

Motivational States Shape the Process of Information Seeking

by

Abigail Hsiung

Department of Psychology and Neuroscience
Duke University

Date: _____

Approved:

R. Alison Adcock, Co-Supervisor

Scott A. Huettel, Co-Supervisor

Tobias Egner

John M. Pearson

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Psychology and Neuroscience in the Graduate School
of Duke University

2023

ABSTRACT

Motivational States Shape the Process of Information Seeking

by

Abigail Hsiung

Department of Psychology and Neuroscience
Duke University

Date: _____

Approved:

R. Alison Adcock, Co-Supervisor

Scott A. Huettel, Co-Supervisor

Tobias Egner

John M. Pearson

An abstract of a dissertation submitted in partial
fulfillment of the requirements for the degree
of Doctor of Philosophy in the Department of
Psychology and Neuroscience in the Graduate School of
Duke University

2023

Copyright by
Abigail Hsiung
2023

Abstract

Seeking information is fundamental not only for making adaptive choices but also for extending our knowledge both individually and as a collective society. Understanding how people organize and carry out seeking information, remains an open and critical challenge for the field of psychology and neuroscience. In this dissertation, I approach this challenge by examining the organization of information seeking through the lens of the motivational states of an individual. I leverage behavioral, computational modeling, and neuroimaging methods to provide a multi-faceted understanding of the relationship amongst motivation, information seeking, and memory formation. In the first study, I used computational models to examine latent sampling strategies as people gathered information to guide subsequent action. These results demonstrated that unlike the optimal strategy, people perform simple accumulations instead of computing expected values, and that unlike a simple heuristic strategy, people dynamically integrate multiple sources of information in lieu of using only one source, highlighting a computationally efficient strategy that balances competing motivations for accuracy and cost minimization. In the second study, I explore how curiosity evaluates information as it unfolds over time, reporting that curiosity can promote the patience to prolong uncertainty and such patience culminates in experiencing discovery. In the third study, I extend these findings to capture the rewarding process of discovery and its impact on memory using functional magnetic resonance imaging. In the fourth study, I investigate how distinct motivational forces can drive information seeking and learning through distinct mechanisms. Collectively, the work presented in this dissertation reveals that our motivational states play a significant role in shaping our information seeking behaviors,

working to balance competing goal states, and pushing us to confront and even enjoy the experience of uncertainty.

Contents

Abstract.....	iv
List of Tables.....	x
List of Figures.....	xi
Acknowledgements.....	xiii
1. Introduction	1
1.1 Delineating Instrumental and Non-Instrumental Information Seeking.....	3
1.2 Reward motivation shapes the value of instrumental information seeking	4
1.2.1 Competing motivations shape the structure of instrumental information seeking.....	5
1.3 Curiosity shapes the value of non-instrumental information.....	8
1.3.1 Curiosity towards the reduction of uncertainty.....	11
1.3.2 Curiosity towards the collection and integration of new knowledge.....	13
1.4 Motivational states influence the structure of non-instrumental information seeking.....	17
1.4.1 Motivational states balance suspense and surprise across information-seeking.....	17
1.4.2 Curiosity shapes sequential search through the development of knowledge.....	20
1.5 Summary and the current dissertation	21
2. Between heuristics and optimality: Flexible integration of cost and evidence during information sampling.....	24
2.1 Introduction	24
2.2 Methods and Materials	27
2.2.1 Participants	27
2.2.2 Procedure.....	28

2.2.3 Computational Modeling	31
2.2.3.1 Optimal Behavior	32
2.2.3.2 Expected Value Computation Strategy.	34
2.2.3.3 Evidence Difference Heuristic Strategy.....	37
2.2.3.4 Sample Number Heuristic Strategy.....	37
2.2.3.5 Evidence-by-Costs Tradeoff (ED-CA) Strategy.	38
2.2.3.6 Model Comparison.....	39
2.3 Results	40
2.3.1 Greater sampling is associated with higher task accuracy, but at the expense of greater cost accumulation.....	41
2.3.2 Models of information sampling strategies.....	43
2.3.3 Moment-by-moment sampling decisions were best predicted by an Evidence- by-Costs (ED-CA) strategy	47
2.3.4 Stopping criteria were non-stationary within a trial but stationary across trials	48
2.3.5 Increasing reward stakes does not encourage use of more optimal strategy	50
2.4 Discussion.....	52
3. Spoiler Alert! Curiosity can prioritize the information gathering process over the outcome	59
3.1 Introduction	59
3.2 Methods	64
3.2.1 Participants.	64
3.2.2 Stimuli and Apparatus.....	64
3.2.3 Task.....	65
3.2.4 Data Analysis.....	69
3.3 Results	70

3.3.1 Study 1	70
3.3.1.1 Curiosity and Information Seeking	71
3.3.1.2 Uncertainty and Information Seeking	73
3.3.1.3 Curiosity and Affect.....	75
3.3.1.4 Study 1 Conclusions.....	76
3.3.2. Study 2.....	77
3.3.2.1 Choice Behavior Results	78
3.3.2.2 Affect Results.....	81
3.3.2.3 Free vs. Predetermined Timing	82
3.3.2.4 Study 2 Conclusions.....	85
3.4 General Discussion.....	86
4. The process of discovery engages distinct neural correlates across unfolding information	93
4.1 Introduction	93
4.2 Methods and Materials	97
4.2.1 Participants	97
4.2.2 Experimental Procedure	98
4.2.3 Statistical analysis of memory data	100
4.2.4 fMRI data acquisition and preprocessing	101
4.2.4.1 Anatomical data preprocessing	102
4.2.4.2 Functional data preprocessing	102
4.2.5 Univariate fMRI analyses.....	103
4.3 Results	105
4.3.1 Behavioral Response Distribution	105

4.3.2 Behavioral Correlates of Memory	107
4.3.3 Univariate Whole Brain Activation	108
4.3.3.1 Explore Phase	109
4.3.3.2 Hypothesis Testing Phase.....	111
4.3.3.3 Comparing Explore and Hypothesis Testing Phases	114
4.3.3.4 Subsequent Memory Effects	117
4.4 Discussion.....	119
5. Complementary roles of curiosity and prediction error in information seeking and learning	126
5.1 Introduction	126
5.2 Methods and Materials	128
5.2.1 Participants	128
5.2.2 Task Design and Procedure	129
5.2.3 Data Analysis	131
5.3 Results	133
5.3.1 Seeking Additional Information.....	133
5.3.2 Memory Updating.....	135
5.4 Discussion.....	138
6. Conclusion	144
6.1 Study Summaries	144
6.2 Future Directions.....	148
6.3 Conclusion	151
References.....	153
Biography	174

List of Tables

Table 1: Effects of video type during the Explore Phase on whole-brain activity.....	111
Table 2: Effects of video type during the Hypothesis Testing Phase on whole-brain activity.	113
Table 3: Effect of trial phase on whole-brain activity for Real videos only.	116
Table 4: Subsequent memory effect on whole-brain activity for Real videos only.....	118

List of Figures

Figure 1: Task Design. Participants completed a modified version of the Information Sampling Task.	30
Figure 2: Sampling tendencies between participants relate to overall task accuracy and total earnings.....	42
Figure 3: Tested information sampling strategies with model predictions and example representations.	46
Figure 4: Evidence-by-Costs (ED-CA) strategy outperforms both simpler heuristics and expected value models in predicting decisions.	48
Figure 5: Evidence difference and number of samples trade off.	49
Figure 6: High reward stakes increase sampling but do not encourage more optimal behavior.	51
Figure 7: Task design and Study 1 and 2 manipulation.....	66
Figure 8: Results from Study 1. Curiosity predicts continuation across Forgo- and Jump-to Resolution conditions.....	73
Figure 9: Results from Study 2. Curiosity predicted continuation even when participants were given agency over the timing of their choices.	79
Figure 10: Comparing information gathering behavior between Study 1 and Study 2....	84
Figure 11: Two session task design.	105
Figure 12: Memory results for session 2. Rates of remembering a video during re-watch (Session 2).....	108
Figure 13: Whole-brain during Exploration Phase (for a full list of regions, see Table 1).	110
Figure 14: Reaching realization preferentially activated reward-related regions and object recognition regions for Real compared to Scribble videos.....	113
Figure 15: Whole-brain activation comparing Exploration versus Hypothesis Testing Phase (for a full list of regions, see Table 3).....	116
Figure 16: Real videos that were later remembered show increased activation in the hippocampus, PHC, and midbrain at encoding.....	118

Figure 17: Task design for testing the influence of curiosity and prediction error on information seeking and memory updating.	131
Figure 18: Curiosity and prediction error offer complementary motivations to seek additional information.....	135
Figure 19: Curiosity and Prediction error independently enhance memory updating. ..	136
Figure 20: Seeking additional information boosts memory updating, particularly for information tested after a delay.	137

Acknowledgements

The work within this dissertation is a reflection of the deeply collaborative nature of science. I would not have made it through graduate school without the support, love, and guidance from a collection of truly incredible people. Although the words here will never be able to express my full gratitude and appreciation, they can call out just how many people it took to get here today.

First and foremost, I'd like to thank Drs. Alison Adcock and Scott Huettel for their mentorship and guidance over the last six years. I have learned so much from you both, not only in how to conduct science but how to tell compelling stories, how to think mechanistically, and how to remain curious amidst the pressures of graduate school. Beyond these lessons, thank you for giving me the freedom to be creative with building new paradigms and to take risks with the questions I wanted to explore. I could not have asked for a better and more complementary set of mentors. I also want to thank John Pearson for his incredible patience and commitment to teaching me math, I never would have believed I could build computational models, but here I am! Finally, I'd like to thank Tobias Egner for lending me his time and expertise during my critical milestones throughout graduate school.

I have been so incredibly fortunate to meet the most amazing people here at Duke. These people have not only been exceptional role models and sources of guidance but have also become some of my closest friends. To my cohort members, Eva Gjorgieva and Leonard Faul, and my lab mate, Matthew Bachman, I don't even know where to begin. Thank you for making graduate school a time of joy and fun. From 8-hour baking adventures to touring the world through wine to learning how to play one

of the most complex and silly games (The House Haunting with Betrayal, on a Hill?), it is because of y'all that I am the scientist and the person I am today. I have been truly gifted with your friendships. To Brenda Yang, who from day one has been a source of inspiration both as a scientist and as a human. Your commitment to pedagogy, accessibility, and design are unparalleled, and have offered an insane amount of guidance throughout my academic (and industry!) journey.

To the members of both the Adcock and Huettel labs, past and present, including Allie Sinclair, Candice Wang, Ben Muzekari, Kelly Eom, John Thorp, Sonakchhi Shrestha, Katie Dickerson, Laura Yost, Shabnam Hakimi, Dianna Amasino, Rosa Li, Samyu Dore, Libby Jenke, Ed McLaurin, Nikki Sullivan, Khoi Vo, and Alex Breslav, y'all have always made lab feel like a safe and amazing place to practice science. A very special thanks to Mai-Anh Vu, you showed me how to confront challenge (and coding errors) with curiosity and to think deeply about diversity and equity in science (also I truly appreciated climbing with you!). To Rachael Wright, you essentially taught me to preprocess fMRI data and I am forever grateful for our late-night sessions that got me through writing this dissertation, and to Jia-Hou Poh, you are an exceptional teacher, one of the smartest humans I know, and just a truly stellar human.

The Center for Cognitive Neuroscience has been a wonderful place for learning science, and I am forever grateful to all of the members that have shaped my thinking along the way. I'd also like to give thanks to staff members Colleen Bauer and Tanya Schreiber for making all of this work possible and easy.

Outside of my mentors here at Duke, I want to recognize that I would not have made it to graduate school without the support and guidance from all of the members of

the Section for the Neurobiology of Fear and Anxiety. To Christian and Monique, thank you for all of your guidance and for providing the warmest, most welcoming space in 15K. To my fellow IRTAs, Bari Fuchs, Joseph Leshin, Jeffrey Liu, and Gabby Alvarez, thank you for making the IRTA experience the absolute best. I am truly fortunate to have had such a collaborative, collective, and fun post-bac experience. Thank you to Adam Gorka and Sam Torrisi for everything from teaching me statistics to explaining the wonders of 7T. And, lastly, a special thank you to Nick Balderston, who beyond teaching me so, so much about how to be a scientist, taught me that karaoke was a great way to reduce the fear of public speaking.

This journey would not have been possible if not for the outpouring support I've received from my wonderful friends. To Amanda Woodall, I cannot thank you enough for being the most wonderful friend, your strength, kindness, and openness to get down in any wild cooking adventure has truly brought me so much joy and solace. To Megan Gallagher, you are an unbelievable source of encouragement, I cannot tell you how many times I've looked to you for inspiration for both compassion and tenacity. To Peter Lostritto, thank you for creating a space for silliness, seriousness, and everything in between. To the people who have helped Durham feel like home, including Reid Shalvoy, Emma Powers, Ryan Robertson, and Allison Pfeffer, thank you for being the absolute greatest.

Finally, I'd also like to thank my family for supporting me throughout this journey. To my parents, Laura and Shang Hsiung, thank you for your encouragement and love as I traversed fashion school to graduate school. You both have been incredible examples of compassion, hard work, sincerity, and perseverance. To my brother Matthew and his

partner, Michelle Seals, thank you for all the wine and game nights, the volleyball and the food. To my in-laws, Steve and Donna Cravath, and Cristen and Ben Clemm, thank you for all of the support and for welcoming me so generously into your family. To my husband, Craig, thank you for being my best and brightest cheerleader and champion, my greatest source of support through all of the challenges, and for being my bestest friend.

1. Introduction

Seeking information is fundamental to navigating and flourishing in an uncertain and ever-changing world. The act of information seeking aids a variety of goals, from increasing the accuracy of our choices to building out our knowledge of the world (D. E. Berlyne, 1966; Garner, 1962; Gold & Shadlen, 2007; Golman & Loewenstein, 2018; Gottlieb & Oudeyer, 2018; Sharot & Sunstein, 2020). Yet the amount of information currently available to us is almost infinite. Such excess conflicts with the limitations that we face as biological agents; we do not have infinite time or resources to explore and interact with all available information. Instead, we must be selective in how we seek out and collect information.

Often, seeking information is not simply a singular decision but a series of choices that unfold across time. Sequential information seeking allows for the information that we collect to build upon itself, helping to form an impression or judgment, or moving us from a simple to a more complex understanding (Gottlieb et al., 2014; Lydon-Staley et al., 2021). However, open questions remain regarding how we decide when we have gathered enough information.

Understanding how people organize and titrate their search behaviors not only has consequences for subsequent actions but also shapes what we end up learning (Gruber & Ranganath, 2019). Our memories are dependent upon the information we expose ourselves to, and such memories can also be impacted by how information is experienced across time. Despite the importance of the interplay between information

seeking and memory formation, few studies have explored how our learning is molded by how we search for information.

In this dissertation, I approach these challenges by examining the motivational state of an individual during ongoing sequential search. Motivation, here, refers to both the higher-level incentive structure behind information gathering (e.g., to build an accurate model of the world) as well as the modulation of cognitive resources (e.g., attention, effort, dynamic neural states) engaged during search (Braver et al., 2014). The overarching goal of my dissertation research is to shed light on how different motivations can prioritize unique patterns of sequential information seeking, producing tailored outcomes both in terms of information gathered and memories formed (Lydon-Staley et al., 2021; Sharot & Sunstein, 2020).

