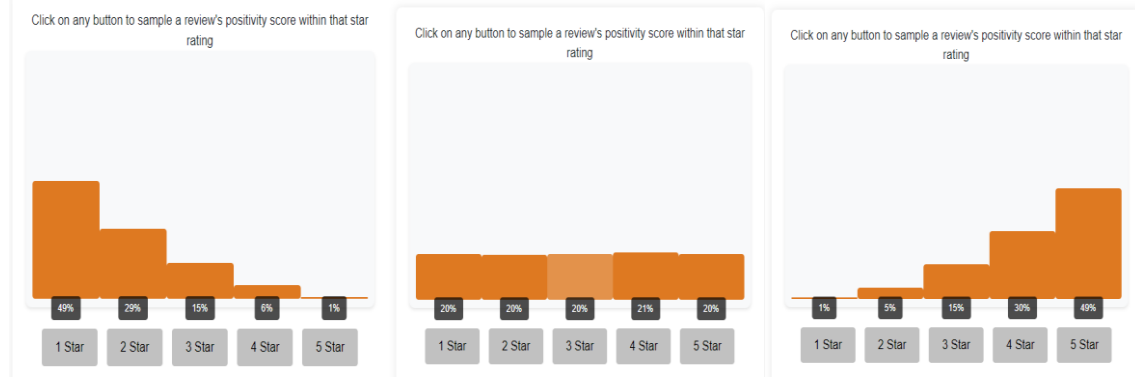
Histograms presented to participants were created by binning the beta distributions into five intervals, with each bin displaying the probability of receiving a sample from that interval if drawn randomly (see Figure 4.1). When participants sampled from a bin, a score was drawn from the corresponding interval in the continuous beta distribution and shown below the histogram.

Procedure and Incentives

Following the pre-registered sample size, we recruited 180 participants through Prolific (2024) to take part in our online experiment. The task was implemented using Otree (Chen et al., 2016). Participants were randomly assigned to either the 10- or 20-sample condition. At the experiment's start, participants were briefed on sampling rules and positivity scores and informed of a potential £8 bonus based on the accuracy of their estimates in one randomly selected trial.

Figure 4.1

Screenshot of different distributional shape conditions.



Note. From left to right: positively skewed condition, uniform condition, and negatively skewed condition. Clicking on each button below draws a sample from the respective interval of the underlying distribution.

Participants began with a practice trial featuring a randomly selected distribution shape to familiarize themselves with the task. After sampling from the histogram, they provided an estimate for the average positivity score and proceeded to the next round. Following the practice trial, participants completed comprehension checks to ensure they understood the task before beginning the main task. In the main task, they completed the three within-subject conditions, with four randomized trials per condition. Upon completing all trials, participants provided demographic information. Following the preregistered exclusion criteria, we excluded trials in which only one bin was sampled. We also excluded participants who provided the same estimation in all trials, as well as any participants with two excluded rounds. This resulted in a final sample of 145 participants ($Mean_{age} = 39$, $SD = 12$; 47% male, 53% female, 50% in the 10-samples condition).

Results

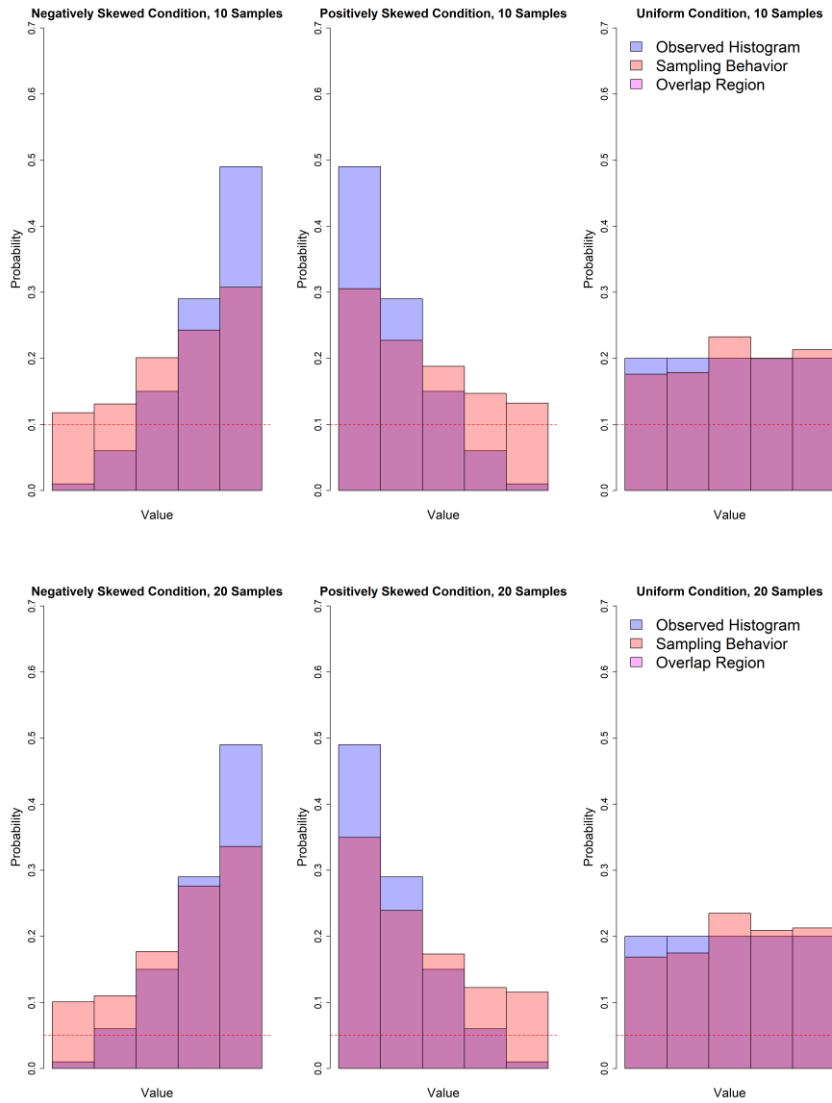
Sampling Bias

Binomial tests were conducted to evaluate the hypotheses regarding participants' sampling biases toward rare outcomes in skewed distributions. In particular, we compared the observed sampling behavior with a simple heuristic of stratified sampling, which is a simple, efficient and unbiased sampling rule in this task (proof in appendix). In the negatively skewed condition, participants spent 11% of their total sample on bin 1, substantially more than the expected 1% ($p < .001$, 95% *CI* [0.10, 0.11]). Similarly, in the positively skewed condition, participants sampled from bin 5 (12%) more than suggested 1% by the stratified-sampling heuristic ($p < .001$, 95% *CI* [0.11, 0.13]). These results provide evidence of a sampling bias toward rare events in skewed distributions. Across both the 10- and 20-sample conditions, participants allocated approximately 10% of their samples to rare events, suggesting a strong and consistent tendency to focus more on these outcomes as compared to stratified sampling.

In the uniform distribution condition, we observed a slight undersampling in bin 1 ($P_{bin1} = 0.17$, 95% *CI* [0.16, 0.18], $p < .001$), and no statistically significant difference in bin 5 (95% *CI* [0.20, 0.22], $p = .063$), providing evidence against a strong sampling bias toward a combination of rare *and* extreme events in this condition. However, we observed a small positive bias in bin 3 of the uniform condition ($P_{bin3} = 0.23$, 95% *CI* [0.23, 0.24], $p < .001$). Although this was not predicted in our hypothesis, the oversampling of bin 3 suggests that participants may have employed a simple rule of thumb: sampling from the middle bin to approximate the mean when the distribution shape was uniform. This demonstrate that participants understand the task quite well and adapt their sampling strategies to solve the task.

Figure 4.2

Comparison of the observed histogram with participants' actual sampling behavior in the 10- and 20-sample conditions.



Note. Each bar of the histogram represents bins 1 to 5, from left to right. The pink area represents the overlap between sampling behavior and the observed histogram, while the other colored areas suggest deviations between the two. The purple area could be interpreted as undersampling, while the orange area suggested oversampling compared to the observed histogram. The red dotted line represents the proportion of samples allocated if each option were sampled exactly once.

Experienced Mean and Estimation Bias

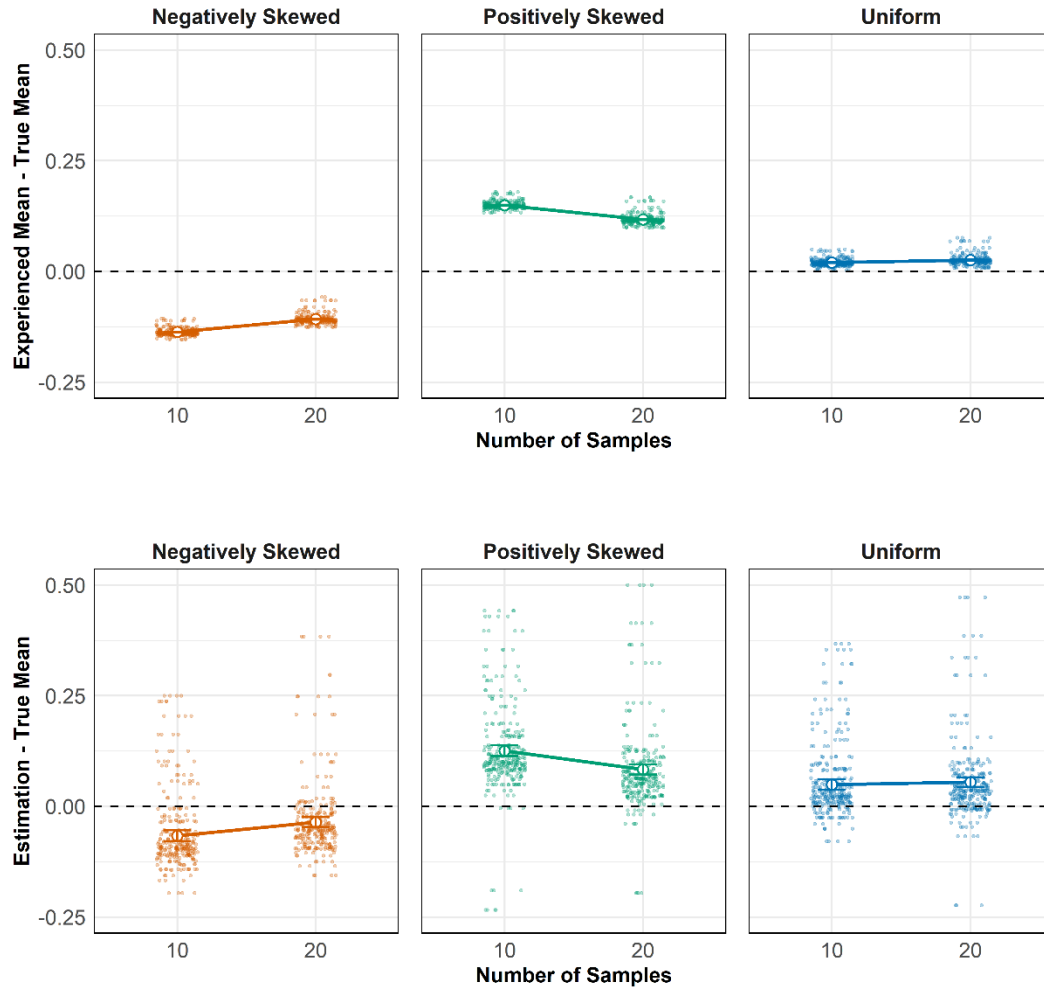
An important aspect of decision-making lies in how individuals perceive and internalize the information obtained through their sampling behavior, particularly in relation to systematic biases across different distributional shapes and levels of search resources. In this section, we focus on our two key hypotheses: the deviation between the experienced mean and the true mean (H2), and the deviations between participants' estimation and the experienced and true means (H3). To investigate these hypotheses, we constructed a series of Bayesian linear mixed-effects models with a common structure—using the same set of predictors but varying in dependent variables. Models were implemented in R (R Core Team, 2021), using the *brms* package (Bürkner, 2021). Each model included random intercepts for participants and fixed slopes for the main effects, and was estimated using the default priors for the Gaussian family in *brms*.

For H2, a Bayesian linear mixed-effects model was used with sampling bias, defined as the deviation of the experienced mean from the true mean of the distribution, as the dependent variable. The predictors included skewness conditions (negatively skewed, positively skewed, and uniform), sample sizes (10 and 20), and their interaction effect. We first report the experienced mean for the 10-sample condition. In the negatively skewed condition, the sampled mean was below the true mean, yielding a negative sampling bias ($M = -0.14$, 95% *CI*: [-0.15, -0.12]). Conversely, in the positively skewed condition, the sampling bias was positive ($M = 0.149$, 95% *CI*: [0.148, 0.150]). The model indicated a small effect of sample size, with larger samples (20 vs. 10) reducing sampling bias in both the negatively and positively skewed conditions. The upper panel in Figure 4.3 shows that the experienced means in both the 10- and 20-sample conditions aligned with our predictions in H2a and H2b: the experienced mean was greater than the true mean in the positively skewed condition (slope above the 0 line) and smaller than the true mean in the negatively skewed condition (slope below the 0 line).

The substantial sample bias in the two skewed conditions suggests that participants would incur a potential incentive loss of 13.5% (£1.08) if they base their estimation solely on the experienced mean. In the uniform condition, positive biases in experienced means were observed in both the 10-sample ($M = 0.02$, 95% *CI*: [0.019, 0.021]) and 20-sample conditions ($M = 0.026$, 95% *CI*: [0.024, 0.027]). This bias in the experienced mean in the uniform condition highlights the importance of sampling behavior, suggesting that even subtle biases in search behavior could lead to a biased experienced mean.

Figure 4.3

The effect of distribution shapes and sample sizes on sampling bias and estimation accuracy



Note. Upper: The effect of distribution shapes and sample sizes on the difference between the experienced mean and the true mean (sampling bias). Lower: The effect of distribution shapes and sample sizes on the difference between participants' estimates and the true mean (estimation accuracy). The y-axis in both plots is standardized to the same scale, ranging from -1 to 1. Error bars represent 95% credible intervals.

In the previous analyses, we examined how biased sampling behavior led to deviations between the experienced mean and the true mean. To test H3, we now compare participants' estimates with both the experienced mean and the true mean. The first analysis focuses on the deviation of estimates from the true mean across different skew conditions (negatively skewed,

positively skewed, uniform), sample sizes (10 and 20), and their interaction effect. In the negatively skewed condition, the deviance of participants' estimates to the true mean was smaller in the 20-sample condition ($M = -0.036$) than in the 10-sample condition ($M = -0.067$, 95% *CI* for difference: $[-0.058, -0.004]$). Likewise, the positively skewed 10-sample condition ($M = 0.126$) showed a larger bias than the 20-sample condition ($M = 0.083$, 95% *CI* for difference: $[0.014, 0.082]$). However, in the uniform condition, differences between the 10-sample ($M = 0.049$) and 20-sample ($M = 0.055$) conditions were not different (95% *CI* for difference: $[-0.021, 0.007]$). These findings support our hypothesis that estimation accuracy improved more in the skewed conditions than in the uniform condition when participants had more samples, although the effect size was modest—an additional 10 samples led to only a 3.6% improvement in accuracy. A linear projection suggests that participants would require approximately 30 samples in the negatively skewed condition and 40 samples in total to achieve an estimate within 5% of the true mean.

Analyzing the difference between experienced means and participants' estimations—referred to as estimation adjustment—revealed that, on average, participants' estimates were 7% closer to the true mean compared to their experienced means in the negatively skewed distribution. (10-sample: $\Delta = -0.07$, 95% *CI* $[-0.08, -0.06]$; 20-sample: $\Delta = -0.07$, 95% *CI* $[-0.08, -0.06]$). In the positively skewed condition, estimates were also closer to the true mean, though to a lesser extent (10-sample: $\Delta = 0.02$, 95% *CI* $[0.01, 0.04]$; 20-sample: $\Delta = 0.03$, 95% *CI* $[0.02, 0.04]$). Although we observed sampling bias in all three distribution conditions, estimation adaptation only occurred in the skewed conditions but not the uniform condition. A potential explanation is that participants were aware of substantial sampling biases in the two skewed conditions but not the relatively small bias in the uniform condition. These results dovetail with our hypothesis regarding how people adjust their estimates based on the experienced mean. In the skewed conditions, these adjustments led to an average increase of 4.8% in accuracy compared to relying solely on the experienced mean. However, they did not fully eliminate the impact of sampling bias. Despite attempts to adjust, participants' estimates still deviated from the true mean by an average of 9.3% in the skewed conditions, highlighting the negative consequence of sampling bias in this context.

A Potential Explanation for Biased Sampling Behavior

Research on search behavior in online and consumer context suggests that exploration increases with lower search costs (Hills & Hertwig, 2010; Hills et al., 2013; Levav et al., 2012). However, participants in previous experiments often searched under uncertainty, having little a-priori information about the outcome distribution. In our experiment, we observed that participants

prioritized sampling all five available outcomes, even though they knew that information in the smallest bin of a skewed distribution contributed little to estimating the true mean. While this sampling pattern introduced a bias in the experienced mean, participants partly adjusted for it when estimating the mean, suggesting that they were aware of their sampling bias. This implies that participants' behavior was not merely an artifact of irrational sampling but rather a strategic approach to information gathering.

A potential explanation for why participants oversampled the rarest bin is that they prioritized knowing the full range of the distribution to estimate the mean. According to a framework called cognitive fencing (Liu & Scheibehenne, working paper), individuals are more certain that all values within an experienced range are possible, while they remained uncertain about the probability of values outside the experienced range. To reduce this uncertainty, they preferred to allocate samples to discovering the full range rather than focusing solely on accuracy and ignoring the smallest bin, leading to a sampling bias in skewed distributions. An analysis of unique buttons sampled at the participant level showed that out of 145 participants, only three did not sample all five buttons in any of their trials across the three environmental conditions, suggesting that this may have been part of their information search strategy. We utilized the cognitive fencing framework, which accounts for both the shape of the objective distribution that participants observed and the cognitive preference for sampling the entire values range, to model the sampling process of participants. The model attributed participants' sampling behavior, represented as a histogram $f(x)$, as a function of $g(x)$, the presented star rating histogram, and $u(x)$, a uniform shaped histogram that captures the inherent tendency to sample the full range:

$$f(x) = (1 - w) * g(x) + w * u(x)$$

The parameter w determines the relative weight of the two histograms, quantifying how much the uniform distribution contributes to the final sampling behavior. We set w with a prior of $N(\mu_w, \sigma_w)$ allowing each participant to have their own parameter. We fitted the model in R using the rstan package (Stan Development Team, 2025) with the No-U-Turn Sampler (NUTS) to estimate w . Sampling was performed with four chains, each running for 2000 iterations, including 1000 warm-up iterations. The estimated weight parameter w had a posterior mean of $\mu_w = 0.51$ (95% CI [0.21, 0.74]), indicating that the observed sampling behavior was driven by a fair contribution between $u(x)$ and $g(x)$. This combination means that even though participants were aware of the shape of the distribution, their actual sampling distribution was influenced by both the uniform and the objective histogram equally. Individual differences in this strategy were captured

by an estimated $\sigma_w = 0.22$ (95% *CI* [0.21,0.22]), suggesting substantial variability in how participants weighted the uniform distribution relative to the objective histogram.

Using the parameter w , we can determine whether the observed estimation adjustments stemmed from a response to different distributional shapes (such as an unfavorable/favorable preference for skewness) or from participants' own strategic adjustments. If participants were aware of their strategy and incorporated it into their estimations, then differences observed across the three distribution conditions would primarily reflect the underlying strategy w . In this case, including w in the model should eliminate the effect of the skewed conditions on estimation adjustment. Conversely, if the distributional shapes independently influenced estimation, its effect should persist even after accounting for w .

To test this, we fitted a Bayesian hierarchical model using absolute sampling adjustment as the outcome variable, with w , distributional shapes, and their interaction as predictors, and participant as a random effect. We specified random intercepts at participants' level and fixed slopes structure for the main effects, and used the default Gaussian family priors in the *brms* package. We found only a main effect of w ($b = 0.15$, $SE = 0.02$, 95% *CI* [0.10, 0.19]), with no effects of the shapes of distribution (95% *CI* [-0.02, 0.02]) or their interaction effect (95% *CI* [-0.05, 0.02]). This supports the idea that while the shape of the distribution influences sampling behavior, its effect on estimation adjustment disappeared when w was taken in account. The extent of adjustment varied according to the magnitude of w , reflecting the premium of sampling the entire range of values, or that participants prioritize sampling all the outcomes, even when it is suboptimal. These results support the hypothesis that participants apply a strategy incorporating sampling all outcomes to have more certainty about the range of values. The estimated influence of w suggests a consistent bias towards an even allocation of samples across bins, particularly when the true distribution is highly skewed.

Conclusion and Discussion

This study examined how distribution shape and sample size impact sampling bias and estimation. We observed significant sampling biases in skewed distributions, with participants disproportionately sampling rare outcomes—bin 1 in the negatively skewed condition and bin 5 in the positively skewed condition. This behavior was consistent across both the 10-sample and 20-sample conditions, leading to biased experienced means due to extra samples of rare outcomes. In addition, we also found a small deviation from the expected sampling behavior in the uniform

distribution condition, which contributed to slight deviations in experienced means in this condition as well. Together, these results highlight the importance of sampling behavior in unveiling the number sequence that participants observed.

In our experiment, increasing the number of samples that could be drawn reduced both sampling- and estimation errors in skewed distributions, with participants adjusting their estimates closer to the true mean. This indicates that a larger sample size provides a more accurate representation of the underlying distribution by mitigating the effects of sampling biases. By examining the true mean, experienced mean, and participants' estimates simultaneously, our study reveals two critical insights regarding the underestimation and overestimation of means in skewed distributions. First, when comparing the true mean with participants' estimates, we found evidence for overweighting rare outcomes, where people overestimate the mean of positively skewed distributions and underestimate the mean of negatively skewed ones. This pattern aligns with prior research on the perception of skewed distributions (Åstebro et al., 2015; Garrett & Sobel, 1999; Olschewski et al., 2024). The overweighting of rare outcomes in our experiment could be directly linked to participants' sampling bias, which overrepresents such outcomes—a finding consistent with previous research on decisions from experience (Hills et al., 2013).

When focusing solely on the experienced sequence of numbers sampled by participants and comparing these with their estimates, our results could be interpreted as if they underweighted rare outcomes. Specifically, participants underestimated the experienced mean in positively skewed conditions and overestimated it in negatively skewed conditions because they adjusted their mean estimates. A possible explanation for this pattern of results is that participants' were aware of their own sampling biases (Olschewski & Scheibehenne, 2023; Soll et al., 2019). This hypothesis is further supported by the fact that in the uniform distribution condition, where sampling bias was minimal, participants' estimates and their experienced mean was more closely aligned.

Overall, our findings contribute to a better understanding of the overweighting and underweighting of rare events by highlighting how conclusions depend on different points of comparison. Furthermore, our results add to the expanding literature on the joint role of behavioral and cognitive factors in shaping human judgment and decision-making. We observed that sampling behavior is sensitive to features of the choice environment—particularly the presence of rare events—which in turn shapes perception. This aligns with earlier work suggesting that how people sample information can bias what they ultimately perceive and decide (Hertwig & Pleskac, 2010). Importantly, our results go beyond the traditional naïve assumption that people simply rely on what

they observe (Juslin et al., 2007). Instead, we find evidence of cognitive adjustment: individuals appear to recognize the limitations in their own sampling behavior and attempt to correct for them, even if their adjustments are only partially successful.

Our findings resonate with studies that examine the interplay between cognition and sampling behavior in judgment and decision-making. While we did not directly explore the role of higher cognitive function such as motivation, other research has shown that motivational factors can influence how people sample information. For example, Biella and Hütter (2024) found that interest-driven and disinterest-driven search strategies lead to asymmetric sampling: individuals tend to search longer when disinterested and terminate search early when they encounter counterevidence, resulting in more objective information gathering. In our study, sampling bias had a stronger net influence on judgement than cognitive adjustment. However, Le Mens and Denrell (2011) went even further in challenging the naive assumption by demonstrating that even rational sampling processes can yield systematic judgment errors, particularly when individuals prioritize alternatives with more interesting or preferable outcomes. This underscores the critical role of cognitive filters, especially in contexts where individuals have personal stakes or strong prior expectations.