Within the studies below, I advance upon these ideas by examining several ways in which information seeking is organized under different motivational states. To contextualize the work within my dissertation, I will first discuss the background literature pertaining to information seeking as it is operationalized and studied across a range of motivations. Specifically, I will provide an overview of how motivational states help shape the organization of sequential information seeking both when information directly aids in a subsequent choice and when information appears to be sought after as a reward in and of itself. Then, in the first study, I investigate how information is collected and integrated towards a judgment under two reward contexts (**Study 1**). In the next study, I examine how the motivational state of curiosity guides preferences for how information unfolds across time and the value of discovery (**Study 2**). Next, I explore the neural correlates underlying the process of discovery to build upon our understanding of why

curiosity promotes patience and how patience enhances memory (**Study 3**). Finally, I delineate the different motivations of curiosity and surprise on driving information seeking and memory formation (**Study 4**). Throughout these studies, I leverage flexible and complex paradigms that give people agency over their choices, allowing for the examination of organic search behaviors. The goal of this work, collectively, is to use behavioral, computational, and neural measures to better characterize the role of motivation in driving sequential information seeking across dynamic and evolving informational states.

1.1 Delineating Instrumental and Non-Instrumental Information Seeking

We engage in search behaviors in a variety of contexts within our everyday lives. One common distinction used to categorize the types of search processes is whether the goal of seeking information is to inform a subsequent choice or not. These two types of information seeking strategies are referred to as instrumental (when information aids subsequent choice) and non-instrumental (when information is not in service of subsequent choice) (Golman & Loewenstein, 2018; Gottlieb & Oudeyer, 2018; Kobayashi & Hsu, 2019; van Lieshout et al., 2020). For example, seeking out information about one's local weather provides guidance toward whether one should pack an umbrella. Conversely, seeking information about the weather in a distant region provides little guidance towards preparing for the day's weather. Instead, such information can be described as having value in and of itself, independent from its ability to aid in adaptive choice (Kidd & Hayden, 2015).

Differentiating instrumental and non-instrumental information seeking also provides a boundary with which to understand motivational states. For the majority of instrumental information seeking processes, the motivational state can be posed as reward motivation, where actions are taken to obtain a preferred outcome (Gottlieb & Oudeyer, 2018). In contrast, non-instrumental information seeking has mostly been linked with the motivational state of curiosity (Gottlieb & Oudeyer, 2018; Kidd & Hayden, 2015; van Lieshout et al., 2020). Throughout this dissertation, I will use this distinction to call attention to key differences between how these different motivations assign value to information and how this valuation influences the process of sequential information seeking. This is not to say that information seeking cannot be comprised of both instrumental and non-instrumental motives (Kobayashi & Hsu, 2019; Sharot & Sunstein, 2020), but that through this separation we can begin to understand the unique contributions each motivation can have in driving and organizing sequential information seeking.

1.2 Reward motivation shapes the value of instrumental information seeking

When we are motivated to obtain rewards, we assess the value of information insofar as it pushes us towards our desired outcome (Edwards & Slovic, 1965; Howard, 1966). Most often, information imparts value by reducing the uncertainty around which actions will most likely lead to reward. Prior work has demonstrated the link between the experience of uncertainty in an economic choice and the choice to seek relevant information (Blanchard & Gershman, 2018; Desender et al., 2018; Kaanders, Nili, et al., 2021; Kobayashi & Hsu, 2019; Kobayashi et al., 2021). Importantly, the value of

information scales with the amount of uncertainty expected to be relieved, such that information seeking choices are more likely to occur when they alleviate a greater amount of uncertainty (Kaanders, Nili, et al., 2021; Kobayashi & Hsu, 2019; Kobayashi et al., 2021) or during more difficult choices (Coenen & Gureckis, 2016; Furl & Averbeck, 2011; Kiani et al., 2014). Notably, in cases where assessing the exact amount of uncertainty is intractable for decision-makers, information seeking is also predicted by a proxy measurement of uncertainty: *confidence* (Desender et al., 2018, 2019; Hausmann-Thürig & Läge, 2008; Pouget et al., 2016; Schulz et al., 2023). Collectively, when motivated to obtain a reward, information is valued through the uncertainty-reducing advantage it provides a decision-maker in obtaining preferred outcomes.

1.2.1 Competing motivations shape the structure of instrumental information seeking

Beyond the initiation of information seeking, we often collect multiple pieces of information before we commit to a choice. In these types of decision processes, samples of evidence are collected sequentially until the decision-maker decides that enough evidence has been gathered to inform a final judgment. Throughout the sampling process, the decision-maker receives no feedback from their environment, which differs from foraging and reinforcement learning processes where information in the form of reward values is provided after each sequential choice. Instead, the decision-maker must determine the quality and quantity of evidence they wish to uncover before a subsequent action is made.

Therefore, we need not only compute the value of a single piece of information but also decide how to organize our information-gathering process. This process can

involve how we decide when we have gathered enough information, how we handle the costs of information gathering, and how we select and implement different strategies of information sampling. Importantly, our motivation to obtain a desired outcome or reward may now stand in competition with motivations to reduce the amount of effort we expend (Inzlicht et al., 2018; Otto & Daw, 2019) or to minimize the computation load (Bossaerts & Murawski, 2017; Shenhav et al., 2017) of deciding whether more information is needed.

When information seeking extends over multiple samples from the environment, a decision-maker must decide how much information is enough to terminate the information gathering process. One mechanism relies on tracking the dynamic change in the value an additional piece of information will impart (Kobayashi et al., 2021). Here, the value of the next piece of information depends on the evidence already collected. If the value of an additional piece is unlikely to further reduce the uncertainty of a choice, then this provides a signal for the end of information seeking (Hausmann-Thürig & Läge, 2008; Kiani et al., 2014). But value is not the only factor in this computation.

The decision to stop information gathering is also influenced by the amount of elapsed time since the start of information gathering regardless of the evidence gathered (Drugowitsch et al., 2012; Kiani et al., 2014; Murphy et al., 2016). The passage of time can evoke non-linear tradeoffs between gathering evidence and the pressure to commit to a choice. (Drugowitsch et al., 2012; Murphy et al., 2016). Computational models suggest that such non-linear decreases could be a result of dynamic changes to the criterion for stopping during evidence accumulation (Murphy et al., 2016). Hauser et al. (2018) found that blocking noradrenaline (NE) lead to decreased information-gathering

behavior by promoting the expression of a decision-related urgency signal (Cisek et al., 2009), which caused the value of continuing to seek information to drop starkly once a threshold was crossed (Drugowitsch et al., 2012; Hauser et al., 2017, 2018; Murphy et al., 2016). The authors speculated that blocking NE could result in a dampened sensitivity to uncertainty (Aston-Jones & Cohen, 2005; Yu & Dayan, 2005) potentially, impacting how the value of information was initially computed.

Other costs beyond the passage of time have also been shown to contribute to the balance between the motivation to obtain reward and the motivation to minimize the resources used to gather and integrate instrumental information. Research investigating the costs of information sampling has examined both experimentally imposed costs (Cisek et al., 2009; Esslinger et al., 2013; Furl & Averbeck, 2011; Hauser et al., 2018; Juni et al., 2016; Malhotra et al., 2017) as well as factoring in implicit, or internally derived costs (Hauser et al., 2018; Kobayashi et al., 2021; Petit et al., 2021). Overall, such costs reduce the value of information, resulting in fewer choices to seek information and earlier terminations. Indeed, decision-makers must strike a balance between the time and resource costs of gathering information with the expected payout for a successful choice.

Assessing how decision-makers strike this balance highlights the importance of understanding what underlying computations are being used to calculate the value of information, the costs of search, and how these quantities are integrated to inform the next action. Central to this computation is an accurate estimate of the uncertainty one faces in a choice (Bach & Dolan, 2012). Under ideal conditions, normative approaches can accurately estimate uncertainty using methods that rely on complex and multi-step

procedures (FitzGerald et al., 2015; Moutoussis et al., 2011). Although possible to deploy, humans are also known to shortcut such intensive computations in favor of simple heuristics (Gigerenzer & Gaissmaier, 2011). Thus, one outstanding question reflects how a decision-maker selects which of many different information-gathering strategies to execute to balance the competing motivations to obtain reward and minimize expenditure. Prior work has documented the existence of tradeoffs between computational complexity compared to simpler heuristics (Bossaerts & Murawski, 2017; Gerd Gigerenzer & Gaissmaier, 2011) but less work has examined the space between the most complex and the simplest of strategies. Furthermore, little work has been done to explore how strategies on the spectrum of complex to simple are selected and implemented during instrumental information gathering. Recent work has begun to demonstrate the use of simplified mechanisms to reconcile these competing motivations (Baumann et al., 2020; Korn & Bach, 2018; Sang et al., 2020), but these studies have not examined how a decision-maker decides without external feedback how much information is required to guide subsequent choice. Therefore, in **Study 1** I address this question by developing strategies along a continuum of simple to complex and evaluate the types of strategies people use to guide their sequential search behavior. Moreover, I expressly test whether the deployment of strategies changes as a function of increasing the available rewards.

1.3 Curiosity shapes the value of non-instrumental information

Information helps us understand the world by guiding our choices toward outcomes that provide the greatest amount of external reward. However, there are often

cases when we seek out or gather information that is not in service of instrumental action. Indeed, humans, as well as other animals, will spend substantial amounts of time and energy seeking information that provides no immediate benefit to obtaining rewards (D. E. Berlyne, 1966; Kidd & Hayden, 2015). Although seemingly intuitive, these behaviors stand in contrast with the notion that information holds value only so much as it helps with a future action. Why would the limited resources of time and energy be spent on uncovering information deemed nonessential to maximizing our current reward outcomes? In this next section, I will review the current understandings of non-instrumental information seeking and the motivational states that guide how and when information is sought after.

To understand the qualities of information that contribute to its value, we must examine the motivational states that lead to seeking non-instrumental information. Broadly, theories of non-instrumental information seeking posit that curiosity can act as a motivational force to guide the search and attainment of information (D. E. Berlyne, 1962, 1966; Gottlieb & Oudeyer, 2018; Kidd & Hayden, 2015; van Lieshout et al., 2020). For the purpose of this dissertation, curiosity is defined as a motivational state that encourages information seeking without regard to the instrumental value of the information sought after. While curiosity represents one of many motivational forces that can encourage non-instrumental information seeking unless explicitly specified, curiosity and information seeking will be used interchangeably. Importantly, however, there is no single agreed-upon definition of curiosity within the literature and the lack of a singular definition signals that curiosity likely takes many forms or flavors that aim to maximize different goals. In his seminal work regarding curiosity, Daniel Berlyne divided the state

of curiosity into several forms, highlighting that the breadth of information of interest could produce different curious states (specific vs. diversive curiosities) as well as the qualities of the information itself (perceptual vs. epistemic curiosities) (D. E. Berlyne, 1966). Throughout his work, Berlyne underpinned the notion that curiosity as a motivational state was fluid in how it arose, how it motivated information search, and how its relief was determined. This provided the groundwork for thinking about curiosity and the information search it directed as a multifaceted state comprising of several components.

Since Berlyne, numerous theories have emerged to delineate the components of curiosity and how each component services different goals during information seeking. Collectively, two main functions of curiosity arise: to reduce environmental uncertainty (Gottlieb & Oudeyer, 2018; Kobayashi et al., 2019; Noordewier & van Dijk, 2020; van Lieshout et al., 2020), and to acquire new knowledge about the world around us (Chater & Loewenstein, 2016; E. G. Liquin & Lombrozo, 2020a; Noordewier & van Dijk, 2020; Zurn & Bassett, 2018). To understand how motivational states of curiosity shape the value of information and thus shape the selection and refinement of sought-after information, we must first unpack how the different functions of curiosity impart value on information and influence choice behavior. Towards this end, I will expand beyond the categorization put forth by Berlyne by contextualizing the different variations of curiosity as they relate to the achievement of different goals.

1.3.1 Curiosity towards the reduction of uncertainty.

A core function of curiosity is the motivation to reduce uncertainty. In many cases, however, uncertainty reduction is the only goal of seeking information, such that the obtained information is wholly non-instrumental and cannot be used outside of the context in which it was experienced. A common paradigm used to assess this is through “observing tasks,” where observers were presented with gambles (e.g., a 20% chance of winning \$1.00 and an 80% chance of winning \$0.00) and could choose whether to receive information about the outcome of these gambles immediately or wait until the outcome was revealed after a delay (Blanchard et al., 2015; van Lieshout et al., 2021, 2018). Critically, information in these paradigms is non-instrumental; observers could not use the information to increase their chances of winning in the future nor did the information hold any inherent value outside the context of an individual outcome (i.e., it cannot be used to understand the task more broadly). Another paradigm type examines the induction and relief of perceptual ambiguity, where people are first shown obscured or blurred images that are then made clear after a delay (Cohanpour et al., 2022; Jepma et al., 2012; Van de Cruys et al., 2021). Across both of these contexts, there are notably behavioral, affective, and neural characteristics that describe this type of curiosity.

Behaviorally, we see that uncertainty-reduction-based curiosity demonstrates a preference towards reducing the uncertainty that provides the largest information gain. In observing experimental tasks, this is reflected in a preference for seeking information about the outcome of a gamble when either outcome is equally likely to happen (i.e., 50/50 chance of winning or losing) (Charpentier et al., 2018; Goh et al., 2021; Jiwa et al., 2021; van Lieshout et al., 2021, 2018; Wang & Hayden, 2019). Across perceptual

curiosity tasks where participants first see ambiguous and then clarified perceptual information, this pattern is reflected in the highest ratings of curiosity coinciding with the highest amounts of uncertainty, either measured as the inverse of confidence ratings (Cohanpour et al., 2022) or as the highest amount of entropy (Van de Cruys et al., 2021). Therefore, when the goal of curiosity is geared towards reducing the state of uncertainty for its own sake, information seeking is steered towards finding the options that will create the largest reduction.

Curiosity that is intended to purely reduce the experience of uncertainty is also coupled with aversive and unpleasant affective states (Bennett et al., 2016; Hsee & Ruan, 2016; van Lieshout et al., 2021). Such affect is corroborated by the neural correlates that track the induction and relief of uncertainty-reduction based curiosity. During the induction of curiosity, when people are first confronted with uncertainty, we see heightened activation in the parietal cortex (van Lieshout et al., 2018), the anterior cingulate cortex (ACC), and the anterior insula (AI) (Jepma et al., 2012), regions typically associated with uncertainty (Huettel et al., 2005; Y. Li et al., 2022) and saliency (Bartra et al., 2013; Seeley et al., 2007). These findings suggest that in these contexts curiosity acts as a detector, heightening arousal and attention towards highly uncertain events. In contrast, during the anticipation and relief of uncertainty, we find activation in areas canonical to rewards processing, including the caudate, putamen, nucleus accumbens (NAcc), and the orbitofrontal cortex (OFC) (Blanchard et al., 2015; Jepma et al., 2012; van Lieshout et al., 2018). In addition, van Lieshout et al. (2018) found increased activation in the right anterior insula corresponding to the relief of curiosity that tracked with the magnitude of positive information prediction errors (Preuschoff et al., 2008),

suggesting that when curiosity is solely concerned with the reduction of uncertainty, the size of that reduction matters. Collectively, when curiosity is focused solely on uncertainty reduction, the need for relief can be viewed as an aversive state pushing information seeking forward to quell the experience of uncertainty.

One potential difference between reducing uncertainty for observing vs. perceptual tasks is the moment of experiencing discovery. Within observing tasks there is uncertainty around the outcome of a gamble but the available options are all known, as the gamble is either won or lost. In perceptual tasks the identity of the obscured image can take various identities, leading to greater uncertainty about the outcome. Thus, the moment of relief when information about the outcome is revealed can happen alongside a moment of discovery (Aha! It *is* a frog). Indeed, past work has shown that the opportunity for discovery in perceptual tasks is coupled with increased activation in the hippocampus (Jepma et al., 2012), which is not shown in parallel gambling-based tasks. However, it remains unexplored whether the opportunity to experience discovery changes when information unfolds across time as opposed to an instantaneous reveal. In **Studies 2** and **3**, I investigate this question using dynamically unfolding drawings to capture the moment of discovery while also assessing the affective and neural components associated with it.