Future Research and Implications

Our findings highlight the crucial role of sampling biases in estimation processes, particularly in skewed distributions, and call for further research into how individuals adapt to such biases, especially when they are less overt. Prior research has examined cognitive factors in decision-making (Olschewski et al., 2022), such as memory (Haines et al., 2022; Sahakian et al., 2023), attention (Pleskac et al., 2021), and the integration of numerical information (Oberholzer et al., 2021; Prat-Carrabin & Woodford, 2022). However, our findings underscore the importance of behavioral factors and their critical role in shaping decision outcomes (Bella-Fernández et al., 2022; Mehlhorn et al., 2015; von Helversen et al., 2018). Behavioral influences, such as sampling biases, have a direct and significant impact on subsequent estimation processes. This provides another perspective on current decision-making theories, which often focus on how information is processed in the brain but treats the informational input as given.

Our work also indicates that cognitive biases can arise at an earlier information sampling stage already. Based on the cognitive fencing framework (Liu & Scheibehenne, working paper), the tendency to oversample the rarest bin in a skewed distribution may have served as a strategy to

reduce uncertainty by exploring the full range of values. The model's parameter helps to explain the observed differences in estimation adjustments across various distribution shapes, emphasizing the interplay between cognitive processes and information search behavior in this task. Additionally, sampling bias alone could not account for the stronger estimation adjustment observed in the negatively skewed condition compared to the positively skewed condition. This result suggests that further investigation into how people perceive and integrate information such as numerical perception (Oberholzer et al., 2021; Olschewski et al., 2022), and information integration strategies (Leuker et al., 2019; Pleskac et al., 2021; Yeung & Summerfield, 2012; Zilker, 2022) may provide valuable insights.

Our findings also have practical implications. For example, in e-commerce, sampling biases, combined with review aggregation algorithms, can distort product perceptions. Consumers often give disproportionate weight to a few negative reviews, which may not accurately represent the broader population, leading to skewed purchasing decisions and potential dissatisfaction (Qahri-Saremi & Montazemi, 2023; Wu, 2013). Similarly, in political voting contexts, the oversampling of rare and extreme opinions can reinforce pre-existing biases, distort perceptions of political candidates, and even sway election outcomes. On social media, where rare events are more likely to go viral, these events can disproportionately shape users' perceptions of reality, emphasize the consequences of fake news or contribute to the spread of false or biased narratives. Although individuals can adjust their opinions and estimations when aware of sampling bias, it remains unclear whether people are aware and can adjust to the bias feedback loop in adaptive systems, recommendation engines, and social media platforms. Without such awareness or intervention, users may remain unaware of the biases shaping their perceptions, resulting in continued misjudgments.

In the context of online information, large language model (LLM) chatbots and AI technologies may be used for reducing biases, but they also hold potential to amplifying bias and misperception. On one hand, LLMs can promote more balanced and objective information search by presenting answers in a comprehensive and impartial manner—even when the user's initial query or search behavior is biased. In this sense, they can serve as a valuable partner or guide, encouraging users to adopt less biased sampling strategies. LLMs benefit from the law of large numbers: by learning from vast and diverse datasets, their knowledge base far exceeds that of any individual, which can contribute to more balanced responses. Yet, they still inherit and reflect

biases present in the training data, sometimes resulting in harmful or morally and practically misleading suggestions (Hanna et al., 2025; Salatino et al., 2025).

While designers often turn to automation to enhance system efficiency and safety, it is important to note that human judgment often becomes even more critical as automation grows in power and ubiquity (Lee & Seppelt, 2023). Therefore, designers of adaptive systems—such as recommendation algorithms and large language models (LLMs)—should account for users’ existing sampling biases in order to support more informed decision-making and help mitigate the influence of biased information. Small interventions, such as providing prompts about the prevalence of certain public opinions online or the representativeness of the information collected so far, or using LLMs to detect LLMs-generated contents could potentially help reduce sampling bias of users or bias caused by the system. For example, these prompts could highlight that extreme opinions are not representative of the general public, or the contents being sampled are generated by other LLMs. Additionally, increasing the level of estimation adjustment during the decision-making process may further mitigate these biases.

Limitations

One limitation is the potential influence of participants’ prior experiences, particularly with positively skewed e-commerce ratings, which may have shaped their sampling and estimation behaviors. For example, frequent exposure to negatively skewed product ratings in e-commerce environments could lead participants to internalize specific real-life problems, such as rating inflations (Aziz et al., 2023; Skreta & Veldkamp, 2009), thereby influencing their judgment in our experimental settings. We selected the context of online shopping because it is familiar to most participants, substantially enhancing task comprehension and reducing rejection rates due to failed comprehension checks.

Additionally, the study focused on two specific sample sizes (10 and 20), which constrains the generalizability of the findings. Increasing the sample size or encouraging more search could potentially reduce sampling biases, but this assumes participants can maintain the same level of attention over a much longer task—an assumption that may not always hold. Because we asked participants to use all available samples before advancing to the next trial, unwanted behaviors such as repeatedly sampling one option to quickly move on to the next trial or randomly clicking might occur. This likely introduced unwanted noise and bias into the data and diminished the reliability of the findings.

Future research could investigate how prior knowledge, such as familiarity with specific rating distributions (e.g., positively skewed ratings common in e-commerce), affects sampling biases and estimation accuracy. Moreover, exploring how learning and experience shape these biases over time could offer valuable insights into the dynamics of human judgment and decision-making. For example, examining how participants adjust their sampling behaviors after receiving feedback on their biases or experiencing varied contexts could help identify effective strategies for mitigating bias across diverse domains, from social media and e-commerce to political decision-making and beyond.

Chapter 5: General Discussion

Overall, this dissertation integrates findings across three projects to advance our understanding of how people search for information and make decisions in complex environments, revealing the cognitive and contextual factors that shape adaptive behavior. In the first project, we re-examined the choice overload phenomenon through the lens of search behavior and found that the strategies people used to search—particularly search breadth—had a strong effect on their search efficiency, measured by Effort-Adjusted Decision Quality (EADQ). This suggests that search breadth tends to improve decision quality, especially when the number of options is small to moderate. However, as the number of options increased (up to 40), the benefit of search breadth diminished, indicating that the cost of managing many options can override its potential advantages. Notably, this decline occurred even in low-noise environments, suggesting that clearer signals do not necessarily simplify decision-making when there are too many options. We also observed that participants underestimated the experienced mean value of the chosen option more in noisy conditions, indicating some awareness of potential bias under noise, but they failed to adjust for the added complexity of larger choice sets. These results highlight the cognitive limits of adaptive search and highlight how environmental structure and set size could influence decision efficiency and perception.

In the second project, the goal was to examine the role and prevalence of the sequential order of search behavior on decision quality and perception. We showed that as the number of samples decreased, participants tended to rely more and more on search breadth as the main search strategy. Furthermore, the interaction between search strategy and choice environment significantly predicted decision quality: in frequent-low environments, where the highest expected value options have frequent but low outcomes, search breadth strategies led to poorer decisions. In contrast, in rare-high environments—where high-value options have rare but large outcomes—search breadth improved decision quality. For estimation accuracy, participants underestimated the observed outcome probabilities when few samples were available, regardless of the choice environment, suggesting they were adapting more to their own strategy than to the environment. To highlight the role of the sequential order of search behavior, various metrics were compared in terms of how well they predicted decision quality. Metrics incorporating sequential order—such as average search length and search breadth—predicted outcomes better than standard measures like searches per option, entropy or inertia. Shuffling the search order diminished the predictive power of sequential-based metrics but not that of non-sequential ones, supporting our assumptions. Overall, the findings

emphasize the critical role of sample size and the dynamic interplay between environment and strategy in shaping decision-making.

In the third project, we aim to investigate how features in the choice environment—such as the shape of the distribution—can influence how people distribute their search behavior. By manipulating the skewness of histograms and having participants sample bins to estimate the means, we highlighted a consistent bias toward sampling rare events in skewed conditions—people tended to oversample rare outcomes far beyond what an unbiased and efficient sampling rule would suggest. While increasing the number of samples modestly reduced estimation bias, we found that participants' self-adjustments had a stronger effect. By either underestimating or overestimating the biased experienced mean, participants succeed in bringing their estimations closer to the true mean in both skewed conditions. However, these adjustments were still insufficient, and the effect of biased sampling behavior remained substantial—highlighting the enduring impact of biased search behavior despite clear incentives for accuracy.

Future Research and Implications

We hope our thesis highlight the crucial role of integrating search behavior into research decision making can contribute greatly to the practical and theoretical aspects of human decision making. The results of the studies highlighted the three interconnected crucial aspects of decision that we suggest should be examined closely together, which is the search behaviors, perception, and choice environments.

Search Behavior

Search behavior not only provides valuable information prior to decision-making but also plays a crucial role in shaping perceptions of the choice environment and the options ultimately selected. Our research contributes to the decision-making literature by highlighting the importance of search behavior in several key areas: reducing search effort in larger choice sets (Kellen et al., 2020; Nobel, 2021), predicting decision outcomes across different environments, and revealing the potentially harmful effects of biased sampling strategies. In the more focused domain of search behavior, our findings underscore the significance of the sequential order of search in shaping both decisions and perceptions. We further emphasize the bidirectional relationship between the choice environment and search behavior—demonstrating that various features of the environment can systematically influence how people search. Our results are consistent with prior research showing

how external factors can alter search patterns (Hills et al., 2015; Mehlhorn et al., 2015; Olschewski et al., 2022).

These insights also have practical implications. For instance, capturing users' search behavior on online shopping platforms can help companies and third-party providers better understand not only their customers' decision-making processes but also the reasons behind decision avoidance or indecision—ultimately leading to improvements in customer satisfaction and search efficiency. In addition to traditional process-tracking tools such as eye tracking or mouse tracking, we suggest that researchers carefully identify which features of search behavior they want to capture, such as search sequences or search direction, to gain deeper insight into consumer behavior and enhance predictive models of preferences and choices. Moreover, our findings offer valuable guidance for practitioners, industry stakeholders, and policymakers aiming to design more effective choice environments. By mitigating biased sampling, reducing unnecessary search effort, and encouraging more informed decisions, such environments can enhance user experience and decision quality.

Looking ahead, we believe future research can benefit significantly from studying how people search—particularly by recording the sequential structure of search behavior—to better understand how individuals form different perceptions of the choice space, arrive at different decisions, and evaluate their outcomes. Our research proposes several methods for integrating sequential search information into metrics, including search breadth and average search length. However, there is room for refinement. Future research should aim to develop more sophisticated metrics that capture the starting and ending points of search sequences, which have been shown to convey critical information about decision processes and participant perceptions (Baumann et al., 2022; Jonas et al., 2001; Zwick et al., 2003). For example, different starting points can anchor comparisons in distinct ways, influencing how subsequent information is evaluated. Similarly, analyzing the final few search steps can shed light on how individuals deliberate among the last remaining options, offering insight into what features ultimately drive choice (Martinovici et al., 2021; Pleskac et al., 2021). Another promising avenue involves analyzing the directionality of search, which can reveal systematic exploration strategies. For instance, prior work has shown that people often search from left to right (Martinovici et al., 2021; Smith & Krajbich, 2019), or initiate search from the most visually salient or attention-grabbing option (Martinovici et al., 2021; Townsend & Kahn, 2014). Incorporating these patterns into improved search metrics offers a more elegant and informative way to understand and predict human decision-making.

A second direction for future research is to theorize and simulate search behavior through algorithmic modeling. While computer scientists have developed various algorithms to efficiently explore novel environments (Hein et al., 2020; Mellers et al., 2023; Nobel, 2021), it remains an open question what reward functions guide human information search. Although several candidate theories exist (Gathergood et al., 2019; Stewart, 2009; Stewart et al., 2005), validating them requires not only building algorithms that closely mimic observed human behavior but also creating experimental paradigms that can causally test how specific reward structures shape search behavior. While this endeavor is considerably more complex, we believe it is ultimately more fruitful for advancing both theoretical understanding and real-world applications in human decision science.

Perception

Perception plays an essential role in shaping how people search for information and make decisions. In our research, we emphasize the distinction between the experienced environment and the true environment, which is central to understanding perception. The true environment represents the actual state of the world—what people ideally want to learn. However, through search behavior and its constraints, people only encounter a subset of the true environment, resulting in what we call the experienced environment. This distinction is not new. Philosophy and prior research have long explored similar ideas—for example, Plato’s allegory of the cave, in which people interpret shadows on a wall as reality (Ferguson, 1922).

However, contrary to some previous research on perception and decision-making, where perception often acted as a distorting lens of the real world, we provide another often-overlooked role of human cognitive ability, the ability to aware their own constraints (Le Mens & Denrell, 2011; Olschewski & Scheibehenne, 2021; Rosenbaum et al., 2022). While behaviors biases existed, people can recognize that their experienced environment is biased and may not fully reflect the true environment in many situations. This metacognitive ability enables them to adjust their judgments in response to biases—such as those arising from small sample sizes or from biased sampling behavior. Our findings contribute to the literature on perception and decision-making by showing how people mentally adjust their estimates from experienced environment to be closer to the true environment. We also demonstrate that examining the experienced environment is crucial in perception research. Without this perspective, researchers might draw incorrect conclusions. For example, in Chapter 4, comparing participants' estimates directly to the true environment could falsely suggest that they underestimate or overestimate the true means in skewed conditions,

suggesting misperception. In fact, participants are often adjusting their estimates away from the experienced mean in ways that bring them closer to the true mean.

Understanding how people perceive information and revise their beliefs has practical value. Companies can use our insights to anticipate how consumers react when they have limited exposure to products. For instance, people tend to adjust more when they have fewer samples per option, which can help explain how expectations shift after first impressions. On the other hand, our research also shows that the negative effects of biased sampling can be so strong that even with adjustment, estimates remain far from the truth. This is especially relevant when people are exposed to rare events—like a shocking news story or a viral post. People may know such events are unrepresentative, but our findings show that when the deviation between the experienced and true environments is large enough, adjustment becomes more difficult. These insights apply to areas like how people interpret negative reviews in e-commerce or why extreme political views gain traction—because both are unrepresentative but receives disproportionate attention. Our research suggests ways to design choice environments that reduce the impact of biased sampling. By doing so, we can create experienced environments that better reflect the true environment, reducing the need for adjustment.

Future work can build on our findings in several ways. One direction is to go beyond measuring final estimates and instead gather more information on how people perceive the entire set of options or the environment. One promising method is the distribution builder task (Hu & Simmons, 2025; Quandt et al., 2022), which allows participants to explicitly describe the distribution they believe they are sampling from. Another direction is to track how perception develops over time. This could reveal how people gradually form the experienced environment through sampling and how their understanding of it evolves. However, these methods can make tasks longer and more complex, so more efficient measurement approaches are needed. Researchers could also explore how people perceive specific attributes of the environment and how those perceptions guide search in different contexts. For example, prior work has looked at how people perceive features like color or size (Bhatnagar & Orquin, 2022; Orquin et al., 2021; Qu et al., 2022), but it is important to study these factors in conjunction with how they are searched, chosen, and remembered in decisions. Finally, we recommend integrating biometric measures—such as pupil dilation, skin conductance, or EEG—to better understand attention and cognitive processes during search. These tools can help reveal what captures attention and causes confusion, offering a deeper view of search dynamics.

Choice Environment

The choice environment is a controlled variable in multiple research topics in decision-making and has a strong impact on both perception and search behavior. The true environment is closely linked to the experienced environment, and both need to be examined side by side to fully understand search and perception.

Our studies offer modest contributions to the literature on choice environments and search behavior by demonstrating that even small design features—such as the number of available options or the skewness of information categories—can significantly influence how individuals allocate their search efforts. In addition, we found that certain features of the choice environment that require learning (e.g., noise levels or environments with rare-high vs. frequent-low outcomes) substantially affect decision-making and perception but have a lesser impact on search behavior. These findings align with prior research on effort allocation, which suggests that people often develop an initial strategy based on early information about their resources and environment, and then stick with it (Levav et al., 2012; Pachur, 2022). For example, the order in which participants encounter different set sizes—seeing a large set first followed by a smaller one—can have a strong and lasting effect on search behavior, as people tend to stick with their initial strategy even when the environment changes.

Insights into the relationship between choice environments and search behavior reveal valuable practical implications. On one hand, it is useful to design choice environments that minimize search effort for customers or users and to remove or redesign factors—such as skewed distributions—that can bias search behavior. For example, our research shows that having more than 40 options can significantly reduce search efficiency, particularly when the larger set size does not offer substantially greater benefits while requiring much more effort to explore. Therefore, conducting analyses that weigh the marginal benefit of each additional option against user search effort can help industries improve search efficiency (Bollen et al., 2010; Nobel, 2021; Zhou et al., 2020). On the other hand, it is important to consider natural search behavior when designing choice environments. One key factor that was not directly controlled in our experiments is search cost between and within options. In our experiments, the cost of searching between and within options was equal—similar to many online environments. However, in many real-world cases, the cost of searching between options can be substantially higher and may lead to very different search strategies. For instance, it is unrealistic to expect people to adopt a breadth-based search strategy when shopping for houses in different parts of a city, due to the high cost of traveling between

locations. In such cases, depth-oriented search strategies are likely preferred. Therefore, these factors must be considered to better understand how people search in specific environments.

Although the choice environment is often tightly controlled in decision-making research, we believe it presents promising future directions for exploring its effects on search behavior and perception. Future research could investigate additional factors that influence search behavior to uncover new and unique patterns of decision-making. Moreover, designing choice environments that selectively influence perception without restricting search behavior could help people make better decisions without imposing artificial or forced constraints. Future research could also focus on aspects of the choice environment that go beyond modifying incentive structures and instead aim to influence search behavior and perception more directly—offering a complementary direction to traditional economic research, which often emphasizes incentivization.

Limitations

There are several important limitations in our research findings that both theoretical and applied researchers should take into account. The first concerns the generalizability of the results. As briefly mentioned, the structure of search costs in the choice environment must be considered when forming hypotheses about the strategies individuals are likely to adopt. In all of our studies, there was no additional cost associated with searching between options, and searching within an option was not substantially less effortful. However, in other environments—such as foraging task or the secretary task (Hills, 2006; von Helversen et al., 2018)—searching the next options often incurs higher costs, either through greater search effort or opportunity costs (i.e., the risk of forgoing a good option). Therefore, we caution against generalizing our findings to contexts that involve fundamentally different search cost structures.

Second, the method of data collection in our studies was entirely computer-based, with participants recruited from Prolific interacting in a two-dimensional, online environment. Environmental factors such as color, layout, and spatial positioning are known to significantly influence search behavior and attentional patterns were not taken into account in our studies. As such, it remains unclear whether our findings would generalize to more realistic settings, such as a three-dimensional physical environment like a supermarket, where individuals search among tangible options. Moreover, the online setting only allows for limited observation of search behaviors—primarily clicking—while more complex behaviors common in real-world contexts, such as pairwise comparisons (e.g., picking up and comparing two items) or placeholder strategies

(e.g., holding one item while comparing it to others), cannot be captured. These behaviors likely play a critical role in naturalistic decision-making.

Lastly, the inherent simplicity of our experimental tasks limits the extent to which our findings can be applied to real-world decision environments. In our studies, option values were often represented numerically and could be directly calculated. However, real-life decisions frequently involve non-additive and qualitative attributes that are difficult to evaluate side by side. Incorporating such qualitative features in future research may enhance ecological validity and provide deeper insight into how attention to different attributes shapes search behavior.

Acknowledgments

I am deeply grateful to my supervisor, Prof. Dr. Benjamin Scheibehenne, for his support, insightful guidance, and constructive criticism throughout my doctoral journey. This thesis would not have been possible without his mentorship.

My sincere thanks also go to Prof. Dr. Christof Weinhardt for serving as my second supervisor and for his helpful feedback, as well as to Prof. Dr. Ulrich Ebner-Priemer for being a member of my examination committee and for his critical reading of this work.

I'm also fortunate to have been part of the interdisciplinary PhD program KD²School, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation). Being in this program gave me the opportunity to engage with research beyond my primary field. I am thankful to the members, associates, and supervisors of KD²School for their support and the enriching discussions that made interdisciplinary research both challenging and rewarding.

A heartfelt thanks goes to the team in the secretary's office—especially Sabine Schneider and Claudia Ritter—for their kind assistance with administrative matters and official documents.

I'm grateful to my colleagues at the Cognition & Consumer Behavior Lab, Hannah Seidler and Tianwei Liu, for their positive influence not only at work but also in my life. Tianwei, thanks for your valuable advice and for always encouraging me to think more deeply. You're a friend I didn't see coming, but one I'm truly glad to have. Thank you, Hannah, for genuinely caring about me during many tough times. My journey would not have been nearly as meaningful or enjoyable without meeting you both.

Lastly, my deepest appreciation goes to my parents, who endured a lifetime of hardship and wars, yet always gave everything to ensure we had a better future. I feel incredibly fortunate and proud to call them my parents, and to know they'll always stand by me no matter what. To my big brother—thank you for always being there with your steady support.

Reference

- Appelhoff, S., Hertwig, R., & Spitzer, B. (2022). Control over sampling boosts numerical evidence processing in human decisions from experience. *Cerebral Cortex*, 33(1), 207–221.
<https://doi.org/10.1093/cercor/bhac062>
- Ástebro, T., Mata, J., & Santos-Pinto, L. (2015). Skewness seeking: risk loving, optimism or overweighting of small probabilities? *Theory and Decision*, 78(2), 189–208.
<https://doi.org/10.1007/s11238-014-9417-4>
- Audibert, J.-Y., Munos, R., & Szepesvári, C. (2009). Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19), 1876–1902.
<https://doi.org/10.1016/j.tcs.2009.01.016>
- Aziz, A., Li, H., & Telang, R. (2023). The Consequences of Rating Inflation on Platforms: Evidence from a Quasi-Experiment. *Information Systems Research*, 34(2), 590–608.
<https://doi.org/10.1287/isre.2022.1134>
- Azzopardi, L. (2021). Cognitive biases in search: A review and reflection of cognitive biases in information retrieval. *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*, 27–37. <https://doi.org/10.1145/3406522.3446023>
- Bartumeus, F., da Luz, M. G. E., Viswanathan, G. M., & Catalan, J. (2005). ANIMAL SEARCH STRATEGIES: A QUANTITATIVE RANDOM-WALK ANALYSIS. *Ecology*, 86(11), 3078–3087. <https://doi.org/10.1890/04-1806>
- Baumann, C., Schlegelmilch, R., & von Helversen, B. (2022). Adaptive behavior in optimal sequential search. *Journal of Experimental Psychology: General*.
<https://doi.org/10.1037/xge0001287>
- Bella-Fernández, M., Suero Suñé, M., & Gil-Gómez de Liaño, B. (2022). Foraging behavior in visual search: A review of theoretical and mathematical models in humans and animals. *Psychological Research*, 86(2), 331–349. <https://doi.org/10.1007/s00426-021-01499-1>