1.3.2 Curiosity towards the collection and integration of new knowledge.

Most ingrained in the colloquial nature of curiosity is its role in building knowledge. Curiosity guides us towards new information, enhancing our understanding of the world around us and our ability to adapt within it. Curiosity towards knowledge

building maps most closely to Berlyne's original formulation of "epistemic curiosity" (D. Berlyne, 08/1954; D. E. Berlyne, 1962). Under the quest for knowledge, curiosity places value on information insofar as it can enhance the learning experience of a person, both in the short term and towards building long-term memories.

Paradigms that investigate the knowledge acquisition flavor of curiosity typically use trivia question-and-answer paradigms to understand what epistemic information drives people's information-seeking behaviors. In these tasks, learners might rate each trivia question on a variety of scales, including how curious they are to know the answer to that question and how confident they are to know the outcome. To understand the impact of curiosity on choice behavior, these tasks will then have learners choose which answers to view and which to skip over. Across a variety of studies, it has been consistently demonstrated that questions that elicit the most curiosity are the ones selected to uncover their answers (Baranes et al., 2015; Dubey et al., 2022; Fastrich et al., 2018; Kang et al., 2009; Lang et al., 2021; Lau et al., 2020; Marvin & Shohamy, 2016; Wade & Kidd, 2019). When we examine how learners assign a rating of curiosity to individual questions, we find that curiosity is related to measures of confidence. Unlike flavors of curiosity aimed towards uncertainty reduction, curiosity under knowledge acquisition guides choices that evoke a *moderate* amount of uncertainty (Baranes et al., 2015; Dubey & Griffiths, 2020a; Kang et al., 2009; Wade & Kidd, 2020). Illustrated by an inverted U-shape, curiosity here is piqued when there is some background knowledge about the trivia question at hand but not too much such that the answer is likely known. Why might curiosity favor intermediate levels of confidence to guide information seeking towards knowledge but maximize uncertainty reduction in other cases? One potential

reason is that under the context of enhancing knowledge, a learner must also think about how such information fits into the broader scope of their knowledge base. Information that is too far away is likely to be forgotten with nothing around it to offer contextualization while information that is too close or known is not offering anything new to the knowledge base (Zhou et al., 2020; Zurn & Bassett, 2018).

In addition to using confidence to evaluate information and guide search, curiosity has also been shown to specifically prefer an understanding of how things work together. For example, seeking information that illuminates a “why” or “how” question appears particularly relevant to how curiosity directs search (Gopnik, 2000; Lau et al., 2020; E. G. Liquin & Lombrozo, 2020a; E. Liquin & Lombrozo, 2017). In explanation curiosity, the role of understanding how sought-after information will influence future learning and utility of such information plays a larger role than in fact-based curiosity (E. G. Liquin & Lombrozo, 2020a) and provides support for the future use of information informing the valuation judgment (Dubey & Griffiths, 2020).

Beyond obtaining new knowledge, curiosity must also participate in boosting memory formation for information of interest in order to enhance the learning experience of a person. Indeed, numerous findings reveal that curiosity promotes successful memory formation for information (Fastrich et al., 2018; Gruber et al., 2014; Halamish et al., 2019; Kang et al., 2009; Marvin & Shohamy, 2016; Poh et al., 2022). Curiosity also enhances memory for incidental information that is unrelated to the target information (Gruber et al., 2014; Halamish et al., 2019). Moreover, curiosity leverages the same mechanisms for memory enhancement as seen in other reward motivations. Under high states of curiosity, there is increased activation in central reward-processing regions,

such as the ventral tegmental area (VTA), NAcc, and other parts of the striatum (Gruber et al., 2014; Kang et al., 2009; Lau et al., 2020; Poh et al., 2022). Curiosity also increases activation in the hippocampus (Gruber et al., 2014; Jepma et al., 2012; Poh et al., 2022), a key region for memory formation (Adcock et al., 2006; Ergorul & Eichenbaum, 2004). Additionally, curiosity appears to facilitate the communication between the VTA and the hippocampus, increasing the functional connectivity between the regions (Gruber et al., 2014; Wittmann et al., 2007) and potentially supporting the emergence of a convergence state for learning in the hippocampus (Poh et al., 2022). In building knowledge of our world, curiosity encourages us to seek out knowledge just beyond our existing reach and promotes the longevity of that knowledge by enlisting the neural architecture to enhance memory formation.

Despite these advances, it remains unclear how curiosity functions along another closely-related motivational state: surprise. Surprise, often occurring when our expectations are violated, can also impact information seeking and memory formation (Metcalf, 2017; Pine et al., 2018; Vogl, Pekrun, Murayama, & Loderer, 2019). While previous work has documented that curiosity can be informed by the experience of surprise (Vogl, Pekrun, Murayama, & Loderer, 2019; Vogl, Pekrun, Murayama, Loderer, et al., 2019), it remains unexplored if such states reflect unique processes towards encouraging information seeking. In **Study 4**, I examine the contributions of curiosity and surprise to tease apart their interactions in driving both information seeking and learning.

1.4 Motivational states influence the structure of non-instrumental information seeking.

So far, the reviewed research on non-instrumental information seeking has focused on single, isolated choices. Investigation into such choices has revealed important features of how curiosity imbues value to information outside of its instrumental use. Yet, information seeking under motivational states like curiosity can also be persistent and enduring. Broadening our scope to understand how curiosity guides sequences of choices provides a foundation to piece together how we organize our own learning outside the context of instructional learning, focusing instead on how we seek and learn to empower our most personal and deep-rooted interests. However, the complexity we undertake by expanding our scientific pursuit is not without challenge. Given the breadth of information available and the unique profile of interests within each person, understanding how curiosity sequentially organizes choices requires creativity in experimental design and patience to tackle complex and messy behavior. The progress put forth thus far in the literature provides fundamental steppingstones to building further research, but we are still in our infancy. The following section will review the nascent research into how motivational states can shape non-instrumental information seeking paths and trajectories, focusing on the roles of suspense, surprise, and affect.

1.4.1 Motivational states balance suspense and surprise across information-seeking.

When examining sequences of non-instrumental information-seeking choices, a person's tolerance to investigating uncertainty appears to be related to the cumulative amount of surprise they have experienced in recent events (Ligneul et al., 2018).

Specifically, experiencing surprising events decreases the value of approaching future uncertainty. Consequently, our experiences of curiosity live in balance with our motivation to maintain a structured and organized view of the world. We draw back and avoid uncertainty when experiencing too much in conflict with what we think we know, but when we feel satisfied in our knowledge state, we invite opportunities to experience uncertainty and surprise (Ligneul et al., 2018; Schwartenbeck et al., 2019). However, the experience of uncertainty is not always aversive and prior work has begun to demonstrate when momentarily avoiding the resolution of uncertainty might actually be preferred even for highly anticipated information.

Unlike examining singular choices in which the path to resolution is marked by a single transition point when information gathering evolves across time the journey towards resolution can take a variety of shapes even if they all arrive at the same outcome. Central to understanding preference over different paths towards resolution is the experience of suspense. Suspense can be characterized as the feelings of tension and excitement that coincide with the anticipation of an event (Ely et al., 2015; Z.-W. Li et al., 2021). Prior work has shown that people prefer narratives, sports, and video games with high compared to low levels of uncertainty because these events are characterized by high suspense (Ely et al., 2015; Klimmt et al., 2009; Z.-W. Li et al., 2021; Su-lin et al., 1997).

Beyond a passive preference, people will also choose to remain in suspense by specifically avoiding opportunities to receive highly anticipated information too early (Boyd-Graber et al., 2013). Recent work has started to link to motivational state of curiosity with the behavior to delay resolution. When people read trivia questions, high

states of curiosity predicted suspending resolution in order to try and guess the answer (Metcalfe et al., 2021), and forcing a delay between question and answer promotes memory formation but only for questions that elicited curiosity (Mullaney et al., 2014).

While these studies provide a foundation for how curiosity organizes sequential information, several questions remain regarding how curiosity shapes preferences for the reveal of information. Notably, how does curiosity change over the course of information gathering and how sensitive is curiosity to developments in one's informational state? In **Study 2**, I demonstrate that curiosity can prioritize the prolonging uncertainty when there is the chance to experience discovery. Through the creation of dynamic, evolving drawings, I explicitly measure the changes in curiosity across unfolding information and demonstrate how the value of information gathering declines once resolution has been reached even if this occurs before a drawing is complete.

Expanding beyond behavioral preferences, in **Study 3**, I also investigate how the experience of resolution and discovery is represented in the brain. The experience of reaching resolution through unfolding information may overlap with such experiences of insight or "Aha!" moments, as the exact timing of resolution across unfolding drawings cannot always be well estimated. Drawing from previous research on the experience of insight or "Aha!" moments, the experience of solving an insight problem corresponds to increased activation across a distributed network of brain regions including the insula (Becker et al., 2020), prefrontal cortex (Aziz-Zadeh et al., 2009; Shen et al., 2018), ACC (Aziz-Zadeh et al., 2009), the hippocampus (Shen et al., 2018; Tik et al., 2018), and reward processing regions including the VTA/SN and the NAcc (Tik et al., 2018). Such regions are likely involved in the search for solutions, the formulation, testing, and

updating of hypotheses, and the rewarding experience of reaching the solution. Additionally, curiosity has also been shown to recruit similar regions, including the hippocampus, VTA/SN, and NAcc (Gruber et al., 2014; Jepma et al., 2012; Kang et al., 2009; Lau et al., 2020; Ligneul et al., 2018; van Lieshout et al., 2018), suggesting that the experience of resolution and anticipating information also share similar representation. Within **Study 3**, I expressly categorize and test how different stages of the information-gathering process relate both to the neural correlates supporting the experience of resolution and to later memory formation. Results from this study can inform how evolving external information maps onto the internal development of discovery.

1.4.2 Curiosity shapes sequential search through the development of knowledge.

When information seeking unfolds over multiple choices, how does a learner decide for how long to stay with a current topic? Although the literature on this topic is scant, work has begun to highlight that sequential information seeking can be guided by growing interest and knowledge within a topic. Specifically, curiosity, a sense of knowledge around the topic, and satisfaction are all key factors that predict the length of engagement a learner will invest in a particular topic (Fastrich & Murayama, 2020; E. G. Liquin & Lombrozo, 2021; Wade & Kidd, 2020). To further break this down, the accumulation of knowledge also increases the subjective state of curiosity, suggesting that curiosity uses prior knowledge as a marker of what information is worth continuing to pursue (Fastrich & Murayama, 2020; Wade & Kidd, 2020). If the experience of gaining new information is particularly satisfying, it can further increase curiosity towards related

questions, suggesting that a positive prediction error between informational expectations and outcomes can signal value in deepening further topical knowledge (E. G. Liquin & Lombrozo, 2021). Lastly, waning curiosity or perceiving the end of available information on a topic can signal disengagement and decisions to end information seeking (Fastrich & Murayama, 2020).

Importantly, however, it remains unclear how curiosity for one piece of information can impact the curiosity felt towards seeking other related information. Moreover, during sequential information gathering we are often confronted with momentary experiences of surprise, but how surprise contributes to continued information seeking is less clear. Understanding the influence of each motivation in driving the next choice can give insight into what subjective states encourage the propagation of information seeking. In other words, when and how can curiosity and surprise self-perpetuate learning? In **Study 4**, I explore how both curiosity and surprise can offer complementary signals as to when to continue and when to end information seeking.

1.5 Summary and the current dissertation

The research above provides a detailed account of how motivation can impact and guide the process of sequential information seeking. Despite these advances, several challenges remain in understanding how we organize our information search and the role of different motivations in shaping these decisions. In this dissertation, I will progress our understanding through two main aims. In the first aim, I investigate how we decide how much information is enough through the lens of instrumental (**Study 1**) and

non-instrumental (**Study 2**) sequential information seeking. In the second aim, I investigate how the experience of seeking information across non-instrumental search influences neural activation (**Study 3**) and learning (**Study 4**).

In **Study 1**, I investigate the underlying strategies people use to balance the competing motivations for accuracy and cost minimization during a sequential information seeking task. Here, I used computational models to develop and test a range of strategies that vary in terms of the informational properties used and the computational complexity to assess if people organize their search behavior in optimal ways and whether this changes as a function of the amount of reward that could be obtained.

In **Study 2**, I examine how the motivational state of curiosity guides preferences for how information unfolds across time and the value of discovery. Through designing and creating novel dynamically evolving videos, I challenge the notion that curiosity primarily prioritizes receiving immediate resolution. Moreover, I characterize the affective state of curiosity, providing evidence that the relationship between curiosity and affect may have consequences for search preferences. Finally, I investigate the extent of this effect through two studies that differ in the amount of agency people are given in deciding how long to remain in uncertainty.

In **Study 3**, I explore the neural correlates underlying the process of discovery to build upon our understanding of why curiosity would promote patience and whether the experience of discovery has consequences for memory. Using fMRI, I show how neural activation patterns change when people experience resolution compared to when resolution never occurs. Furthermore, I track the development of resolution through

different phases of the information-gathering process, demonstrating that as information builds, the neural processes involved also change. Lastly, I show how resolution relates to memory outcomes when people experience the videos again.

In **Study 4**, Using an adaptation of a trivia Q&A paradigm, I unpack how information seeking and memory formation can be uniquely driven by curiosity and surprise. I also demonstrate how seeking additional information has implications for memory formation before, immediately, and after a delay. Through this work, I demonstrate that different motivations can drive similar behaviors through distinct mechanisms. Ultimately, this collection of work expands our knowledge of how motivational states titrate sequential information seeking to produce fitted search experiences as well as memories

2. Between heuristics and optimality: Flexible integration of cost and evidence during information sampling

The content from this chapter is verbatim from the below preprint and has only been reformatted for this dissertation: Hsiung, A., Pearson, J. M., Poh, J.-H., Hakimi, S., Alison Adcock, R., & Huettel, S. A. (2022). Between heuristics and optimality: Flexible integration of cost and evidence during information sampling. In *bioRxiv* (p. 2022.05.17.492355). <https://doi.org/10.1101/2022.05.17.492355>

2.1 Introduction

Before making important decisions, humans often collect information about the likely outcomes of different choice options. Consider the choice between two popular restaurants in a new city. Collecting information about both restaurants can increase the likelihood of a positive dining experience but also carries costs (e.g., time spent on evidence accumulation increases the likelihood that the options will become unavailable). Effective decision making thus requires information sampling strategies that balance accuracy and sampling costs – and understanding this balance remains a critical topic for decision science (Averbeck, 2015; Blanchard & Gershman, 2018; Cohen et al., 2007; G. Gigerenzer & Goldstein, 1996; Gold & Shadlen, 2007; Kolling et al., 2012; H. A. Simon, 1990). In particular, how is information about evidence and costs transformed into the decision to *sample* information or *stop*?

Current normative models of sequential information sampling posit that an optimal information sampling strategy should compare the expected values of available actions (continue sampling, choose option A, or choose option B) before selecting the action with the highest expected value (Coenen & Gureckis, 2016; Furl & Averbeck,

2011; Hauser et al., 2017, 2018; Moutoussis et al., 2011). Computing these expected values requires several steps. The decision maker must first determine the expected value of stopping by calculating the probability that each option is correct given the available evidence collected thus far. This estimate must, then, be multiplied by the available reward minus the costs accrued. To determine the value of continuing, the decision maker must estimate the expected values of potential future states, as if an additional sample was drawn. Accurate estimation of these future expected values necessitates an extensive backward induction process (Bellman, 1957) that must be updated with each new sample drawn (Arrow et al., 1949; Furl & Averbeck, 2011; Hauser et al., 2017, 2018; Moutoussis et al., 2011).