- Besedeš, T., Deck, C., Sarangi, S., & Shor, M. (2015). Reducing Choice Overload without Reducing Choices. *Review of Economics and Statistics*, 97(4), 793–802.
https://doi.org/10.1162/REST_a_00506
- Bhatnagar, R., & Orquin, J. L. (2022). A meta-analysis on the effect of visual attention on choice. *Journal of Experimental Psychology: General*, 151(10), 2265–2283.
<https://doi.org/10.1037/xge0001204>
- Biella, M., & Hütter, M. (2024). Navigating the social environment: linking motivations, impressions, and behaviors via a sampling approach to trustworthiness. *Personality and Social Psychology Bulletin*, 1461672241273237. <https://doi.org/10.1177/01461672241273237>
- Bollen, D., Knijnenburg, B. P., Willemsen, M. C., & Graus, M. (2010). Understanding choice overload in recommender systems. *Proceedings of the Fourth ACM Conference on Recommender Systems - RecSys '10*, 63. <https://doi.org/10.1145/1864708.1864724>
- Bürkner, P.-C. (2021). Bayesian Item Response Modeling in R withbrms and Stan. *Journal of Statistical Software*, 100(5). <https://doi.org/10.18637/jss.v100.i05>
- Buturak, G., & Evren, Ö. (2017). Choice overload and asymmetric regret. *Theoretical Economics*, 12(3), 1029–1056. <https://doi.org/10.3982/TE2037>
- Castellano, S. (2019). Putting mechanisms in foraging theory: the role of computational mechanisms in optimal decision making. *Animal Behaviour*, 153, 159–169.
<https://doi.org/10.1016/j.anbehav.2019.04.015>
- Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2), 129–136. [https://doi.org/10.1016/0040-5809\(76\)90040-X](https://doi.org/10.1016/0040-5809(76)90040-X)
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97.
<https://doi.org/10.1016/j.jbef.2015.12.001>
- Cohen, D., & Teodorescu, K. (2022). On the effect of perceived patterns in decisions from sampling. *Decision*, 9(1), 21–42. <https://doi.org/10.1037/dec0000159>

- Dellaert, B. G. C., Baker, T., & Johnson, E. J. (2017). Partitioning sorted sets: overcoming choice overload while maintaining decision quality. *SSRN Electronic Journal*.
<https://doi.org/10.2139/ssrn.3082254>
- Descamps, A., Massoni, Ss., & Page, L. (2016). Knowing when to stop and make a choice, an experiment on optimal sequential sampling. *SSRN Electronic Journal*.
<https://doi.org/10.2139/ssrn.2883847>
- Desender, K., Boldt, A., & Yeung, N. (2018). Subjective confidence predicts information seeking in decision making. *Psychological Science*, 29(5), 761–778.
<https://doi.org/10.1177/0956797617744771>
- Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2018). On the Relationship between Cognitive Ability and Risk Preference. *The Journal of Economic Perspectives : A Journal of the American Economic Association*, 32(2), 115–134. <https://doi.org/10.1257/jep.32.2.115>
- Elabed, G., & Carter, M. R. (2015). Compound-risk aversion, ambiguity and the willingness to pay for microinsurance. *Journal of Economic Behavior & Organization*, 118, 150–166.
<https://doi.org/10.1016/j.jebo.2015.03.002>
- Everett, S. P., & Byrne, M. D. (2004). Unintended effects: Varying icon spacing changes users' visual search strategy. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 695–702. <https://doi.org/10.1145/985692.985780>
- Ferguson, A. S. (1922). Plato's simile of light (continued). part II. the allegory of the cave. *The Classical Quarterly*, 16(1), 15–28. <https://doi.org/10.1017/S0009838800001956>
- Fleischhut, N., Artinger, F. M., Olschewski, S., & Hertwig, R. (2022). Not all uncertainty is treated equally: Information search under social and nonsocial uncertainty. *Journal of Behavioral Decision Making*, 35(2). <https://doi.org/10.1002/bdm.2250>
- Garrett, T. A., & Sobel, R. S. (1999). Gamblers favor skewness, not risk: Further evidence from United States' lottery games. *Economics Letters*, 63(1), 85–90. [https://doi.org/10.1016/S0165-1765\(99\)00012-9](https://doi.org/10.1016/S0165-1765(99)00012-9)

- Gathergood, J., Loewenstein, G. F., Quispe-Torreblanca, E., & Stewart, N. (2019). Investor attention, reference points and the disposition effect. *SSRN Electronic Journal*.
<https://doi.org/10.2139/ssrn.3422790>
- Haines, N., Kvam, P. D., & Turner, B. (2022). *Explaining the description-experience gap in risky decision-making: Learning and memory retention during experience as causal mechanisms*.
<https://doi.org/10.31234/osf.io/4u5j3>
- Hanna, M. G., Pantanowitz, L., Jackson, B., Palmer, O., Visweswaran, S., Pantanowitz, J., Deebajah, M., & Rashidi, H. H. (2025). Ethical and bias considerations in artificial intelligence/machine learning. *Modern Pathology*, 38(3), 100686.
<https://doi.org/10.1016/j.modpat.2024.100686>
- Hein, A. M., Altshuler, D. L., Cade, D. E., Liao, J. C., Martin, B. T., & Taylor, G. K. (2020). An algorithmic approach to natural behavior. *Current Biology*, 30(11), R663–R675.
<https://doi.org/10.1016/j.cub.2020.04.018>
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539.
<https://doi.org/10.1111/j.0956-7976.2004.00715.x>
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: why small samples? *Cognition*, 115(2), 225–237. <https://doi.org/10.1016/j.cognition.2009.12.009>
- Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience. Do our patterns of sampling foreshadow our decisions? *Psychological Science*, 21(12), 1787–1792.
<https://doi.org/10.1177/0956797610387443>
- Hills, T. T., Noguchi, T., & Gibbert, M. (2013). Information overload or search-amplified risk? Set size and order effects on decisions from experience. *Psychonomic Bulletin & Review*, 20(5), 1023–1031. <https://doi.org/10.3758/s13423-013-0422-3>
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2010). The central executive as a search process: priming exploration and exploitation across domains. *Journal of Experimental Psychology: General*, 139(4), 590–609. <https://doi.org/10.1037/a0020666>

- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., & Cognitive Search Research Group. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19(1), 46–54. <https://doi.org/10.1016/j.tics.2014.10.004>
- Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, 30(1), 3–41. https://doi.org/10.1207/s15516709cog0000_50
- Hutchinson, J. M. C., Wilke, A., & Todd, P. M. (2008). Patch leaving in humans: can a generalist adapt its rules to dispersal of items across patches? *Animal Behaviour*, 75(4), 1331–1349. <https://doi.org/10.1016/j.anbehav.2007.09.006>
- Hu, B., & Simmons, J. P. (2025). Different methods elicit different belief distributions. *Journal of Experimental Psychology: General*, 154(2), 476–496. <https://doi.org/10.1037/xge0001655>
- Jing, L., & Cheo, R. (2013). House money effects, risk preferences and the public goods game. *Economics Letters*, 120(2), 310–313. <https://doi.org/10.1016/j.econlet.2013.04.037>
- Johnson, E. J., & Tversky, A. (1984). Representations of perceptions of risks. *Journal of Experimental Psychology: General*, 113(1), 55–70. <https://doi.org/10.1037/0096-3445.113.1.55>
- Jonas, E., Schulz-Hardt, S., Frey, D., & Thelen, N. (2001). Confirmation bias in sequential information search after preliminary decisions: an expansion of dissonance theoretical research on selective exposure to information. *Journal of Personality and Social Psychology*, 80(4), 557–571. <https://doi.org/10.1037//0022-3514.80.4.557>
- Jung, M. H., Moon, A., & Nelson, L. D. (2020). Overestimating the valuations and preferences of others. *Journal of Experimental Psychology: General*, 149(6), 1193–1214. <https://doi.org/10.1037/xge0000700>
- Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: a naïve sampling model of intuitive confidence intervals. *Psychological Review*, 114(3), 678–703. <https://doi.org/10.1037/0033-295X.114.3.678>
- Kadwe, P. P., Sklenar, A. M., Frankenstein, A. N., Levy, P. U., & Leshikar, E. D. (2022). The influence of memory on approach and avoidance decisions: Investigating the role of episodic memory in social decision making. *Cognition*, 225, 105072. <https://doi.org/10.1016/j.cognition.2022.105072>

- Kellen, D., Steiner, M. D., Davis-Stober, C. P., & Pappas, N. R. (2020). Modeling choice paradoxes under risk: From prospect theories to sampling-based accounts. *Cognitive Psychology*, 118, 101258. <https://doi.org/10.1016/j.cogpsych.2019.101258>
- Kinjo, K., & Ebina, T. (2015). Paradox of choice and consumer nonpurchase behavior. *AI & Society*, 30(2), 291–297. <https://doi.org/10.1007/s00146-014-0546-7>
- Kuleshov, V., & Precup, D. (2014). Algorithms for multi-armed bandit problems. *ArXiv*. <https://doi.org/10.48550/arxiv.1402.6028>
- Lee, J. D., & Seppelt, B. D. (2023). Design for Human-Automation and Human-Autonomous Systems. In S. Y. Nof (Ed.), *Springer handbook of automation* (pp. 433–455). Springer International Publishing. https://doi.org/10.1007/978-3-030-96729-1_19
- Leuker, C., Pachur, T., Hertwig, R., & Pleskac, T. J. (2019). Do people exploit risk–reward structures to simplify information processing in risky choice? *Journal of the Economic Science Association*, 5(1), 76–94. <https://doi.org/10.1007/s40881-019-00068-y>
- Levav, J., Reinholtz, N., & Lin, C. (2012). The Effect of Ordering Decisions by Choice-Set Size on Consumer Search. *The Journal of Consumer Research*, 39(3), 585–599. <https://doi.org/10.1086/664498>
- Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: on the “naivety” assumption in sampling explanations of judgment biases. *Psychological Review*, 118(2), 379–392. <https://doi.org/10.1037/a0023010>
- Loecher, M. (2021). The Perils of Misspecified Priors and Optional Stopping in Multi-Armed Bandits. *Frontiers in Artificial Intelligence*, 4, 715690. <https://doi.org/10.3389/frai.2021.715690>
- Long, X., Sun, J., Dai, H., Zhang, D., Zhang, J., Chen, Y., Hu, H., & Zhao, B. (2021). Choice Overload with Search Cost and Anticipated Regret: Theoretical Framework and Field Evidence. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3890056>
- Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014). Extreme Outcomes Sway Risky Decisions from Experience. *Journal of Behavioral Decision Making*, 27(2), 146–156. <https://doi.org/10.1002/bdm.1792>

- Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: is the description-experience gap in risky choice limited to rare events? *Plos One*, 6(6), e20262. <https://doi.org/10.1371/journal.pone.0020262>
- Mallpress, D. E. W., Fawcett, T. W., Houston, A. I., & McNamara, J. M. (2015). Risk attitudes in a changing environment: An evolutionary model of the fourfold pattern of risk preferences. *Psychological Review*, 122(2), 364–375. <https://doi.org/10.1037/a0038970>
- Malone, T., & Lusk, J. L. (2019). Mitigating choice overload: an experiment in the U.S. beer market. *Journal of Wine Economics*, 14(01), 48–70. <https://doi.org/10.1017/jwe.2018.34>
- Martinovici, A., Pieters, F. G. M. (Rik), & Erdem, T. (2021). Attention trajectories predict brand choice. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3393021>
- Mason, A., Ludvig, E. A., Spetch, M. L., & Madan, C. R. (2024). Rare and extreme outcomes in risky choice. *Psychonomic Bulletin & Review*, 31(3), 1301–1308. <https://doi.org/10.3758/s13423-023-02415-x>
- Mazyar, H., van den Berg, R., Seilheimer, R. L., & Ma, W. J. (2013). Independence is elusive: set size effects on encoding precision in visual search. *Journal of Vision*, 13(5). <https://doi.org/10.1167/13.5.8>
- Ma, I., Sanfey, A. G., & Ma, W. J. (2020). The social cost of gathering information for trust decisions. *Scientific Reports*, 10(1), 14073. <https://doi.org/10.1038/s41598-020-69766-6>
- McShane, B. B., & Böckenholt, U. (2018). Multilevel Multivariate Meta-analysis with Application to Choice Overload. *Psychometrika*, 83(1), 255–271. <https://doi.org/10.1007/s11336-017-9571-z>
- Mechera-Ostrovsky, T., Heinke, S., Andraszewicz, S., & Rieskamp, J. (2022). Cognitive abilities affect decision errors but not risk preferences: A meta-analysis. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-021-02053-1>
- Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3), 191–215. <https://doi.org/10.1037/dec0000033>