Prior work on information sampling has documented that humans sample information sub-optimally, attending to extraneous information (Juni et al., 2016) or through biased weighting of sampling costs (Cisek et al., 2009; Furl & Averbeck, 2011; Hauser et al., 2018). Yet, computing and updating the value of continuing to sample evidence may require significant computational resources, especially for complex decisions, as in sequential information sampling (Bossaerts & Murawski, 2017; Bossaerts et al., 2019; Payzan-LeNestour & Bossaerts, 2011). Indeed, evidence suggests that humans forgo using intensive updating computations, such as Bayesian inference (Charness & Levin, 2005; Gigerenzer & Goldstein, 1996; Payzan-LeNestour & Bossaerts, 2011; Steyvers et al., 2009), even for simpler decisions (Cassey et al., 2016). One factor that might increase the likelihood that humans are willing to expend resources for more optimal computations is the reward value for a correct decision (Bennett et al., 2019; Manohar et al., 2015) but this has not been explored within the context of information sampling. Thus, it remains unclear whether strategies that rely on

computations of expected value reflect human information sampling and whether the use of such computations depends on reward context.

If humans do not follow the computations of an optimal decision maker, what determines when they stop sampling? Early accounts proposed that information search relies on simplified heuristic strategies guided by bounded rationality (Conlisk, 1996; Gigerenzer & Goldstein, 1996; Shah & Oppenheimer, 2008; H. A. Simon, 1990; Herbert A. Simon, 1955; Tversky & Edwards, 1966). In these strategies, a set of rules is established to guide both the process of information acquisition (i.e., what information should be attended to and incorporated as evidence) and the decision to stop sampling. Such heuristic strategies minimize cognitive resource expenditure by leveraging declarative rules; however, by definition, these strategies are less flexible and less adaptable to changes in incoming information or changes in context. While recent accounts have demonstrated support for heuristic-style strategies in information gathering (Baumann et al., 2020; Korn & Bach, 2018; Sang et al., 2020), these studies examined information sampling in contexts where individuals received ongoing feedback about their choices, and it remains unclear whether behavior follows heuristic strategies when individuals must collect and integrate information without ongoing feedback – as in the case for many real-world decisions.

In the present study, we investigated whether humans rely on optimal or heuristic strategies (or their combination) during information sampling, and whether their strategies changed as a function of the reward at stake. We tested several potential models of strategic information sampling that varied in the information used and how that information contributed to the decision process. We found that participants' behavior was best explained by a simple yet flexible strategy in which humans tracked a linear

combination of both the evidence in favor of each category and the accrued costs from sampling – but did not rely on a declarative rule or estimations of expected values. This strategy explained a key pattern we observed in sampling behavior: evidence and costs traded off within but not across trials such that as costs accumulated over a trial, the strength of evidence needed for stopping decreased linearly. Moreover, we found that high-reward contexts neither improved optimality nor impacted which strategy best accounted for participants' decisions. Our results demonstrated how humans implement simple yet flexible information sampling strategies to balance competing motivations for accuracy and cost minimization.

2.2 Methods and Materials

2.2.1 Participants

Participants (N = 105, Mean age = 26.14, SD = 4.79, 69 female) were recruited from the Durham community using flyers and online postings. Our demographic breakdown included 37 participants who identified as White/Caucasian, 47 identified as Asian/South Asian, 14 identified as Black/African American, 3 identified as Hispanic/Latinx, and 4 identified as multi-racial/ethnic. To participate, individuals had to 1) be within the age range of 18-50 years old, 2) have no history of neurological injury or disorders (including seizures and epilepsy), and 3) be fluent in English. Eleven participants were excluded from all analyses, three due to computer error and eight due to having unusable sampling data (failed to sample more than once on over 25% of trials), leaving a final total of 94 participants. All participants received informed consent under the guidelines of the university's Institutional Review Board.

2.2.2 Procedure

At the outset of each experimental session, participants provided informed consent, received task instructions (see Supplementary Methods: Task Instructions) before practicing the experimental task (Fig. 1). Participants returned to the laboratory approximately 24h later for to complete a surprise memory test for the images sampled during the first experimental task. Results for the memory test can be found in the Supplement (see Supplementary Methods: Memory Task, Descriptions, and Findings) but will not be discussed in the main manuscript.

Participants performed a modified version of the Information Sampling Task (Clark et al., 2006) displayed using PsychoPy 2.7 (Peirce et al., 2019). Participants were told that on each trial there was a box that contained 25 images from one of two possible domains: scenes and objects. Each domain had two categories (scenes: indoor or outdoor, objects: living or non-living) and each image belonged exclusively to one category. Each trial contained images from only one domain. Images were all naturalistic photos collected from Google Image searches and scaled to the same size in pixels.

On each trial, participants were tasked with identifying the underlying majority category for a given domain. Participants could sample images from the box serially until they felt they had enough evidence to select a majority category (max of 25 images per trial). Participants were told that there would always be a majority category, but they were not told the true proportions of each image category and were instructed that the proportions could change between trials. The true proportion was kept constant at 60/40 for majority/minority categories (i.e., 15 of the 25 images would be from the majority category). The order of the images was randomized. Participants performed trials under high (\$5.00) and low (\$1.00) reward stakes. Incorrect responses in both stakes

conditions resulted in a reward of \$0.00 for that trial. In addition, participants incurred a cost for each sample they made (2% of the max reward they could earn for that trial). Thus, participants had to balance their confidence in identifying the true majority against accruing sampling costs.

At the start of each trial, a *cue screen* (2000 ms) appeared, informing participants of the image category judgment (e.g. indoor vs. outdoor or living vs. non-living) as well as the monetary reward available for a correct response (e.g. Correct Response = \$1.00/Correct Response = \$5.00, before sampling costs). They then viewed the *sample screen*, whereupon they had the option to either sample an image or make a final choice as to what category they thought predominated on the trial. If they chose to sample (by selecting the down arrow key), one image would immediately appear in the middle of the screen for 2500 ms (*image screen*). After the image disappeared, participants were returned to the *sample screen*. Images did not stay visible to participants after the 2500 ms presentation; thus, participants had to remember past images to guide their choices. At each instance of the sample screen, participants had 15 seconds to make a choice before they automatically advanced to the next trial, with the previous trial being marked as incorrect. This happened on approximately 0.003% of trials across all participants (17 out of 5004 trials).

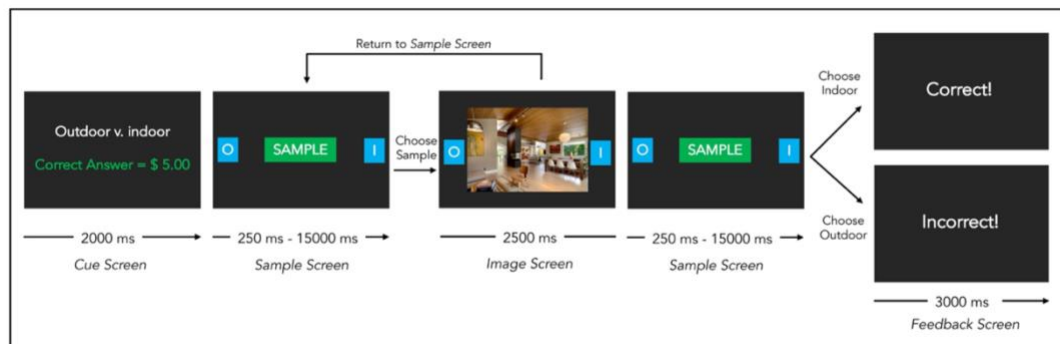


Figure 1: Task Design. Participants completed a modified version of the Information Sampling Task.

Participants were free to sample as few or as many images as they deemed necessary to guess the more prevalent category. When participants decided to stop sampling, they indicated their decision about which category they felt predominated on that trial by choosing the box (by pressing either the right or left arrow key) that was associated with that category, which were displayed on either side of the sample button throughout the trial. After participants made their final choice, they were shown a *feedback screen* (2000 ms) that displayed if their guess matched the true majority in the box (e.g. "Correct"/ "Incorrect").

Participants completed 48 trials in the task. Trials were fully counterbalanced such that they saw an equal number of trials from either category, and each category was equally represented in both high and low reward stakes. Additionally, each category had the same overall probability of winning. To ensure incentive compatibility, participants were paid for 4 trials, randomly chosen. Because the task was self-paced and participants varied in how many images they collected, the session length ranged from 11 minutes to 47 minutes (Mean time: 24.86 minutes, SD: 7.63 minutes).

2.2.3 Computational Modeling

To understand how participants determined when to switch from gathering information to selecting a final choice, we compared participants' behavior using a series of computational models. We first measured how close each participant's stopping choices were to the Ideal Observer (model-predicted optimal choices). We then fit each subject's behavior to four sampling strategies. The first strategy, Expected Value Urgency Threshold (EV-UT), relied on expected value computations to inform choices. We used an adaptation of this strategy similar to Hauser et al., (2018), that suggested humans integrate costs non-linearly. In this strategy, the threshold to transition from sampling to selecting an option was both non-stationary and inconstant across the number of images collected. The second strategy, Evidence Difference Heuristic (ED), was a heuristic that assumed participants tracked the continuous signed difference in evidence between the two categories towards a threshold (e.g., "I sample until one category has 4 more than the other"). This strategy suggests that participants approach information sampling insensitive to the number of images sampled and implies that the stopping boundary is stationary and constant across sampling. The third strategy, Costs Accrued Heuristic (CA), was another heuristic that assumed, participants used the continuous number of samples drawn and only a binary representation of the difference in evidence to inform choice (e.g. "I sample 5 images and then choose the majority). Similar to the first, this strategy maintained a stationary threshold that triggered a decision to select an option but implied that the magnitude of evidence mattered less. The last strategy, Evidence-by-Costs Tradeoff (ED-CA), was a combination of the two heuristic approaches, such that participants used both continuous representations of the difference in evidence between the categories and the number of images collected to

inform their choices. This strategy reflected a linear threshold that decreased as the number of samples collected increased, representing a non-stationary but constant tradeoff between the two informational sources. Detailed descriptions of the strategies are outlined below. In all models, choices were assumed to be probabilistic and were all fit using a SoftMax function. To emphasize, participants were given the following information: each box on each trial contains a total of 25 unique images, the maximum reward value for a trial is either \$5.00 or \$1.00 and the cost per image is a constant 2% of the maximum reward available on a trial (\$0.10 for \$5.00, \$0.02 for \$1.00), the proportion of images from either category is specifically withheld and participants are told that the proportion may change on a trial-by-trial basis.

2.2.3.1 Optimal Behavior

To compare participant sampling behavior to that of an Ideal Observer, we first calculated the optimal stopping points using a model adapted from Hauser et al. (2018). Because the proportion of reward to costs was equivalent for high vs. low stake trials, the computations and optimal stopping points are the same across reward conditions. After each sample (N_{samp}), the optimal agent compares the value of stopping given the current evidence against the value of continuing to sample. In order to determine the value of each action, the agent computes the probability of success in selecting the correct category given the current evidence (i.e., the number of indoor (n_i) and outdoor samples $(N_{samp} - n_i)$ collected thus far). Because the true underlying distribution of indoor to outdoor images is unknown, the optimal agent must also estimate the underlying distribution (q) from which the samples are being drawn from. Then, it must

compute the probability of success under each possible proportion of majority to minority images weighted by the likelihood that that is the true distribution (Eq. 1.1, 1.2). We set prior beliefs, $\bar{\alpha}$ and $\bar{\beta}$, about the true underlying distribution equal to 1.

$$P(\text{Indoor}|n_i, N) = \sum_{\text{Indoor} > \frac{N_{tot}}{2}}^{N_{tot} - (N - n_i)} \int_0^1 P(\text{Ind}|q, n_i, N) P(q|n_i, N) dq \quad (1.1)$$

$$= \sum_{\text{Indoor} > \frac{N_{tot}}{2}}^{N_{tot} - (N - n_i)} \binom{N_{tot} - N}{\text{Ind} - n_i} \frac{B(Y + \alpha, N_{tot} - Y + \beta)}{B(\alpha + n_i, \beta + N - n_i)} \quad (1.2)$$

The expected value of stopping is then computed by taking the probability of success of stopping multiplied by the reward (\$5.00 for high stakes, \$1.00 for low stakes) minus the accrued costs, c , per sample (\$0.10 for high stakes, \$0.02 for low stakes) (Eq. 2).

$$\begin{aligned} Q(\text{Indoor}|n_i, N) &= R_{correct} P(\text{Indoor}|n_i, N) - cost * N_{samp} \\ Q(\text{Outdoor}|n_o, N) &= R_{correct} P(\text{Outdoor}|n_o, N) - cost * N_{samp} \end{aligned} \quad (2)$$

The expected value of stopping is then compared to the expected value of continuing to sample. To compute the expected value of continuing to sample, the optimal agent calculates the expected value of stopping for each state using backward induction to solve for the Bellman equation (Bellman, 1957). Briefly, the expected value of continuing at timepoint 25 is equal to 0 because no additional samples can be drawn. Thus, the expected value at timepoint 25 is equal to the expected value of stopping given all available evidence. Given a behavioral policy that always chooses the highest valued action, the value of all possible states at timepoint 24 (and prior timepoints) can

then be calculated using backward induction. Thus, for each possible state, the expected value of continuing, averages over all potential future states, weighting them by the likelihood that that state will be reached (Eq. 3). s' represents the next immediate state, which can either reveal another indoor image ($i = 1$) or an outdoor image ($i = 0$).

$$Q(\text{Continue}|n_i, N) = \sum_{\text{all possible } s'} P(s'|n_i, N)V(s') \quad (3)$$

To examine how participants' behavior compared to optimal behavior, we binned each trial for each participant as either optimal, under-sampled, or over-sampled based on where each stopping decision fell compared to optimal. Because of the cost and reward structure of this specific task, optimal behavior followed an easily verbalized heuristic of "sample until a difference of 3 is achieved." This heuristic fits with "fast and frugal" criteria of being computationally simple and relying on only a fraction of available information but still performing optimally (Gigerenzer & Goldstein, 1996; Todd & Gigerenzer, 2000). Thus, optimal behavior could be achieved through multiple routes of computation.

$$c_{\text{per step}} = R * 0.02 * n$$

$$c_{\text{per step}} = R * 0.02 * n * t \quad (4, 5, 6)$$

$$c_{\text{per step}} = \frac{0.02 * R}{1 + e^{-10(n-p)}}$$

2.2.3.2 Expected Value Computation Strategy.

We examined a probabilistic modification of the optimal strategy. This strategy relied on the same Bayesian updating and inference to estimate the probability of success given the available evidence but was adapted to allow for inherent noise in

participant decision making as well as to test different cost accrual mechanisms (see Supplementary Methods: Expected Value Model Formulation). Prior research has documented that human deviation from optimality could arise from the accumulation of costs that are different from the specified objective sampling costs (Cisek et al., 2009; Ditterich, 2006; Hauser et al., 2017, 2018). We therefore tested whether the cost per step (c_{step}) was being subjectively accrued in either a linear (Eq. 5), or non-linear (sigmoidal, Eq. 6) manner, and if these outperformed the use of objective costs (Eq. 4). In equations 5, t represents the subjective scaling of objective costs. In equation 6, p represents the sample number where costs begin to accumulate. In all equations, R represents the reward condition, 0.02 represents the percentage of the max reward, which equates to the objective cost per sample. Overall, our non-linear cost accrual outperformed our other two models of cost (see Supplementary Methods: All Sampling Strategies).

To test the different models of cost, we isolated the impact of costs to the choice to continue sampling. To do so, we updated the action values for choosing each final option as well as the value of continuing to sample as such.

$$Q(Indoor|n_i, N) = R_{correct}P(Indoor|n_i, N) \tag{7}$$

$$Q(Outdoor|n_o, N) = R_{correct}P(Outdoor|n_o, N)$$

$$Q(Continue|n_i, N) = -c_{per\ step} + \sum_{s'=\begin{cases} n_i+x \\ N+1 \end{cases}}^{x=[0,1]} P(s'|n_i, N)V(s') \tag{8}$$

These expected values were then transformed into probabilities using the following Softmax function with inverse temperature parameter, β , and irreducible noise parameter, ξ (Eq. 9). Importantly, this first family of models relied on the assumption that humans were still performing the underlying Bayesian operations to determine their policies, albeit with noise in their choice process.

$$\pi(\text{Choice}|n_i, N) = \frac{e^{Q(\text{Cont.}|n_i, N)\beta}}{e^{Q(\text{Cont.}|n_i, N)\beta} + e^{Q(\text{Ind.}|n_i, N)\beta} + e^{Q(\text{Out.}|n_i, N)\beta}} (1 - \xi) + \frac{\xi}{3} \quad (9)$$

For all of the models tested within the Expected Value Computation framework, participant data was best fit by the subjective non-linear cost model, giving rise to the Expected Value Urgency Threshold strategy and replicating previous work (Hauser et al. 2018). Given our two reward contexts, we also tested whether participants adapted this strategy based on the reward available for that trial. To do so, we tested three separate modifications of the subjective non-linear Expected Value Computation strategy. In our first model, we fit separate models for each reward condition for each participant, suggesting that participants could have completely different parameter values for each reward condition. In our second model, we fit one model for both reward conditions and included a parameter that scaled the reward value for low-reward trials to be between \$1.00 and \$5.00, suggesting that the parameter values for both conditions could be equivalent, but participants were still sensitive to the difference in reward outcomes. Our last model either through the same model under just one of the high reward condition models treated all trials as operating under the high-reward conditions and fit one set of parameters for all trials. This was our best fitting model, as such the

best model from the Expected Value Computation strategy was one that included a subjective non-linear cost accrual and treated high and low-rewarded trials as the same (see Supplementary Methods: All Sampling Strategies).

2.2.3.3 Evidence Difference Heuristic Strategy.

Our second model, Evidence Difference Heuristic (ED), was a heuristic that assumed participants tracked the continuous signed difference in evidence between the two categories towards a threshold (e.g., “I sample until one category has 4 more than the other”). This strategy suggests that participants approach information sampling insensitive to the number of images sampled and implies that the stopping boundary is stationary and constant across sampling (Baumann et al., 2020; Korn & Bach, 2018; Shah & Oppenheimer, 2008; Herbert A. Simon, 1955; Tversky & Edwards, 1966). To fit this heuristic, we adapted the rule into a probabilistic account that used the signed difference in evidence drawn to predict choice. The signed difference in evidence between the current majority and minority in the samples collected at each timepoint was submitted to a multinomial SoftMax regression along with a subject-specific intercept, β_0 , in order to produce a probability for each action (Eq. 10).

$$P(Y_i = \text{Choose Indoor}) = \frac{e^{\beta_0 + \beta_{in} * X_{in}}}{1 + \sum_{k=1}^{K-1} e^{\beta_k * X_k}} \quad (10)$$

2.2.3.4 Sample Number Heuristic Strategy.

Our third model, Sample Number Heuristic (SN), was another heuristic that assumed, participants used the continuous number of samples drawn and only a binary representation of the difference in evidence to inform choice (e.g. “I sample 5 images and then choose the majority). Similar to the first, this strategy maintained a stationary

threshold that triggered a decision to select an option but implied that the magnitude of evidence mattered less. Identical to our Evidence Difference Heuristic Strategy, to fit this heuristic, we adapted the rule into a probabilistic account that used the number of samples drawn and a binarized difference in evidence to predict choice. These variables were submitted to a multinomial SoftMax regression along with a subject-specific intercept, β_0 , in order to produce a probability for each action (Eq. 11).

$$P(Y_i = \text{Choose Indoor}) = \frac{e^{\beta_0 + \beta_{in} * X_{in} + \beta_{samp} * X_{samp}}}{1 + \sum_{k=1}^{K-1} e^{\beta_k * X_i}}, \text{ where } X_{in} = 0 \text{ or } 1 \quad (11)$$

2.2.3.5 Evidence-by-Costs Tradeoff (ED-CA) Strategy.

Our fourth model was built on the assumption that participants' decisions to continue sampling or stop and commit to a category could be described by a strategy that depended on multiple forms of information but did not require the computational complexity of optimal strategies. Specifically, the Evidence-by-Costs Tradeoff (ED-CA) strategy, was a combination of the above heuristic strategies, such that participants used both continuous representations of the difference in evidence between the categories and the number of images collected to inform their choices. This strategy reflected a linear threshold that decreased as the number of samples collected increased, representing a non-stationary but constant tradeoff between the two informational sources. To fit this model, both sample number and the signed difference in evidence were submitted to a multinomial SoftMax regression along with a subject-specific intercept, β_0 , in order to produce a probability for each action (Eq. 12).

$$P(Y_i = \text{Choose Indoor}) = \frac{e^{\beta_0 + \beta_{in} * X_{in} + \beta_{samp} * X_{samp}}}{1 + \sum_{k=1}^{K-1} e^{\beta_k * X_i}} \quad (12)$$

Similar to the Expected Value Computation model, we were also interested in whether participants treated information similarly across the different reward conditions. To test if reward significantly changed model fits, we tested two different model iterations. First, to test if participants were using completely difference parameter estimates for the different reward conditions, we split participant trials into high- and low-reward trials and fit each subset of trials to our SoftMax multinomial regression (Eq. 12). We then compared the difference in parameter values for high- vs. low-rewards. Second, to test if reward independently modified choice but did not impact the weight of individual information quantities, we added an additional reward parameter into the original SoftMax multinomial regression. Interestingly, parameter values were comparable across the two reward conditions and adding reward as an independent parameter did not improve model fit beyond Eq. 12. Thus, the best fitting model from the ED-CA Tradeoff strategy was one that also treated high and low-rewarded trials as the same (see Supplementary Methods: All Sampling Strategies).

2.2.3.6 Model Comparison.

For each model, we optimized the parameters to maximize the log likelihood for each participant individually. We used SciPy's standard optimize.minimize function to minimize the negative loglikelihood of the observed choices. Parameters for our Optimal Stopping were bounded based on previous studies (ρ : [0-25] for sigmoidal, p : [0, 0.2] for linear, β : [1,10], ξ : [0, 0.5]) (Hauser et al., 2018) and both our Heuristic and Evidence-by-Cost models were bounded based on preliminary mixed effects multinomial

regression [β_1 (samples drawn): [-1,5], β_2 (evidence difference): [-4,8], β_3 (reward context): [-5,8]). In every case, we ensured the best fitting parameters each fell within these boundaries. We fit each participant 10 times per model to ensure convergence and stability of best fitting parameters.

To compare participants' fits from our models, we first took the top performing models from each strategy if a strategy had more than one iteration before examining cross-group comparisons. All models in the final group were compared using both Akaike Information Criterion (AIC) (Akaike, 1974) and Bayesian Information Criterion (BIC) (Schwarz, 1978) scores. To examine patterns of best model fits on the group level we ran a repeated measures ANOVA to determine if participant-specific AIC or BIC scores differed significantly amongst models. Distributions of AIC and BIC scores per top performing can be found in the Supplement (see Supplementary Note 3, Figure 2).

2.3 Results

We tested participants on a modified version of the Information Sampling Task (Fig. 1) (Clark et al., 2006). Participants viewed a series of images randomly drawn from a pool of 25 images. The pool contained images from two categories (e.g., indoor vs. outdoor scenes), with one category comprising 60% of the images and the other comprising 40%. Importantly, participants were unaware of the true proportions of each image category for each trial, although they were told that there would always be a majority category. Participants attempted to identify which category was more prevalent on each trial, under either high (\$5.00) or low (\$1.00) reward stakes for correct answers. Each image participants chose to draw came with a sampling cost of 2% of the maximum reward on that trial (i.e. \$0.10 for a \$5.00 trial; \$0.02 for a \$1.00 trial). Thus,

participants had to balance competing goals: sampling more images could increase the accuracy of their guesses, but they would win less reward overall due to the increasing cost accrued (see Supplementary Methods: Task Instructions).

2.3.1 Greater sampling is associated with higher task accuracy, but at the expense of greater cost accumulation

We first investigated how well participants performed the task. Participants correctly identified the majority category 79% of the time (SD = 10%) (Fig 2a). Across the entire task, participants accumulated an average of \$95.07 (SD = \$10.52) but were only paid for a randomly selected subset of those trials (see Supplementary note 1 and Supplementary Fig. 1). On average, participants viewed 7.82 images (SD = 2.89 images, Range = 1 – 24 images) and reached an average difference in evidence between the currently held majority and minority category of 2.61 images (SD = 0.62, Range = 0 – 8 images) before selecting a majority category (Fig. 2a-c).

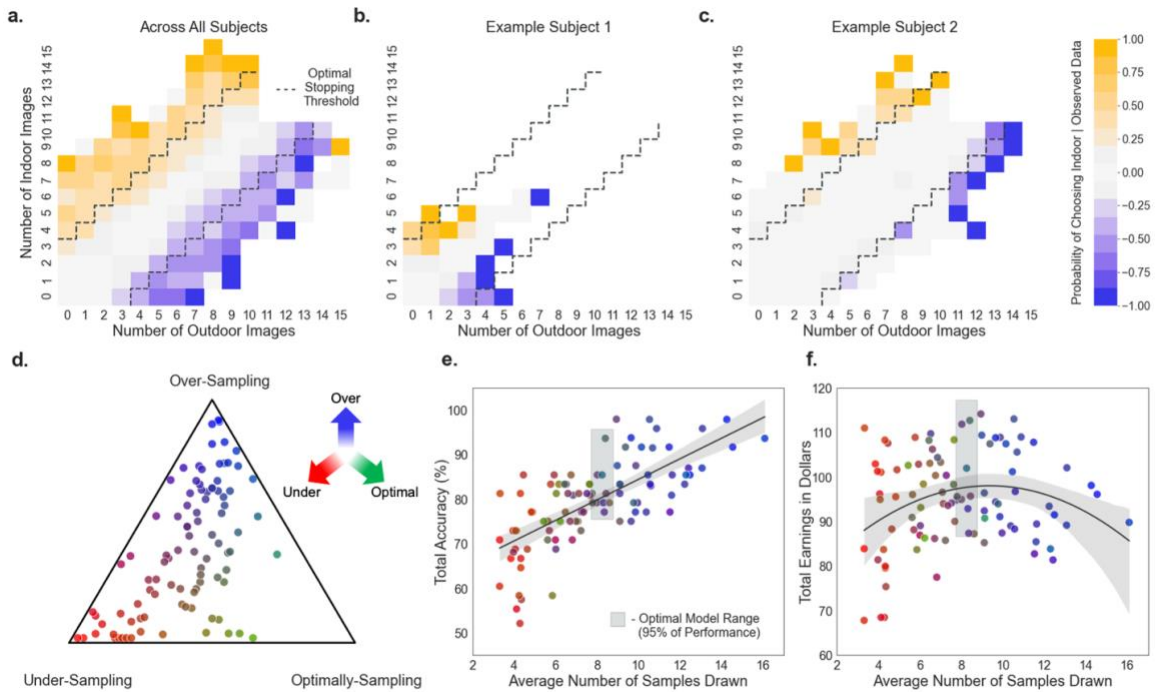


Figure 2: Sampling tendencies between participants relate to overall task accuracy and total earnings.

We next examined whether performance outcomes were predicted by participants' relative tendency to sample. Across participants, a higher average number of samples predicted better task accuracy (Fig. 2e) ($F(2, 91) = 43.5, p < 0.001, R^2 = 0.48, 95\% \text{ CI } [1.47, 3.09]$). Similarly, the average number of samples also predicted total earnings via an inverted-U shaped function (Fig. 2f) ($F(3, 90) = 3.239, p = 0.026, R^2 = 0.07, 95\% \text{ CI (quadratic) } [-49.91, -8.39]; \chi^2(91, N = 94) = 7.785, p = 0.006$). Thus, participants whose sampling was, on average, much lower or much higher than average tended to have lower overall earnings. We repeated this analysis at the trial-level within participants and found consistent results. The number of samples drawn on a given trial predicted both accuracy ($\beta_{samples} = 0.0818, t = 7.49, p < 0.001, 95\% \text{ CI } [0.06, 0.10]$) and

earnings ($\beta_{samples} = -6.49, t = -4.38, p < 0.001, 95\% \text{ CI } [-9.22, -3.73]$; $\beta_{samples^2} = -4.44, t = -3.46, p < 0.001, 95\% \text{ CI } [-6.98, -1.93]$). Overall, a higher average number of samples collected prior to stopping resulted in better task accuracy but also greater accumulation of sampling costs – leading to lower earnings.

To compare participant performance to an optimal decision maker (Fig. 2e-f), we computed optimal choices using an Ideal Observer model (see Supplementary Methods: Ideal Observer Model). We used this model to label each trial for each participant as optimal (matched the choice made by the Ideal Observer), under-sampled (stopped sampling earlier than the Ideal Observer), or over-sampled (stopped sampling after the Ideal Observer). We then created composite scores for each participant to assess how close participants were to optimal behavior (Fig. 2d). Collectively, participants performed worse than the Ideal Observer, both in accuracy (two-sided t-test: $t(186) = 6.33, p < 0.001, M_{ideal} = 85.6\%, SD = 5.2\%$) and in earnings (two-sided t-test: $t(186) = 6.02, p < 0.001, M_{ideal} = \$101.97, SD = \$7.67$), but they did not differ significantly from the Ideal Observer in the average number of samples drawn (two-sided t-test: $t(186) = 1.38, p = 0.169, M_{ideal} = 8.24 \text{ images}, SD = 5.21$). Upon further examination, however, this was due to participants either over- or under-sampling relative to optimality ($F(5, 530) = 20.54, p < 0.001$), indicating that participants are either estimating optimal behavior poorly or relying on a different process to establish stopping criteria.

2.3.2 Models of information sampling strategies

Our next analysis investigated how participants sampled and integrated information towards a decision. We first identified four potential information sampling strategies that relied on a range of optimal and heuristic approaches (see Supplemental

Methods: All Sampling Strategies). Our first strategy was a probabilistic modification of the optimal Ideal Observer (see Supplementary Methods: Expected Value Model Formulation). This strategy relied on the same Bayesian updating and inference to estimate the expected value of each option but was adapted to allow for inherent noise in participant decision making as well as to account for cost accrual mechanisms that deviated from objective costs accumulation (Cisek et al., 2009; Ditterich, 2006; Furl & Averbeck, 2011; Hauser et al., 2018) (see Methods). The cost mechanism that best accounted for participant behavior was similar to the strategy from Hauser et al., (2018), in which subjective costs were accumulated nonlinearly, representing a growing urgency across information sampling to select a final option (Cisek et al., 2009). As such, the Expected Value Urgency Threshold (EV-UT) strategy predicted that participants would sample until the value of an option surpassed the value of continuing to gather information, and that the value of continuing would sharply decline after the urgency threshold was met (Fig. 3, EV-UT).