- Mejía, V. D., Aurier, P., & Huaman-Ramirez, R. (2021). Disentangling the respective impacts of assortment size and alignability on perceived assortment variety. *Journal of Retailing and Consumer Services*, 59, 102386. <https://doi.org/10.1016/j.jretconser.2020.102386>
- Mellers, B. A., McCoy, J. P., Lu, L., & Tetlock, P. E. (2023). Human and Algorithmic Predictions in Geopolitical Forecasting: Quantifying Uncertainty in Hard-to-Quantify Domains. *Perspectives on Psychological Science*, 17456916231185340. <https://doi.org/10.1177/17456916231185339>
- Nobel, N. (2021). Recommender Systems: Friend (of Choice) or Foe? A Large-scale Field Experiment in Online Shopping Platforms. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3932914>
- Oberholzer, Y., Lindskog, M., & Scheibehenne, B. (2021). *The influence of the place value system on symbolic number perception*. <https://doi.org/10.31234/osf.io/9f3z2>
- Olschewski, S., Luckman, A., Mason, A., Ludvig, E. A., & Konstantinidis, E. (2022). *The Future of Decisions from Experience: Connecting Real-World Decision Problems to Cognitive Processes*. <https://doi.org/10.31234/osf.io/7csbt>
- Olschewski, S., Newell, B. R., Oberholzer, Y., & Scheibehenne, B. (2021). Valuation and estimation from experience. *Journal of Behavioral Decision Making*, 34(5), 729–741. <https://doi.org/10.1002/bdm.2241>
- Olschewski, S., & Scheibehenne, B. (2021). *What's in a Sample? How Sampling Information Affects Epistemic Uncertainty and Risk-Taking*.
- Olschewski, S., & Scheibehenne, B. (2023). *What's in a sample? epistemic uncertainty and metacognitive awareness in risk taking*. <https://doi.org/10.31234/osf.io/hf8dm>
- Olschewski, S., Spektor, M. S., & Le Mens, G. (2024). Frequent winners explain apparent skewness preferences in experience-based decisions. *Proceedings of the National Academy of Sciences of the United States of America*, 121(12), e2317751121. <https://doi.org/10.1073/pnas.2317751121>
- Oppewal, H., & Koelemeijer, K. (2005). More choice is better: Effects of assortment size and composition on assortment evaluation. *International Journal of Research in Marketing*, 22(1), 45–60. <https://doi.org/10.1016/j.ijresmar.2004.03.002>

- Orquin, J. L., Lahm, E. S., & Stojić, H. (2021). The visual environment and attention in decision making. *Psychological Bulletin*, 147(6), 597–617. <https://doi.org/10.1037/bul0000328>
- Pachur, T. (2022). Strategy selection in decisions from givens: Deciding at a glance? *Cognitive Psychology*, 136, 101483. <https://doi.org/10.1016/j.cogpsych.2022.101483>
- Park, J.-Y., & Jang, S. (Shawn). (2013). Confused by too many choices? Choice overload in tourism. *Tourism Management*, 35, 1–12. <https://doi.org/10.1016/j.tourman.2012.05.004>
- Pleskac, T. J., Yu, S., Grunevski, S., & Liu, T. (2021). *Attention biases preferential choice by enhancing an option's value*. <https://doi.org/10.31234/osf.io/n3ghb>
- Polman, E. (2012). Effects of self-other decision making on regulatory focus and choice overload. *Journal of Personality and Social Psychology*, 102(5), 980–993. <https://doi.org/10.1037/a0026966>
- Prat-Carrabin, A., & Woodford, M. (2022). Efficient coding of numbers explains decision bias and noise. *Nature Human Behaviour*, 6(8), 1142–1152. <https://doi.org/10.1038/s41562-022-01352-4>
- Prolific. (2024). *Prolific*. Prolific. <https://www.prolific.com>
- Qahri-Saremi, H., & Montazemi, A. R. (2023). Negativity bias in the diagnosticity of online review content: the effects of consumers' prior experience and need for cognition. *European Journal of Information Systems*, 32(4), 717–734. <https://doi.org/10.1080/0960085X.2022.2041372>
- Quandt, J., Figner, B., Holland, R. W., & Veling, H. (2022). Confidence in evaluations and value-based decisions reflects variation in experienced values. *Journal of Experimental Psychology: General*, 151(4), 820–836. <https://doi.org/10.1037/xge0001102>
- Qu, C., DeWind, N. K., & Brannon, E. M. (2022). Increasing entropy reduces perceived numerosity throughout the lifespan. *Cognition*, 225, 105096. <https://doi.org/10.1016/j.cognition.2022.105096>
- Rahman, M. A. (2025). *Comparative Analysis of Run-Length Encoding Techniques: Standard, Modified, 2D, and Bitwise Approaches for Efficient Data Compression with Numerical Example*.
- Rajsic, J., Wilson, D. E., & Pratt, J. (2015). Confirmation bias in visual search. *Journal of Experimental Psychology. Human Perception and Performance*, 41(5), 1353–1364. <https://doi.org/10.1037/xhp0000090>

- Reed, D. D., Reed, F. D. D., Chok, J., & Brozyna, G. A. (2011). The “tyranny of choice”: choice overload as a possible instance of effort discounting. *The Psychological Record*, 61(4), 547–560. <https://doi.org/10.1007/BF03395776>
- Reutskaja, E., Cheek, N. N., Iyengar, S., & Schwartz, B. (2022). Choice Deprivation, Choice Overload, and Satisfaction with Choices Across Six Nations. *Journal of International Marketing*, 30(3), 18–34. <https://doi.org/10.1177/1069031X211073821>
- Rosenbaum, D., Glickman, M., Fleming, S. M., & Usher, M. (2022). The Cognition/Metacognition Trade-Off. *Psychological Science*, 33(4), 613–628. <https://doi.org/10.1177/09567976211043428>
- R Core Team. (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna. <https://www.R-project.org>
- Sahakian, A., Gayet, S., Paffen, C. L. E., & Van der Stigchel, S. (2023). Mountains of memory in a sea of uncertainty: Sampling the external world despite useful information in visual working memory. *Cognition*, 234, 105381. <https://doi.org/10.1016/j.cognition.2023.105381>
- Salatino, A., Prével, A., Caspar, E., & Bue, S. L. (2025). Influence of AI behavior on human moral decisions, agency, and responsibility. *Scientific Reports*, 15(1), 12329. <https://doi.org/10.1038/s41598-025-95587-6>
- Saltsman, T. L., Seery, M. D., Ward, D. E., Lamarche, V. M., & Kondrak, C. L. (2021). Is satisficing really satisfying? Satisficers exhibit greater threat than maximizers during choice overload. *Psychophysiology*, 58(1), e13705. <https://doi.org/10.1111/psyp.13705>
- Scheibehenne, B., Greifeneder, R., & Todd, P. M. (2010). Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload. *The Journal of Consumer Research*, 37(3), 409–425. <https://doi.org/10.1086/651235>
- Schmidt, U., & Zank, H. (2008). Risk aversion in cumulative prospect theory. *Management Science*, 54(1), 208–216. <https://doi.org/10.1287/mnsc.1070.0762>
- Schulz-Hardt, S., Frey, D., Lüthgens, C., & Moscovici, S. (2000). Biased information search in group decision making. *Journal of Personality and Social Psychology*, 78(4), 655–669. <https://doi.org/10.1037/0022-3514.78.4.655>

- Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: happiness is a matter of choice. *Journal of Personality and Social Psychology*, 83(5), 1178–1197. <https://doi.org/10.1037//0022-3514.83.5.1178>
- Sethuraman, R., Gázquez-Abad, J. C., & Martínez-López, F. J. (2022). The effect of retail assortment size on perceptions, choice, and sales: Review and research directions. *Journal of Retailing*, 98(1), 24–45. <https://doi.org/10.1016/j.jretai.2022.01.001>
- Skreta, V., & Veldkamp, L. (2009). Ratings shopping and asset complexity: A theory of ratings inflation. *Journal of Monetary Economics*, 56(5), 678–695. <https://doi.org/10.1016/j.jmoneco.2009.04.006>
- Smith, S. M., & Krajbich, I. (2019). Gaze amplifies value in decision making. *Psychological Science*, 30(1), 116–128. <https://doi.org/10.1177/0956797618810521>
- Soll, J. B., Palley, A., Klayman, J., & Moore, D. (2019). Overconfidence in probability distributions: people know they don't know but they don't know what to do about it. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3435793>
- Song, M., Bnaya, Z., & Ma, W. J. (2019). Sources of suboptimality in a minimalistic explore-exploit task. *Nature Human Behaviour*, 3(4), 361–368. <https://doi.org/10.1038/s41562-018-0526-x>
- Speekenbrink, M. (2022). Chasing unknown bandits: uncertainty guidance in learning and decision making. *Current Directions in Psychological Science*, 31(5), 419–427. <https://doi.org/10.1177/09637214221105051>
- Spektor, M. S., & Wulff, D. U. (2023). *Strategic information search in decisions from experience*. <https://doi.org/10.31234/osf.io/7ru58>
- Stewart, N., Brown, G. D. A., & Chater, N. (2005). Absolute identification by relative judgment. *Psychological Review*, 112(4), 881–911. <https://doi.org/10.1037/0033-295X.112.4.881>
- Stewart, N. (2009). Decision by sampling: the role of the decision environment in risky choice. *Quarterly Journal of Experimental Psychology*, 62(6), 1041–1062. <https://doi.org/10.1080/17470210902747112>

- Sunstein, C. R. (2003). Terrorism and probability neglect. *Journal of Risk and Uncertainty*, 26(2–3), 121–136. <https://doi.org/10.1023/A:1024111006336>
- Suzuki, M., & Yamamoto, Y. (2021). Characterizing the influence of confirmation bias on web search behavior. *Frontiers in Psychology*, 12, 771948. <https://doi.org/10.3389/fpsyg.2021.771948>
- Townsend, C., & Kahn, B. E. (2014). The “Visual Preference Heuristic”: The Influence of Visual versus Verbal Depiction on Assortment Processing, Perceived Variety, and Choice Overload. *The Journal of Consumer Research*, 40(5), 993–1015. <https://doi.org/10.1086/673521>
- van Opheusden, B., Kuperwajs, I., Galbiati, G., Bnaya, Z., Li, Y., & Ma, W. J. (2023). Expertise increases planning depth in human gameplay. *Nature*, 618(7967), 1000–1005. <https://doi.org/10.1038/s41586-023-06124-2>
- von Helversen, B., Mata, R., Samanez-Larkin, G. R., & Wilke, A. (2018). Foraging, exploration, or search? On the (lack of) convergent validity between three behavioral paradigms. *Evolutionary Behavioral Sciences*, 12(3), 152–162. <https://doi.org/10.1037/ebs0000121>
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599–637. <https://doi.org/10.1111/cogs.12101>
- Walters, D. J., Ülkümen, G., Tannenbaum, D., Erner, C., & Fox, C. R. (2023). Investor behavior under epistemic vs. aleatory uncertainty. *Management Science*, 69(5), 2761–2777. <https://doi.org/10.1287/mnsc.2022.4489>
- Wang, F., Wang, M., Zheng, Y., Jin, J., & Pan, Y. (2021). Consumer vigilance and choice overload in online shopping. *International Journal of Electronic Commerce*, 25(3), 364–390. <https://doi.org/10.1080/10864415.2021.1943189>
- Wang, R., & Sahin, O. (2017). The impact of consumer search cost on assortment planning and pricing. *Management Science*. <https://doi.org/10.1287/mnsc.2017.2790>
- Weber, E. U. (2010). Risk attitude and preference. *Wiley Interdisciplinary Reviews. Cognitive Science*, 1(1), 79–88. <https://doi.org/10.1002/wcs.5>

- Wulff, D. U., Hills, T. T., & Hertwig, R. (2015). How short- and long-run aspirations impact search and choice in decisions from experience. *Cognition*, 144, 29–37.
<https://doi.org/10.1016/j.cognition.2015.07.006>
- Wu, P. F. (2013). In search of negativity bias: an empirical study of perceived helpfulness of online reviews. *Psychology and Marketing*, 30(11), 971–984. <https://doi.org/10.1002/mar.20660>
- Yeung, N., & Summerfield, C. (2012). Metacognition in human decision-making: confidence and error monitoring. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 367(1594), 1310–1321. <https://doi.org/10.1098/rstb.2011.0416>
- Yun, J. T., & Duff, B. R. L. (2017). Is utilizing themes an effective scheme? Choice overload and categorization effects within an extensive online choice environment. *Computers in Human Behavior*, 74, 205–214. <https://doi.org/10.1016/j.chb.2017.04.038>
- Zhou, T., Wang, Y., Yan, & Tan, Y. (2020). Spoiled for Choice? Personalized Recommendation for Healthcare Decisions: A Multi-Armed Bandit Approach. *ArXiv*.
<https://doi.org/10.48550/arxiv.2009.06108>
- Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2020). The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*, 127(5), 719–748.
<https://doi.org/10.1037/rev0000190>
- Zilker, V. (2022). Stronger attentional biases can be linked to higher reward rate in preferential choice. *Cognition*, 225, 105095. <https://doi.org/10.1016/j.cognition.2022.105095>
- Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: not enough or too much? *Marketing Science*, 22(4), 503–519.
<https://doi.org/10.1287/mksc.22.4.503.24909>

Appendix

Supplementary Material Chapter 2

Experiment Instructions, Instruction Checks, and Attention Check

Instruction about the task:

Welcome

In a city, there are several restaurants, each attracting different numbers of customers on average. As an investor, you seek the busiest restaurant to invest in, one that attracts more customers on average than the rest.