The next two strategies were probabilistic adaptations of two common heuristics. The first, Evidence Difference Heuristic (ED), was a heuristic that assumed participants tracked the continuous signed difference in evidence between the two categories towards a threshold (e.g., “I sample until one category has 4 more than the other”). This strategy predicted that participants approach information sampling insensitive to the number of images sampled and implies that the stopping boundary is stationary and constant across sampling (Fig. 3, ED). The second, Costs Accrued Heuristic (CA), was a heuristic that assumed, participants used the continuous number of samples drawn as a proxy for the costs accrued from sampling and only a binary representation of the difference in evidence to inform choice (e.g. “I sample 5 images and then choose the

majority). Similar to the first, this strategy predicted a stationary threshold that triggered the end of sampling, but now, bound to the number of samples drawn, implying that the magnitude of evidence mattered less (Fig. 3, CA).

The last tested strategy, Evidence-by-Costs Tradeoff (ED-CA), was a combination of the two heuristic approaches, such that participants used continuous representations of both the difference in evidence between the categories and the number of images collected to inform their choices. These two quantities were then linearly combined towards a threshold. This strategy predicted that as the number of samples collected increased, the difference needed between the evidence in favor of one category over the other would decrease, representing a non-stationary but constant tradeoff between the two informational sources (Fig. 3, ED-CA).

Across these four strategies, there were striking visual distinctions in the types of choices each predicted (Fig. 3, Model Example). Furthermore, models were distinguishable from one another, as shown by model recovery such that choices generated by each model were best fit by the model that generated them (see Supplemental Methods: Parameter Recovery and Model Recovery). Thereby, we were not only able to detect descriptive differences in sampling strategies but were also able to draw inferences about the underlying process guiding information sampling (Fig. 3).

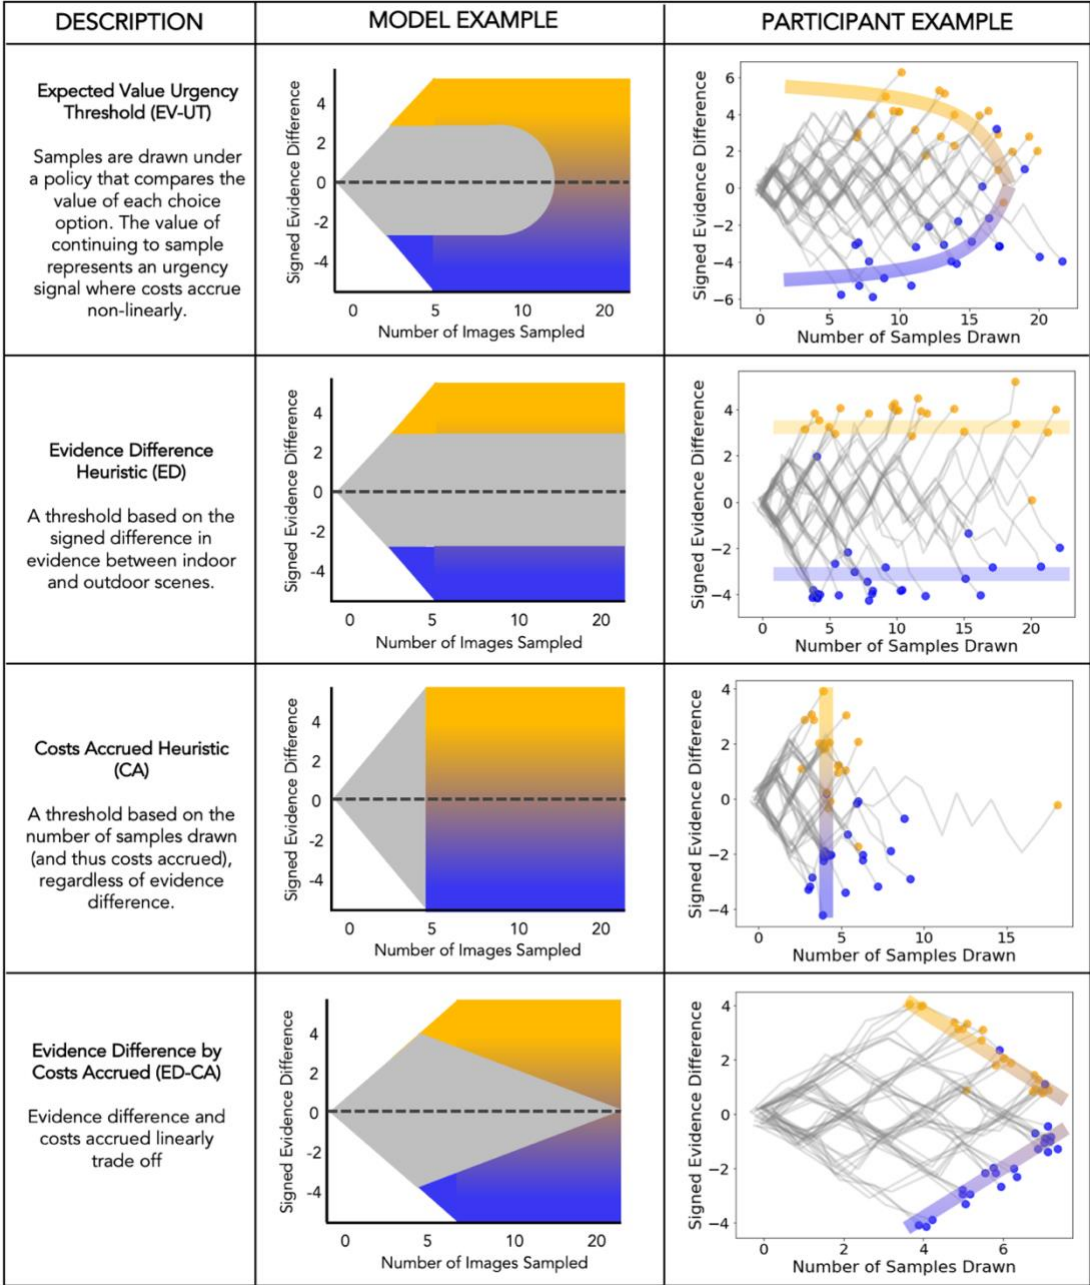


Figure 3: Tested information sampling strategies with model predictions and example representations.

2.3.3 Moment-by-moment sampling decisions were best predicted by an Evidence-by-Costs (ED-CA) strategy

We next fit each participant's data to each of the four sampling strategies outlined above. Participants overall were best fit by the Evidence-by-Costs Tradeoff (ED-CA), which used continuous representations of both the difference in evidence between the categories and the number of images collected to inform their choices. This strategy outperformed the three other proposed strategies such that of our 94 participants, 83 were best fit by the ED-CA strategy (Fig. 4a). A repeated-measures ANOVA confirmed that the ED-CA strategy had significantly lower BIC scores ($F(3, 279) = 81.17, p < 0.001$) compared to the CA strategy ($t(93) = -14.65, p < 0.001$), the ED strategy ($t(93) = -10.92, p < 0.001$), and the EV-UT strategy ($t(93) = -11.32, p < 0.001$) (Fig. 3b-d). There was no difference in fit between the ED strategy and the EV-UT strategy ($t(93) = 1.17, p = 0.24$). The CA strategy fit the worst of all four strategies tested (ED and CA: $t(93) = -4.83, p < 0.001$; EV-UT and CA: $t(93) = -5.40, p < 0.001$).

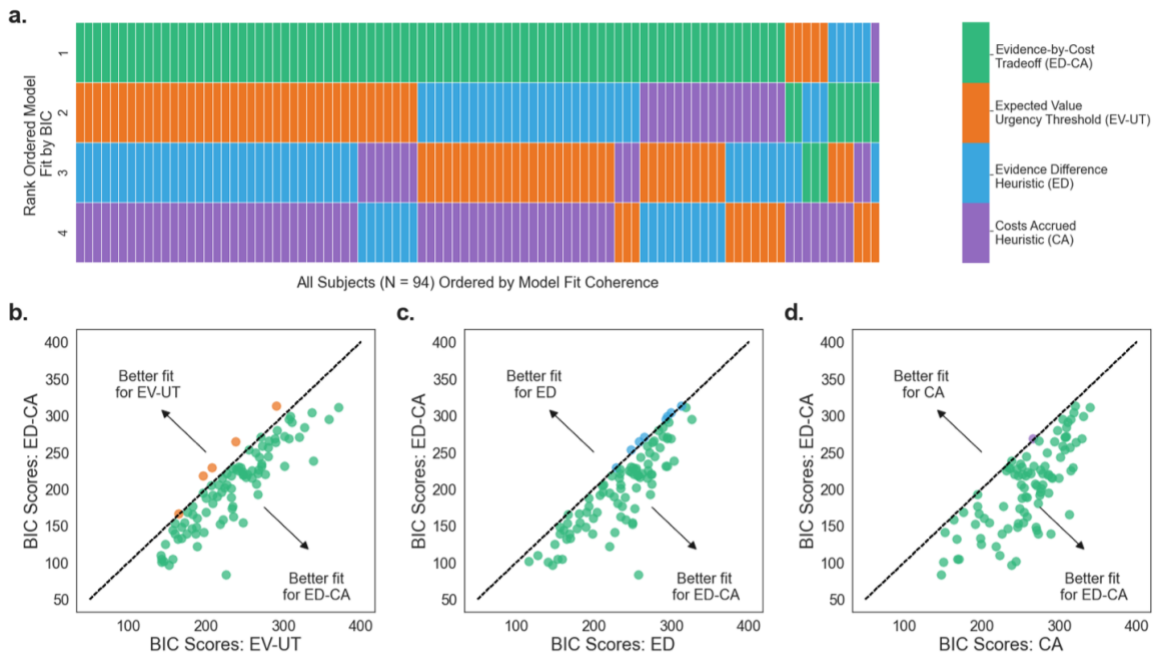


Figure 4: Evidence-by-Costs (ED-CA) strategy outperforms both simpler heuristics and expected value models in predicting decisions.

Results from the ED-CA strategy revealed that participants were sensitive to the costs of increased sampling (i.e., number of samples drawn) and the magnitude of evidence favoring one category over the other (i.e., signed difference in samples). We next sought to test whether there was a linear tradeoff between the difference in evidence for each category and the number of samples.

2.3.4 Stopping criteria were non-stationary within a trial but stationary across trials

The Evidence-by-Costs (ED-CA) strategy aggregates the currently available information and compares it to a decision threshold that reflects a participant-specific level of information needed to terminate sampling. A prediction of this strategy is that, over the course of a trial, the contributions of total samples drawn and evidence difference can trade off against each other. We found that the number of samples drawn

negatively predicted evidence difference at the time of stopping ($\beta_{samples}: -0.0722$, $t(66.6) = -6.93$, $p < 0.001$ CI [-0.09, -0.05]) and that the majority of our participants (N=78/94) evinced this negative relationship (Fig. 5). This result indicates that the more images sampled (and thus the higher the costs), the less the evidence difference must be in order to stop sampling, consistent with accounts of collapsing boundaries in evidence accumulation (Ditterich, 2006; Drugowitsch et al., 2012; Malhotra et al., 2017; Murphy et al., 2016). We tested whether this relationship was better fit with the inclusion of a quadratic term, but that inclusion did not significantly improve model fit ($\chi^2(1, 94) = 2.94$, $p = 0.09$). Interestingly, the initial threshold for stopping is similar to that of the Ideal Observer model but quickly decreases as the trial progresses.




Figure 5: Evidence difference and number of samples trade off.

We also examined whether stopping criteria changed across trials in the task. We first investigated whether there were changes in stopping criteria (i.e., changes in participant's average sampling number or evidence difference prior to stopping) or changes in stopping consistency between the first and second half of the task (i.e., changes in the standard deviation of samples drawn or evidence difference). None of the dependent measures were significantly different from the first to the second half of the experiment (all $p > 0.05$). Additionally, there were no significant differences in accuracy from the first to the second half of the task ($p > 0.05$), indicating that performance was relatively stable from trial to trial. Moreover, none of the dependent measures were significantly different as a function of the outcome of the previous trial (all $p > 0.05$). Collectively, these results suggest that behavior in this task was stable across trials.

2.3.5 Increasing reward stakes does not encourage use of more optimal strategy

Lastly, we investigated whether an increase in reward stakes shifted strategy use across our participants. We hypothesized that under high reward contexts, participants might expend more computational resources to sample more optimally (Bennett et al., 2019; Manohar et al., 2015), thus participants' behavior would be more similar to the Ideal Observer or be better fit by a strategy that relies on optimal value estimation (e.g., the EV-UT strategy). To test this, we split trials by reward condition within each participant and fit each subset of behavior to the ED-CA strategy and to the EV-UT strategy. We found that overall, the ED-CA strategy provided the best fit for the majority of our participants ($N=69$) across both high-reward and low-reward contexts. Of the participants that had different fits between high-stakes and low-stakes trials, there was not a consistent pattern of change (EV-UT strategy provided a better fit under high-

reward trials for 7 participants and provided a better fit under low-reward trials for 12 participants) (See Supplementary Note 4, Fig. 3 for more detail).

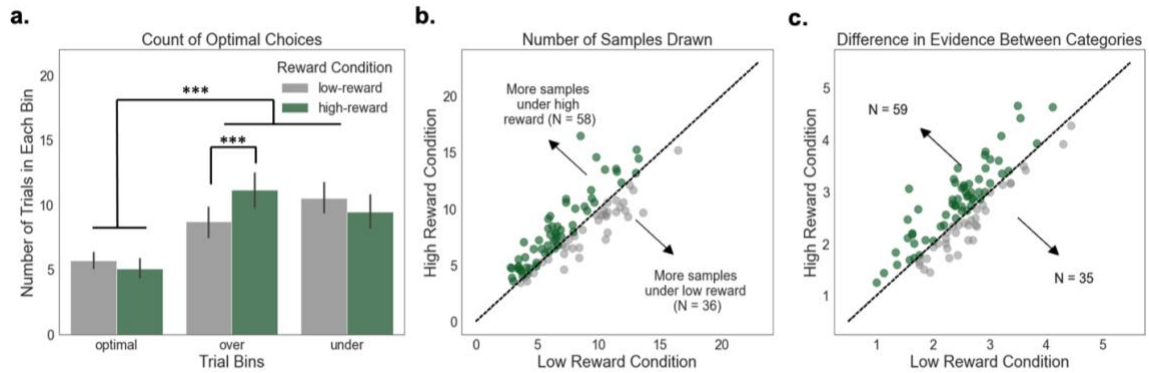


Figure 6: High reward stakes increase sampling but do not encourage more optimal behavior.

Similarly, changes in sampling as a function of reward condition did not encourage optimal behavior compared to the Ideal Observer. We found no difference in the proportion of trials in which participants sampled optimally across reward conditions ($t(441) = 0.847, p = 0.958$). Instead, we see an increase in the relative number of over-sampling trials in high-reward contexts, compared to low-reward contexts ($t(446) = 3.068, p = 0.023$) (Fig. 6a). Follow-up trial-by-trial analyses confirmed that participants

sampled more images under high-reward compared to low-reward ($\beta_{reward} = 0.503, t = 4.762, p < 0.001$ CI [0.29, 0.71]) (Fig. 6b) and required a larger evidence difference to stop ($\beta_{reward} = 0.147, t = 4.428, p < 0.001$ CI [0.08, 0.21]) (Fig. 6c). Despite the

adjustments in sampling behavior under high-reward compared to low-reward, there was no difference in performance outcomes between the two conditions ($\beta_{reward} = 0.02, t = 0.25, p = 0.803$). High-reward trials, thereby, did not encourage more optimal choices but

instead, drove participants to simply adjust the threshold for stopping within their existing strategy. These adjustments, however, ultimately did not improve performance on high-reward trials. Our reward manipulation provides support that participants adopt a global policy for stopping behavior – and adjust that policy in different reward conditions – but do not systematically deploy different policies in different reward contexts.

2.4 Discussion

Balancing the costs of gathering information with the desire to accurately choose outcomes represents a fundamental challenge for human decision-making. Various theoretical models of information sampling have been developed to explain how humans address this challenge, but these models tend to either emphasize resource-intensive optimal computations or efficient but rigid heuristics. Our results support an alternative perspective: information sampling relies on a computationally simple and flexible strategy that accumulates task-relevant information until a threshold is reached. This Evidence-by-Costs strategy also predicted other key features of sampling behavior, including the tradeoff such that more evidence was required when costs were low but that less evidence was required when accrued costs were high. In addition, we found that changing the reward value for a correct choice did not encourage more optimal behavior nor a switch to using more computationally expensive estimations. Collectively, we concluded that humans flexibly adapt the information accumulation process to balance competing motivations for accuracy and cost minimization.

Previous work has demonstrated that humans exhibit subjective biases in the computation and integration of accruing costs that limit the optimality of behavior (Cisek et al., 2009; Coenen & Gureckis, 2016; Furl & Averbeck, 2011; Hauser et al., 2017,

2018). However, these policies still assume that humans engage in the computational expensive estimations of the value of stopping and the value of continuing in the same manner as an Ideal Observer (Bossaerts & Murawski, 2017; Payzan-LeNestour & Bossaerts, 2011). The Evidence-by-Costs strategy provides a starkly simpler rule that reduces computational demands through limits both on the number of steps involved and the amount of information held in memory (Bossaerts & Murawski, 2017; H. A. Simon, 1990). It also provides a better account of behavior: instead of computing the probability of success in accordance with Bayesian inference, participants relied on a mechanism that estimated success by tracking the evidence difference between the two categories of images.

The use of an Evidence-by-Costs strategy was evident in the trial-by-trial behavior of our participants. Consistent with prior work (Ditterich, 2006; Malhotra et al., 2017; Murphy et al., 2016), we found that there was a tradeoff between two quantities: as the number of samples (costs accrued) increased, the smaller the evidence difference needed to stop sampling. The existence of this tradeoff could reflect two possibilities. First, as more samples are drawn, individuals could be relying on smaller but more reliable differences between the majority and minority category. Although we cannot rule this possibility out completely, it is unlikely given that participants were kept blind to the true underlying distribution and thus the proportion of the samples not yet drawn was also unknown. Alternatively, this tradeoff could mean that the stopping threshold may be reached from different linear combinations of two quantities: a large number of samples or a large evidence difference, and individuals actively use both to inform stopping. The result is that as sampling continues and costs are accrued, people

begin to prioritize different information – moving from an initial bias toward larger evidence differences toward a later bias against costs.

Our results extend prior research positing tradeoffs between evidence and sampling costs, akin to the speed-accuracy tradeoff, to determine choice. Murphy et al. (2016) saw this tradeoff under conditions in which participants were pressed for time, yet our results suggest that this trade-off can exist even without external time pressures. Additionally, Malhotra et al. (2017) found that the tradeoff between costs and evidence only existed when participants completed intermixed trials of varying difficulty, yet we show that the trade-off persists even when difficulty is fixed across trials. It is possible that our task, similar to that of Malhorta et al. (2017), had increased uncertainty because the distribution of images from either category was unknown to our participants. This could have encouraged a strategy that did not rely on a single information factor (e.g., evidence) to inform stopping. Additionally, our tradeoff appears linear in form, as opposed to exponential or sigmoidal, which are indicative of a growing urgency signal (Cisek et al., 2009). Unique to our results is that this tradeoff was predicted by the specific strategy (Evidence-by-Costs) individuals used to arbitrate between sampling and stopping.

We also found that increasing the financial stakes for a trial did not encourage adoption of a strategy that relies on optimal value estimation nor did it increase the likelihood of optimal behavior (i.e., participants did not move closer to the stopping decisions that would be made by an Ideal Observer (Achtziger et al., 2015)). Instead, we saw that participants were still best fit by the Evidence-by-Costs strategy. Behaviorally, participants in the high-reward condition over-sampled, effectively raising the stopping threshold without changing the underlying computations; both drawing more samples

and waiting until they achieved a larger difference between categories. This stands in contrast with prior work that suggested increasing monetary stakes can increase one's motivational state, thereby encouraging the use of more optimal but expensive strategies (Bennett et al., 2019; Manohar et al., 2015) .

Why might increased monetary stakes encourage behavioral adjustments but not push participants towards optimal behavior? One possible explanation is that conditions of equal difficulty but with higher monetary stakes could increase the effort put into the trial because of lower opportunity costs (Otto & Daw, 2019; Shenhav et al., 2013), but still not warrant the adoption of a completely new and expensive strategy. This is likely the case given that our task had intermixed high and low reward trials, such that participants would have to continually switch between strategies, which introduces cognitive costs (Luwel et al., 2009). Alternatively, our task was more difficult than prior information sampling paradigms because participants had to remember their previous samples – and thus the added memory demands might have deterred adoption of a more resource-intensive strategy. Although prior work has suggested that working memory capacity can impact the amount and use of information during information acquisition (Rakow et al., 2010), we did not measure the working memory capacities of our participants. Future research will need to explore how working memory demands and the cost of switching information sampling strategies shape stopping policies.

This study raises an important set of questions regarding how individuals determine their idiosyncratic thresholds. Evidence-by-Costs sampling provides two unique informational components that contribute to stopping – and individual model fits revealed variability in how participants weighed information about evidence and costs. Specifically, variability in the weighting of evidence could reflect varying levels of

confidence required for stopping across individuals, as seen in other work (Hausmann-Thürig & Läge, 2008). In line with prior work (Hauser et al., 2017, 2018; Juni et al., 2016; Otto & Daw, 2019; Petitet et al., 2021), we also see individual differences in sensitivity to accruing costs. Future work towards encouraging more optimal behavior can leverage our approach by specifically targeting informational components that most contribute to a person's sub-optimal sampling behavior. For example, an individual who consistently over-samples might do so because they are more sensitive to accuracy (a higher threshold starting point) or because they are less sensitive to accruing costs (a shallower trade-off slope). The ability to arbitrate between potential sources of error could provide a more targeted prescription to ameliorating the cause of over-sampling.

Moreover, our findings emphasize key directions for understanding sampling strategies themselves. First, additional research should identify and delineate strategies that do not completely conform to either heuristics or optimal behavior. A recent study (Korn & Bach, 2018), demonstrated the use of both heuristic and optimal strategies (but not a combination of the two) across a foraging task, providing insights into factors that shape strategy selection; for example, higher levels of experienced uncertainty may push sampling toward optimality. Similarly, in our current study, the Evidence-by-Costs strategy did not specifically integrate components from the optimal strategy but was sensitive to the same information sources. Cataloging a more complete space of sampling strategies will advance our understanding of how humans select what information to attend to and how that information is transformed into potential actions.

Future research should also explore how people determine what strategies to implement in different contexts. During information sampling, individuals not only decide how to balance sampling costs with accuracy but also contend with balancing the costs

and benefits of exerting control (Shenhav et al., 2013). Previous information sampling accounts have examined the impact of contexts such as changes in task difficulty (Coenen & Gureckis, 2016; Malhotra et al., 2017) and changes in sampling costs (Hauser et al., 2018; Juni et al., 2016) in altering sampling behavior but have not specifically examined if these contexts changed the underlying strategy. We examined the context of varying reward stakes on information sampling and found that while individuals maintained the same underlying strategy between both contexts, reward increased the overall information that people gathered. Our results differ from that observed in a reinforcement learning task, where reward stakes resulted in a switch to a more intensive but optimal strategy (Bennett et al., 2019). Additional investigation is needed into how individuals use context to evaluate when to switch information sampling strategies and when to adapt an existing strategy. Prior work has indicated that humans can learn to use context to determine strategy selection (Lieder & Griffiths, 2017; Payne et al., 1988; Rieskamp & Otto, 2006); for example, experience with a problem leads to adoption of more heuristic strategies and can even direct selection amongst different heuristics (Rieskamp & Otto, 2006). Although we did not find any changes in strategy in the current study, our task involved a longer sampling process and fewer sampling episodes, thus making it harder for participants to explore a variety of strategies. Future work will need to investigate how much experience individuals need in order to use contextual factors to inform both the selection and implementation of information sampling strategies.

Information sampling is a complex but ubiquitous challenge for decision makers. In the present study, we show that humans confront this challenge by adopting a strategy that balances the efficiency of heuristics but with increased flexibility.

Specifically, our results demonstrate that unlike optimal strategies, humans are performing simple accumulations instead of computing expected values, and unlike heuristic strategies, humans are dynamically integrating information instead of using rigid rules. Future work expanding how humans build such flexible strategies and how individual differences determine the relative weighting of different elements of those strategies (e.g., reward sensitivity) will provide further insight into the mechanisms by which bounded rationality guides decision-making processes.

3. Spoiler Alert! Curiosity can prioritize the information gathering process over the outcome

The content from this chapter is verbatim from the below preprint and has only been reformatted for this dissertation: Hsiung, A., Poh, J.-H., Huettel, S. A., & Adcock, R. A. (2022). *Spoiler Alert! Curiosity encourages patience and joy in the presence of uncertainty*. <https://doi.org/10.31234/osf.io/x5hgc>

3.1 Introduction

Curiosity, the desire to know, is a fundamental driver of human behavior. Curiosity encourages us to explore the unknown. It guides our behavior toward reducing the uncertainty in our environment (D. Berlyne, 08/1954; Dubey & Griffiths, 2020a; Gottlieb & Oudeyer, 2018; Loewenstein, 1994; Sharot & Sunstein, 2020; van Lieshout et al., 2020) and towards the pursuit of new knowledge (Gruber & Ranganath, 2019; Murayama et al., 2019; Noordewier & van Dijk, 2020). Across a variety of studies that have examined the subjective state of curiosity, a clear pattern emerges; when the alternative is to forgo information for the immediate future, curiosity directs choices to obtain information (Blanchard et al., 2015; Jach et al., 2021; Kang et al., 2009; Lau et al., 2020; Marvin & Shohamy, 2016; van Lieshout et al., 2018). Curiosity predicts the choice to obtain information even when it is costly, implying a value outside of its instrumental or external reward value (Bennett et al., 2016; Blanchard et al., 2015; Brydevall et al., 2018; FitzGibbon et al., 2021; Goh et al., 2021; Hsee & Ruan, 2016; Kang et al., 2009; Kobayashi & Hsu, 2019; Lau et al., 2020; Rodriguez Cabrero et al., 2019). Such experimental work has advanced our understanding of when and why curiosity guides

choice but has also encouraged a one-dimensional view of the relationship between curiosity and information seeking choices.

Accounts of curiosity as directing choice towards information that resolves states of uncertainty (e.g., seeking out outcomes or answers) stand in contrast with real-world examples where anticipation and uncertainty are intentionally prolonged. Consider, for example, preparing to see a highly anticipated movie. From a framework of curiosity as uncertainty reduction, we might predict that if you are curious about the movie, you would be motivated to actively seek out and learn about the outcome immediately (e.g., by reading summaries of its plot) (Blanchard et al., 2015; Kang et al., 2009; Loewenstein, 1994; Marvin et al., 2020; Rodriguez Cabrero et al., 2019; van Lieshout et al., 2020). Contrary to this prediction, people often prefer extended uncertainty and suspense (Ely et al., 2015; Li et al., 2021), such as in narratives (Brockbank et al., 2022; Gopnik, 2000; E. G. Liquin & Lombrozo, 2020; Schwartenbeck et al., 2019), close games in sports (Su-lin et al., 1997), and video games (Klimmt et al., 2009; Lomas et al., 2017). People even specifically avoid early resolution (Maxwell 2021; Johnson and Rosenbaum 2015) by adopting self-control mechanisms to decrease the likelihood of information exposure (Boyd-Graber et al. 2013; Bao et al. 2021). The desire to avoid information that resolves uncertainty prematurely is so ubiquitous that we explicitly label such information via “spoiler alerts”.

What can account for the discrepancy between the experimental findings on curiosity-induced information seeking and these real-world behaviors? A critical difference lies in how information seeking is assessed in behavioral experiments. Often, the central focus has been on the information itself, without consideration for the process of how such information is collected. This produces paradigms that treat the transition

from curiosity induction to resolution as a step-function, where the jump from uncertainty to certainty happens all at once (i.e., the reveal of a trivia answer or the outcome of an anticipated gamble) (Bennett et al., 2016; Dubey & Griffiths, 2020a; Jepma et al., 2012; Kang et al., 2009; Marvin & Shohamy, 2016; Van de Cruys et al., 2021; van Lieshout et al., 2018). Even under circumstances where the delay between the induction and relief of curiosity is manipulated, the outcome of information is still delivered altogether (Iigaya et al., 2016; Lang et al., 2021; Noordewier & van Dijk, 2017).

Yet in our daily lives, information is often accumulated over time (e.g., reading a book), and the transition from curiosity to relief is represented more as a gradually evolving function. Expanding our perspective on how we test information-seeking opens investigations into understanding how curiosity functions in more dynamic information gathering experiences. Indeed, work has begun to explore how curiosity impacts the organization of information seeking over longer timescales, highlighting differences in trait-level curiosity in constructing distinct informational networks (Jach et al., 2021; Lydon-Staley et al., 2021) and the role of agency in discovering answers (Metcalf et al., 2021). However, this work still leaves open the question of whether curiosity prioritizes seeking advanced information in contexts where information is continuously unfolding and one's informational state is dynamically determined.

Moreover, expanding the breadth of information seeking examined under curiosity has the potential to uncover how different affective states during the experience of curiosity can direct distinct preferences for how information is revealed. Prior work has documented different affective experiences of curiosity both at the trait-level (Kashdan et al., 2018; Litman, 2010) and at the state-level (Jach et al., 2021; Loewenstein, 1994; Noordewier & van Dijk, 2020; Ryakhovskaya et al., 2021; van Lieshout et al., 2021),

illustrating that curiosity can take on exploratory and playful forms as well as urgent, aversive forms. Understanding whether curiosity is characterized by more positive or negative affective states during continuous information gathering can provide insight into why curiosity might favor different paths towards reaching the outcome.

We developed a novel dynamic perceptual task in which participants watched videos of line drawings as they slowly resolved into objects. Across video watching, participants could periodically choose to stop the video early or keep watching the drawing unfold. To investigate curiosity-induced preferences for how outcomes were revealed, we manipulated the consequence of stopping a video early between participants. In one group, stopping a video meant receiving no further information about the drawing (*Forgo Resolution*), akin to previous studies examining the choice between receiving information or not (Kang et al., 2009; Marvin & Shohamy, 2016; van Lieshout et al., 2020). This condition was thus directly comparable to tasks in which resolution was a discrete event, to determine whether curiosity experienced over continuous information gathering would similarly value the outcomes of receiving or not receiving information. In our other group, stopping early meant participants could immediately jump to the end of a video, revealing the full drawing without having to wait (*Jump-to Resolution*). In this condition, we could test whether curiosity was sensitive to how an individual arrived at the same outcome of receiving information, here a completed drawing. Specifically, would high states of curiosity encourage outcome-focused information seeking, jumping to the end to avoid waiting and receive immediate resolution? Or would curiosity instead promote process-focused information seeking, delaying resolution in order to watch the video unfold? In both conditions, we also

explored what affective states mapped onto curiosity as well as monitored how confident participants were about the identity of the drawings.