A national report shows that only **20%** of restaurants in a city are considered busy compared to other local restaurants.

To identify a busy restaurant, you can compare its daily customers count to other restaurants in the city.

Each restaurant attracts a different average number of customers, but the daily customers count fluctuates around their average.

This means that sometimes a busy restaurant on a slow day will have fewer customers than an ordinary restaurant on a good day.

You can watch the daily customers count of each restaurant in a city until you're confident enough to invest.

However, you can only invest in one restaurant per city.

The bonus you receive will depend on the **average customer count** of the restaurant you invested in.

A restaurant attracting more customers on average will make you more money.

Additionally, we will request you to provide an estimate of the average number of customers expected to visit your restaurant.

You will receive a reward of 0.1 euro each time your estimate falls within 10% of the actual average.

In each round, you will come to a new city to find the best restaurant in that city.

Please note that each city has a **different population size** so the best restaurant in a small city might be less busy than a normal restaurant in a large city.

The number of restaurants in each city also varies, but it is unrelated to the population size.

Thus, to maximize your payout, it is crucial not to compare restaurants across different cities. Instead, focus on identifying the best restaurant in each city individually.

Please read this instruction carefully because we will ask some questions to test how well you understand the task.

When you are ready, click **Next** to see an **example trial** of the task.

Instruction checks:

Question 1:

There are 100 restaurants in city A. Based on the national report, how many restaurants in this city are busier than the rest?

Answers:

- A. 15
- B. 10
- C. 20
- D. 30

Question 2:

My reward in this task is determined by:

Answers:

- A. The number of customers visited my invested restaurant in 1 random day

- B. The number of customers visited the best restaurant in the city
- C. The average number of customers visited my invested restaurant
- D. The smallest number of customers visited my invested restaurant

Question 3:

Can a busy restaurant has less customers than a normal restaurant?

Answers:

- A. No, a busy restaurant always has more customers
- B. Yes, a busy restaurant can have less customers on average, but have more customers on a good day than a normal restaurant
- C. Yes, a busy restaurant can have less customers on a slow day, but have more customers on average than a normal restaurant
- D. None of the answers above are correct

Question 4:

The observed customers count of a restaurant in 5 random days are: 30, 40, 40, 30, 20. Which number below is closest to the average of this number sequence?

Answers:

- A. 10
- B. 20
- C. 30
- D. 60

Note: The correct answer is C for all four questions. Only participants who answer more than two questions correctly will be included in the data analysis.

Figure 6.2.1

A screenshot from an example trial of the experiment in chapter 2

Example trial: There are 35 restaurants

20% of restaurants in city 1 are busier than the rest.

Click on any restaurant button to observe how many customers visited that restaurant in one random day. You can do this until you feel confident enough to invest.

Restaurant 1

Restaurant 2

Restaurant 3

Restaurant 4

Restaurant 5

Restaurant 6

Restaurant 7

Restaurant 8

Restaurant 9

Restaurant 10

Restaurant 11

Restaurant 12

Restaurant 13

Restaurant 14

Restaurant 15

Restaurant 16

Restaurant 17

Restaurant 18

Restaurant 19

Restaurant 20

Restaurant 21

Restaurant 22

Restaurant 23

Restaurant 24

Restaurant 25

Restaurant 26

Restaurant 27

Restaurant 28

Restaurant 29

Restaurant 30

Restaurant 31

Restaurant 32

Restaurant 33

Restaurant 34

Restaurant 35

Choose a restaurant to invest in:

▼

How many customers do you expect to visit your restaurant on average?

Click Next to submit your investment. You cannot change your choice after this.

Next

Analysis Results

Table 6.2.1

Bayesian Model Results for EADQ in low-noise condition

Fixed Effect	Estimate	SE	95% CI Lower	95% CI Upper
Intercept	-0.85	0.22	-1.28	-0.41
Search breadth	2.17	0.27	1.63	2.72
Set Size = 10	-0.23	0.31	-0.84	0.37
Set Size = 20	-0.17	0.29	-0.75	0.41
Set Size = 40	-0.35	0.30	-0.94	0.23
Search breadth \times Set Size = 10	-0.35	0.38	-1.11	0.40
Search breadth \times Set Size = 20	-1.01	0.38	-1.76	-0.27
Search breadth \times Set Size = 40	-0.93	0.40	-1.71	-0.15

Table 6.2.2*Bayesian Model Results for EADQ in high-noise condition*

Fixed Effect	Estimate	SE	95% CI Lower	95% CI Upper
Intercept	-0.66	0.17	-1.00	-0.33
Search breadth	1.75	0.22	1.31	2.19
Set Size = 10	-0.20	0.19	-0.56	0.17
Set Size = 20	-0.18	0.19	-0.54	0.19
Set Size = 40	-0.14	0.20	-0.52	0.26
Search breadth \times Set Size = 10	-0.15	0.26	-0.65	0.35
Search breadth \times Set Size = 20	-0.64	0.26	-1.16	-0.12
Search breadth \times Set Size = 40	-1.17	0.27	-1.70	-0.64

Table 6.2.3

The effect of number of samples on chosen option, set size, and noise on estimation accuracy

Fixed Effect	Estimate	SE	95% CI Lower	95% CI Upper
Intercept	-2.31	0.56	-3.40	-1.21
Number of samples on chosen option	0.04	0.01	0.02	0.06
Set Size = 10	-0.05	0.71	-1.42	1.34
Set Size = 20	-0.74	0.71	-2.14	0.65
Set Size = 40	-1.22	0.71	-2.63	0.17
Low-noise	1.16	0.50	0.19	2.14

Supplementary Material Chapter 3

Experiment Instructions, Instruction Checks, and Attention Check

Instruction about the task:

Welcome!

Here, you'll find various slot machines in a casino. Your objective is to identify the machine with the **highest expected value**.

Each slot machine has only two **outcomes**: you either win or lose.

Each slot machine has its own unique winning reward, which can range from very large to very small. However, if you lose, you will always receive nothing. Additionally, each machine has its own **winning probabilities**. Some machines have very rare wins, while others have more common wins.

The machine with the highest expected value is the one that consistently delivers the **highest and most frequent winning outcomes**.

For instance, consider Machine A, which has a 50% chance of winning 50 GBP and a 50% chance of not winning anything.

Machine B offers a 50% chance of winning 20 GBP and a 50% chance of winning nothing.

Machine A is deemed superior as it boasts a higher **expected value**.

You can calculate the expected value by **multiplying the winning amount by the probability of winning** (how frequently you encounter the winning amount). For instance, Machine A has an expected value of $50 * 50\% = 25$, while Machine B has an expected value of $20 * 50\% = 10$. Therefore, Machine A has higher expected value.

You can only choose one machine each trial, and the casino keeps the winning probabilities secret. However, we will provide you with a number of samples to observe the outcomes of the machines.

Utilize these samples to learn more about the machines and select the one you believe has the highest expected value. Please make sure to use all the samples before moving on to the next round.

The **expected value** of your chosen slot machine is your potential reward for that round. We will **randomly select a participant** to receive their potential reward from a **random trial**, providing everyone with a chance to win up to 8 GBP.

Please read this instruction carefully because we will ask some questions to test how well you understand the task. When you are ready, click **Next** to see an **example trial** of the task.

Instruction checks:

Question 1:

What is the main objective of the casino task?

- A) To lose money
- B) To observe slot machines
- C) To identify the machine with the highest expected value
- D) To keep winning probabilities secret

Question 2:

What are the two possible outcomes for each slot machine?

- A) Win or win
- B) Lose or lose
- C) Win something or gain nothing
- D) All of the above

Question 3:

Who might get their potential reward in this task?

- A) The person with the lowest score
- B) The person with the highest score
- C) A random person
- D) The person with an average score

Note: The correct answer is C for all three questions. Only participants who answer more than two out of three questions correctly will be included in the data analysis.

Figure 6.3.1

A screenshot from a trial of the experiment in chapter 3

Round: 1/8

[Click](#) on any button to observe an outcome of a slot machine. You have 20 samples in total.

The number of samples left: 20

<div>Machine 1</div> <div></div>	<div>Machine 2</div> <div></div>	<div>Machine 3</div> <div></div>	<div>Machine 4</div> <div></div>
<div>Machine 5</div> <div></div>	<div>Machine 6</div> <div></div>	<div>Machine 7</div> <div></div>	<div>Machine 8</div> <div></div>
<div>Machine 9</div> <div></div>	<div>Machine 10</div> <div></div>		

Choose a slot machine to play for real:

----- ▾

How many times do you expect to win if you get 100 draws from this slot machine?

You still have samples left. Use all your samples to move on to the next round

How to calculate entropy

Let c_1, c_2, \dots, c_k be the number of samples (or selections) spent on each option in the search task, where $k = 10$ is the total number of available options.

First, we calculate the proportion of samples allocated to each option:

$$p_i = c_i / (c_1 + c_2 + \dots + c_k)$$

Then, we compute the entropy H of the search sequence using the Shannon entropy formula:

$$H = - \sum p_i \times \log_2(p_i)$$

This formula measures how evenly the samples are distributed across the options so that: If samples are concentrated on one option (e.g., all on one), entropy is low. Meanwhile, if samples are spread evenly across options, entropy is high.

Where:

p_i is the relative frequency (proportion) of samples for option i ,

- The sum is taken over all options,
- \log_2 is the logarithm base 2,
- H represents the uncertainty or diversity in the search behavior.

How to calculate average search length

Given a search sequence consisting of a series of samples from available options, the sequence is divided into R consecutive runs. A run is a group of consecutive repeated selections of the same option.

Let the lengths of these runs be denoted as l_1, l_2, \dots, l_R .

Then the average search length, denoted as \bar{L} , is defined as:

$$\bar{L} = (l_1 + l_2 + \dots + l_R) / R$$

In other words, you add up the lengths of all runs and divide by the total number of runs.

Supplementary Material Chapter 4


Experiment Instructions, Instruction Checks, and Attention Check

Instruction about the task:

Welcome to the Product Review Estimation Task!

In this task, you will evaluate various products based on their star ratings and reviews left by previous customers.

In each trial, you will see a graph showing how previous customers have rated a product using a **star rating** system, ranging from 1 star (worst) to 5 stars (best).



Here, 49% of customers have rated this product 5 stars.

You can click on any button under each bin of the graph to sample a review from people who rated that star for the product.

Each time you click on a button, a reviewer's **positivity score** will appear.

The positivity score reflects the tone of the review's wording about the product.

It ranges from 0 (no positive wording) to 100 (completely positive review).

People who rated the same star might still have different feelings about the product, so the positivity scores of the same star rating still vary.



Reviewer's Positivity Score: 83.87

After clicking on the 5-star button, a reviewer's positivity score will show up below.

Additionally, positivity scores can vary across different trials even when the star ratings might seem similar.

For example, in some trials, products may have many 5-star ratings but still have lower positivity scores than previous products if the reviews are not enthusiastic about them.

Your goal is to estimate the average positivity score of the product by sampling the star ratings.

However, you have a limited number of samples available, so use them wisely.

Your accuracy in estimating the average positivity score will determine your potential reward for that round, which can range from 2 to 8 GBP.

We will randomly select a participant to receive their potential reward from a random trial, giving everyone a chance to win a bonus.

Please read these instructions carefully, as we will ask some questions to test your understanding of the task.

When you are ready, click **Next** to see a **practice trial** of the task.

Instruction checks:

Question 1: What does the star ratings tell you about the product?

- A) The total number of reviews received.
- B) The average price range of the product.
- C) The user ratings of the product.
- D) The geographical distribution of users.

Question 2: How is the positivity score defined in the context of the reviews?

- A) The length of each review in characters.
- B) A score ranging from 0 to 100, indicating the degree of negativity in review wordings.
- C) A score ranging from 0 to 100, indicating the degree of positivity in review wordings.
- D) The number of stars given in each review.

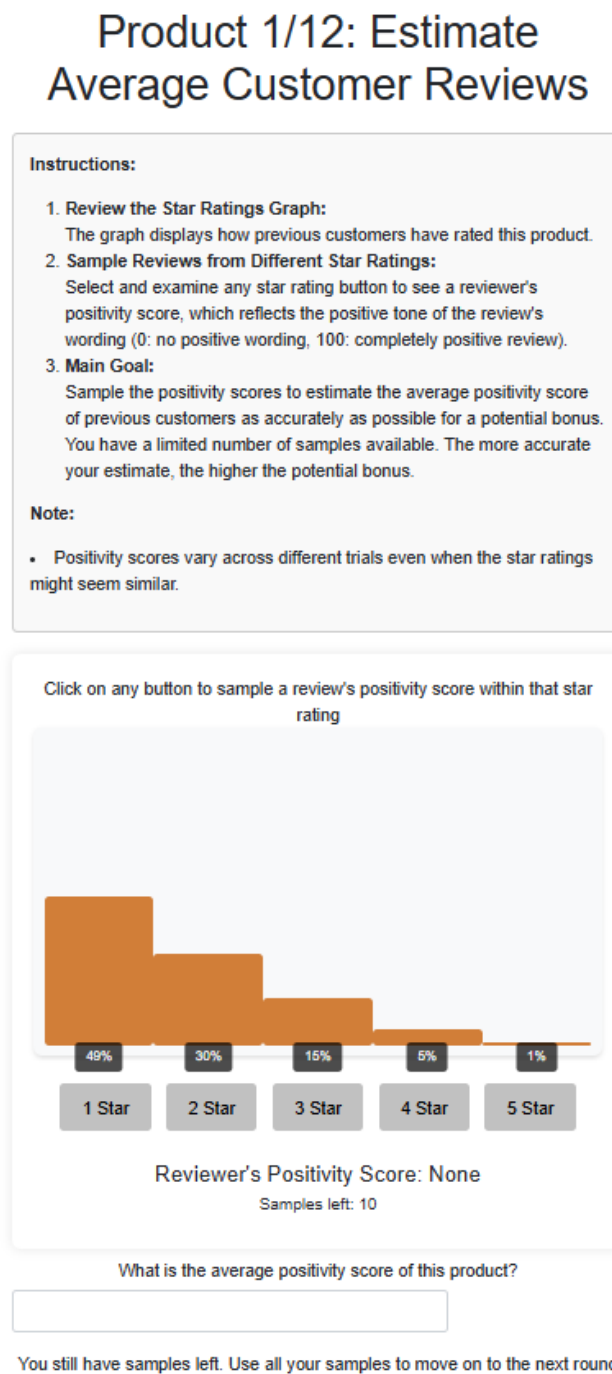
Question 3: What is the main goal of sampling the positivity score according to the instructions?

- A) To estimate the number of positive reviews received.
- B) To determine the product's price range.
- C) To estimate the average positivity score of previous users for a potential bonus.
- D) To compare the product with competitors.

Note: The correct answer is C for all three questions. Only participants who answer more than two out of three questions correctly will be included in the data analysis.

Figure 6.4.1

A screenshot from a trial of the experiment in chapter 4



Additional Analysis Chapter 4

Proof that stratified sampling results in an unbiased and efficient estimate of the true mean.

There are 5 bins i with their μ value: $\mu_1, \mu_2, \mu_3, \dots, \mu_5$

Each bin i has their probability PMF: $p_1, p_2, p_3 \dots p_5$

The experienced mean of the distribution can be calculated by:

$$\mu = \sum p_i \mu_i$$

After collecting the samples, the experienced mean $\hat{\mu}$ is computed as the average of all the drawn samples. Suppose we draw n_i samples from bin i , the experienced mean after N samples is:

$$\hat{\mu} = \sum \frac{n_i \hat{\mu}_i}{N}$$

Proof that stratified sampling is unbiased:

If $\hat{\mu}$ is an unbiased estimator of μ then:

$$\hat{\mu} = \mu = E(\hat{\mu})$$

$$= E\left(\sum p_i \hat{\mu}_i\right)$$

$$= \sum E(p_i \hat{\mu}_i)$$

Using stratified sampling means that distributing n_i so that $\frac{n_i}{N} \sim p_i$ then:

$$\lim_{\frac{n_i}{N} \rightarrow p_i} \sum E\left(\frac{n_i}{N} \hat{\mu}_i\right) = \sum E(p_i \hat{\mu}_i) = \mu$$

Proof that stratified sampling is efficient:

We aim to minimize the variance of the experienced mean $\text{Var}(\hat{\mu})$ for an efficient estimator:

$$\text{Var}(\hat{\mu}) = \text{Var}(\sum (\hat{\mu}_i \hat{p}_i))$$

We consider μ_i and p_i are independent or uncorrelated for simplicity.

$$\text{Var}(\hat{\mu}) = \sum \text{Var}(\hat{\mu}_i \hat{p}_i) = \sum \hat{p}_i^2 \text{Var}(\hat{\mu}_i) = \sum \left(\frac{n_i}{N}\right)^2 \text{Var}(\hat{\mu}_i)$$

We also have $\text{Var}(\hat{\mu}_i) = \frac{\sigma_i^2}{n_i}$

$$= \sum \left(\frac{n_i}{N}\right)^2 \frac{\sigma_i^2}{n_i} = \sum \frac{n_i}{N^2} \sigma_i^2 = \frac{1}{N} \sum \frac{n_i}{N} \sigma_i^2$$

where:

- $\text{Var}(\hat{\mu}_i)$ is the variance of the experienced mean in bin i
- \hat{p}_i is $\frac{n_i}{N}$, the proportion of samples taken from bin i ,
- N is the total number of samples, such that $\sum n_i = N$.

As N is fixed so the efficient sampling strategy to distribute n_i to minimize $\text{Var}(\hat{\mu}_i)$ is:

$$n_i \propto p_i \cdot \sigma_i^2$$

However, given that participants have no information about σ_i^2 , the efficient strategy is reduced to:

$$n_i \propto p_i$$

It is important to note that under uniform distribution conditions, there are two equally efficient sampling strategies: distributing samples evenly across all bins or sampling exclusively from the middle bin. However, in practice, people tend to distribute samples evenly across bins to mitigate real-world uncertainties, such as potential discrepancies between the true underlying beta distribution and the observed histogram.

Analysis Results of H1

Table 6.4.1*Uniform condition*

Test	Number of Successes	Number of Trials	p-value	95% Confidence Interval	Probability of Success
Bin 1	1453	8490	< .001	0.163 to 0.179	0.171
Bin 5	1807	8490	.003	0.204 to 0.222	0.213
Bin 3	1988	8490	< .001	0.225 to 0.243	0.234

Table 6.4.2*Left-Skewed Condition*

Test	Number of Successes	Number of Trials	p- value	95% Confidence Interval	Probability of Success
Bin 1	899	8440	< .001	0.100 to 0.113	0.107

Table 6.4.3*Right-Skewed Condition*

Test	Number of Successes	Number of Trials	p- value	95% Confidence Interval	Probability of Success
Bin 5	1034	8530	< .001	0.114 to 0.128	0.121

Analysis Results of H2**Table 6.4.4***The effects on sampling error*

Condition	Samples	Mean Sampling Error	SD	SE	95% CI Lower	95% CI Upper
Left-skewed	10	-0.137	0.0084	0.0005	-0.138	-0.136
Left-skewed	20	-0.108	0.0119	0.0007	-0.109	-0.107
Uniform	10	0.020	0.0084	0.0005	0.019	0.021
Uniform	20	0.026	0.0123	0.0007	0.024	0.027

Condition	Samples	Mean Sampling Error	SD	SE	95% CI Lower	95% CI Upper
Right-skewed	10	0.149	0.0085	0.0005	0.148	0.150
Right-skewed	20	0.117	0.0126	0.0008	0.116	0.119

Table 6.4.5

Bayesian Mixed Model Results for Sampling Error

Effect	Estimate	SE	95% CI Lower	95% CI Upper
Intercept	-0.14	0.01	-0.15	-0.12
Uniform	0.16	0.01	0.14	0.17
Right-Skewed	0.29	0.01	0.27	0.30
Samples = 20	0.03	0.01	0.01	0.05
Uniform \times Samples = 20	-0.02	0.01	-0.05	0.00

Effect	Estimate	SE	95% CI Lower	95% CI Upper
Right-Skewed \times Samples = 20	-0.06	0.01	-0.09	-0.04
Residual <i>SD</i> (sigma)	0.10	0.00	0.10	0.10

Table 6.4.6

Descriptive Statistics of Estimation Accuracy by Condition and Sample Size

Condition	Samples	Mean Estimation Accuracy	SD	SE	95% CI Lower	95% CI Upper
Left-Skewed	10	-0.067	0.107	0.0063	-0.079	-0.054
Left-Skewed	20	-0.036	0.094	0.0056	-0.047	-0.025
Uniform	10	0.049	0.105	0.0062	0.037	0.061
Uniform	20	0.055	0.095	0.0057	0.044	0.066
Right-Skewed	10	0.126	0.105	0.0062	0.114	0.138
Right-Skewed	20	0.083	0.098	0.0058	0.072	0.095

Table 6.4.7*Bayesian Model Results for Estimation Accuracy*

Fixed Effect	Estimate	SE	95% CI Lower	95% CI Upper
Intercept	-0.07	0.02	-0.10	-0.04
Uniform	0.12	0.01	0.09	0.14
Right-Skewed	0.19	0.01	0.17	0.22
Samples = 20	0.03	0.02	-0.01	0.08
Uniform \times Samples = 20	-0.03	0.02	-0.06	0.00
Right-Skewed \times Samples = 20	-0.08	0.02	-0.11	-0.04

Eidesstattliche Versicherung

gemäß § 13 Abs. 2 Ziff. 3 der Promotionsordnung des Karlsruher

Instituts für Technologie für die KIT-Fakultät für Wirtschaftswissenschaften

1. Bei der eingereichten Dissertation zu dem Thema ***The Influence of Search Behavior on Perception and Decision-Making*** handelt es sich um meine eigenständig erbrachte Leistung.
2. Ich habe nur die angegebenen Quellen und Hilfsmittel benutzt und mich keiner unzulässigen Hilfe Dritter bedient. Insbesondere habe ich wörtlich oder sinngemäß aus anderen Werken übernommene Inhalte als solche kenntlich gemacht.
3. Die Arbeit oder Teile davon habe ich *bislang nicht* an einer Hochschule des In- oder Auslands als Bestandteil einer Prüfungs- oder Qualifikationsleistung vorgelegt.
4. Die Richtigkeit der vorstehenden Erklärungen bestätige ich.
5. Die Bedeutung der eidesstattlichen Versicherung und die strafrechtlichen Folgen einer unrichtigen oder unvollständigen eidesstattlichen Versicherung sind mir bekannt. Ich versichere an Eides statt, dass ich nach bestem Wissen die reine Wahrheit erklärt und nichts verschwiegen habe.

Hanoi, den 21.07.2025

Thai Quoc Cao