We compared our Forgo Resolution and Jump-to Resolution conditions across two experiments. In Study 1, the timing of the choices to continue watching or stop early was externally imposed (Predetermined Timing) and in Study 2, participants had agency over exactly when to make these choices (Free Timing). In both experiments, we showed that curiosity predicted the choice to continue watching videos when the alternative was to forgo further information, consistent with prior literature demonstrating a preference for information compared to no information (Forgo Resolution). Crucially, curiosity also predicted the choice to continue watching videos when the alternative was to accelerate resolution (Jump-to Resolution), indicating that under curiosity participants preferred delaying resolution to watch the information unfold. In addition, we found that measures of confidence and certainty predicted stopping videos early, indicating that uncertainty about the identity of the drawing was a driving force for engaging with these videos, rather than an alternate reward value such as aesthetic experience. The subjective experience of curiosity was closely associated with the affective state of enjoyment and with tension but not with frustration. Lastly, when comparing Study 1 and 2, we found that when participants had agency over the timing of their choices, they were more likely to wait when they were uncertain and more likely to stop early once they identified the drawing, emphasizing the value of reaching resolution during continuous information gathering. These results support the conclusions that curiosity can promote choices to prolong uncertainty, that this type of curiosity is characterized by positive affect, and that prolonging uncertainty to experience the natural unfolding of information can drive non-instrumental information seeking.

3.2 Methods

3.2.1 Participants.

Our study recruited a total of 2153 US adults (Study 1: N = 1,076, Study 2: N = 1077) through the online labor marketplace, Prolific, in September of 2020. Participants completed the entire study online. 109 participants were excluded in total, 27 due to technical errors with the task that impeded reading or matching data files with Prolific IDs, 4 due to completing the task twice, 14 participants had unidentifiable guesses (e.g. “ahexgfe”), 17 participants could not identify more than 80% of the drawings (final answer = “I don’t know” or “none”), and 48 participants due to not providing a Mid-Video Rating on 90% of trials (Study 2 only). This resulted in a final sample of 2,043 participants (Study 1: N = 1,033, Study 2: N = 1,010). Within our sample, 48% of our participants identified as women, 48% identified as men, and 4% identified as gender non-conforming or preferred not to say. The mean age of our participants was 32 years (SD, 11.4 years; range, 19 - 76 years). Our racial demographic breakdown showed that 69% of our participants identified as White, 13% as Asian, 8% as Black, and 10% as either mixed or other. Our study was estimated to take 35 minutes (or 55 minutes for the two-part study that included memory, see SI Appendix C) and participants were paid a flat \$5.75 for completion of the study (or \$9.00 for the two-part study). All participants received informed consent under the guidelines of our university’s Institutional Review Board.

3.2.2 Stimuli and Apparatus.

Single line drawings were collected through Google Image searches. Each drawing consisted of a single item that could be named using a single word. Items came

from a variety of categories, including objects (e.g., guitar, headphones), food (e.g., taco, pepper), persons (e.g., ballerina, golfer), and animals (e.g., tiger, eel). Video stimuli were then hand created and standardized through iMovie. Videos were grayscale and adjusted to fit the center of the screen. All videos were calibrated to last between 20 and 30 seconds ($M = 24$ seconds, $SD = 3$ seconds) (SI Appendix D: for examples). To program the task, we used the JavaScript framework Vue.js. We hosted the experiment on a private virtual web server using psiTurk (Gureckis et al., 2016), which participants could access through their browser.

3.2.3 Task

After informed consent and prior to starting the task, participants were randomly assigned to one of two Stop-Outcome conditions: Forgo Resolution and Jump-to Resolution that determined the consequence of choosing to end a video early. In the Forgo Resolution condition (Study 1: $N = 501$, Study 2: $N = 485$), participants were told that if they stopped the trial, they would receive no further information about the drawing. In the Jump-to Resolution condition (Study 1: $N = 532$, Study 2: $N = 525$), participants were told that stopping the video early would jump them to the end of the video where they would see the completed drawing immediately. Participants were told they would watch animated videos of single-line drawings unfold over time. They were informed of the expected length of each video as well as how many videos they were to view (minimum of 25 videos). Participants were told that during video watching they would be expected to answer questions about the drawing as it was unfolding (Mid-Video Decisions). Answering all questions was required to advance through the task. Across all versions, participants were told that their choices and guesses would not impact their compensation. Changing and updating guesses were described as an expected part of

the video watching process across all conditions and participants would not be penalized for wrong guesses (SI Appendix E: task instructions).

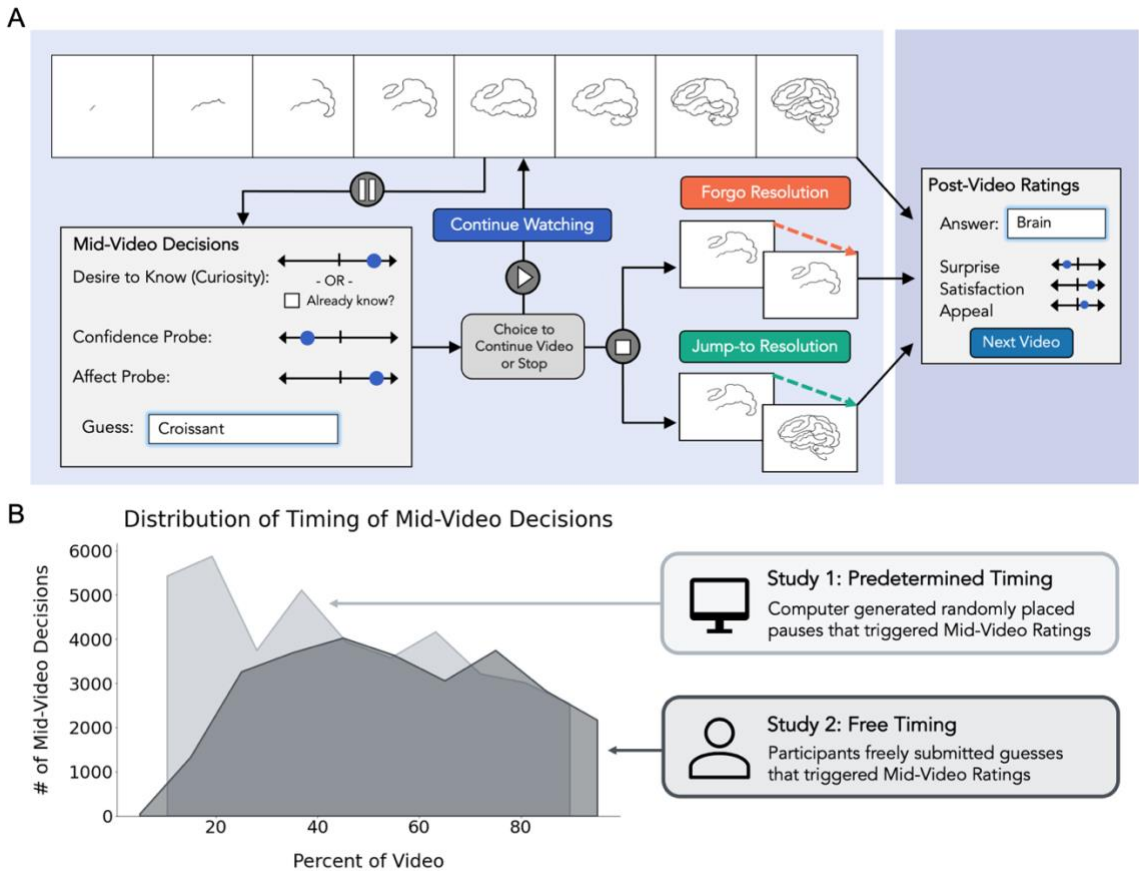


Figure 7: Task design and Study 1 and 2 manipulation

Each trial started with a video box in the center of the webpage (width: 620 pixels, height: 400 pixels). Videos played automatically at the start of each trial. In Study 1, participants were informed that there would be randomly placed pauses, between 1-3 per video, where they would be asked to answer the Mid-Video Decisions (Predetermined Timing). In Study 2, throughout video play, a textbox was present underneath the video. Participants were told to guess as soon as they thought of an

identity for the drawing and to update that identity at least once per video (Free Timing). Once participants submitted a guess, they answered the other Mid-Video Decisions.

At Mid-Video Decisions, participants were asked to report on their curiosity (“How much do you want to know the outcome of the drawing?”), one of two randomly assigned questions gauging their confidence about the drawing’s identity (“how confident right now”, “how much more time needed to know”), one of three randomly assigned affective states (enjoyment, tension, or frustration) and what they thought the identity of the drawing was given their current information (Guess). Curiosity, confidence, and affect were reported on sliding scales that ranged from 1 - 100 (for the time-needed confidence question the scale was 0 - 30 seconds). For our curiosity question only, if participants were certain they knew what the identity of the drawing was, instead of reporting on the scale, they could check a box labeled “Already Know”, providing a measure of subjective certainty. Each question was displayed on top of the sliding scale for that question. Only the lower and upper bounds of the scale were labeled (e.g., Not at all/Very much). For guesses, we instructed participants to use the most descriptive and common single English word (e.g., “For example, if you see an image of a hamster, write ‘hamster’ instead of ‘animal’ or ‘Winter White Dwarf hamster.’”). Participants were required to submit a guess that was greater than 1 character.

After providing responses to all 4 questions, participants then chose whether they wanted to continue watching the video or stop the video and move to the next trial. If a participant chose to continue watching a video on their first Mid-Video rating, on the next Mid-Video rating, they again were presented with the option to continue or stop. Across all Mid-Video Decisions, participants chose to continue watching a video 73% of the time. When participants did choose to stop, it was on average, 60% or 14.8 seconds

through a video. At the end of a trial, regardless of whether participants watched the video through or ended the trial early, participants were presented with the Post-Video ratings. Participants were asked to report on their video satisfaction, their surprise in the video outcome, the aesthetic appeal of the drawing, and their anticipated curiosity for the upcoming trial. Similar to the Mid-Video Decisions, all Post-Video ratings were reported on sliding scales that ranged from 1 – 100 with each question displayed on top. Only the lower and upper bounds of the scale were labeled (e.g., Not at all/Very much). Participants also provided a final answer as to what they thought the identity of the drawing was in a text box underneath the video. Feedback about the identity of the drawing was not provided to participants at any time. Participants thereby relied on their own sense of certainty to name the drawings. Each trial ended with the option to move to the next trial or replay the current video. During replay, participants simply watched the video, no pauses or guessing were permitted. Replay choices were made on 1.6% of trials and will not be discussed further.

Participants received 2 alerts to mark their progress through the task, one at the midway point and one with 5 videos left. At the end of the required 25 videos, all participants were given the option to terminate the video guessing portion of the study or watch up to 10 extra videos before moving to the questionnaires. Participants were informed that watching extra videos would not change their compensation. Each additional video trial functioned identically to the main 25 video trials except that after answering the Mid-Video Decisions or the Post-Video ratings, participants now also had the option to terminate the video guessing portion of the task. 32% of our participants chose to watch at least one additional video ($M = 2.72$ additional videos, $SD = 2.90$ videos, Range: 1 - 10 videos).

3.2.4 Data Analysis

Before conducting all statistical analyses, predictors were either scaled or factorized. For this manuscript, we collapsed the two confidence questions into a single subjective confidence marker. We also standardized all participant guesses and final answers (see SI Appendix F). In short, final answers were scored as correct if the answer matched the name of the drawing, if it matched a name in a list of alternative identification, or if the answer was provided after the participant had seen the completed drawing (either from watching the video all of the way through in the Forgo Resolution condition or anytime in the Jump-to Resolution condition). Across all final answers provided, accuracy was 98%. Guess Accuracy was then determined by comparing each participant's mid-video guesses to their final answer. Guess accuracy across all Mid-Video Decisions was 47.9%.

Analyses were conducted with R v4.0.3 and RStudio v1.3.1093. We created linear and logistic mixed effects regression models with the lme4 package (Bates et al., 2014) and obtained p-values with the lmerTest package (Kuznetsova & Brockhoff, 2017). All mixed effects models included random effects with both participant and video as random intercepts. All models included all possible two-way interactions with the predictor for Stop-Outcome since this manipulation was the main focus for these studies. For all statistical analyses, we report standardized β values (and 95% confidence intervals (CIs) around the slope estimate). All models also included covariates for each task version a participant was in. Figures were produced in Python with Seaborn and Matplotlib.

3.3 Results

3.3.1 Study 1

To examine how curiosity directs information seeking when information is continuously unfolding, we asked 1033 participants to watch animated videos of single-line drawings unfold over time (*See Methods & Materials for information on participant exclusions and demographics*). Participants were informed that each video would be between 20 – 30 seconds and they would watch a minimum of 25 videos. Participants were informed that there would be randomly placed pauses throughout the video where they would be asked to answer Mid-Video Decisions (Fig. 1A). The timing of the pauses was determined by the computer and participants experienced between 1-3 per video (Fig. 1B, Predetermined Timing). Participants were asked to report on their curiosity, their confidence about the drawing's identity, and one of three randomly determined affective states (enjoyment, tension, frustration) assigned between-subjects on sliding scales (*See Methods and Materials*). They were also asked to provide a guess for what they thought the identity of the drawing was given their current informational state.

After providing responses to all 4 questions, participants then chose whether they wanted to continue watching the video or stop the video and move to the next trial. We manipulated the Stop-Outcome between participants, such that in one group (Forgo Resolution, N = 501), stopping the video ended information gathering and moved subjects to the Post-Video ratings (SI Appendix A: Post-Video results). In the other group (Jump-to Resolution, N = 532), stopping the video jumped participants to the end of the video where they saw the final drawing fully revealed before moving to the Post-Video ratings. If a participant chose to continue watching a video on their first Mid-Video Decision, on the next Mid-Video Decision, they were again presented with the option to

continue or stop. Participants were told that their choices and guesses would not impact their compensation. Changing and updating guesses were described as an expected part of the video watching process and participants would not be penalized for wrong guesses.

3.3.1.1 Curiosity and Information Seeking.

We first investigated whether our continuous information gathering task would replicate curiosity-guided choice behavior from previous discrete information gathering tasks. We fit a mixed-effects logistic regression with both participant and video as random intercepts to determine if high states of curiosity would predict continuing to watch a video when the alternative was to forgo the rest of the video and move to the next trial (Forgo-Resolution). Within the model, we included participant's curiosity ratings, confidence ratings, guess accuracy (determined by comparing each participant's Mid-Video guesses to their final guess provided in the Post-Video ratings), and the percent time elapse since the start of the video. Trial number was included as a regressor of no interest. We hypothesized that, consistent with prior research, high states of curiosity would predict continuing to watch a video since this was the only path towards receiving resolution (Fig. 2B). Indeed, we found that high states of curiosity predicted continuing to watch a video when the alternative was to forgo collecting more information about the drawing ($\beta = 0.65$, 95% CI [0.58, 0.72], $z = 17.46$, $p < 0.001$) (Fig. 2C). This suggests that curiosity encourages uncertainty resolution during continuous information gathering, similar to that observed during discrete information seeking.

We next sought to investigate the role of curiosity in predicting information seeking when the choice options were between continuing to watch the video or jumping to the end and resolve uncertainty immediately (Jump-to Resolution). For this analysis,

we hypothesized that high states of curiosity would either predict a decreased probability of continuing to watch a video, demonstrating a sensitivity to the costs of delaying receiving the outcome and thus a preference for the fastest path towards resolution (outcome-focused), or that curiosity would predict an increased probability of continuing to watch a video, signaling that experiencing the process of information unfolding confers value (process-focused) (Fig. 2B). We fit an identical mixed-effects logistic regression as our Forgo Resolution condition and found that even when participants had the option to receive immediate resolution, higher curiosity increased the likelihood of continuing with a video ($\beta = 0.45$, 95% CI [0.39, 0.52], $z = 13.93$, $p < 0.001$) (Fig. 2C). These results suggest that people do not always seek immediate resolution for their curiosity, and that curiosity can instead encourage prioritizing the process of information gathering over the outcome.

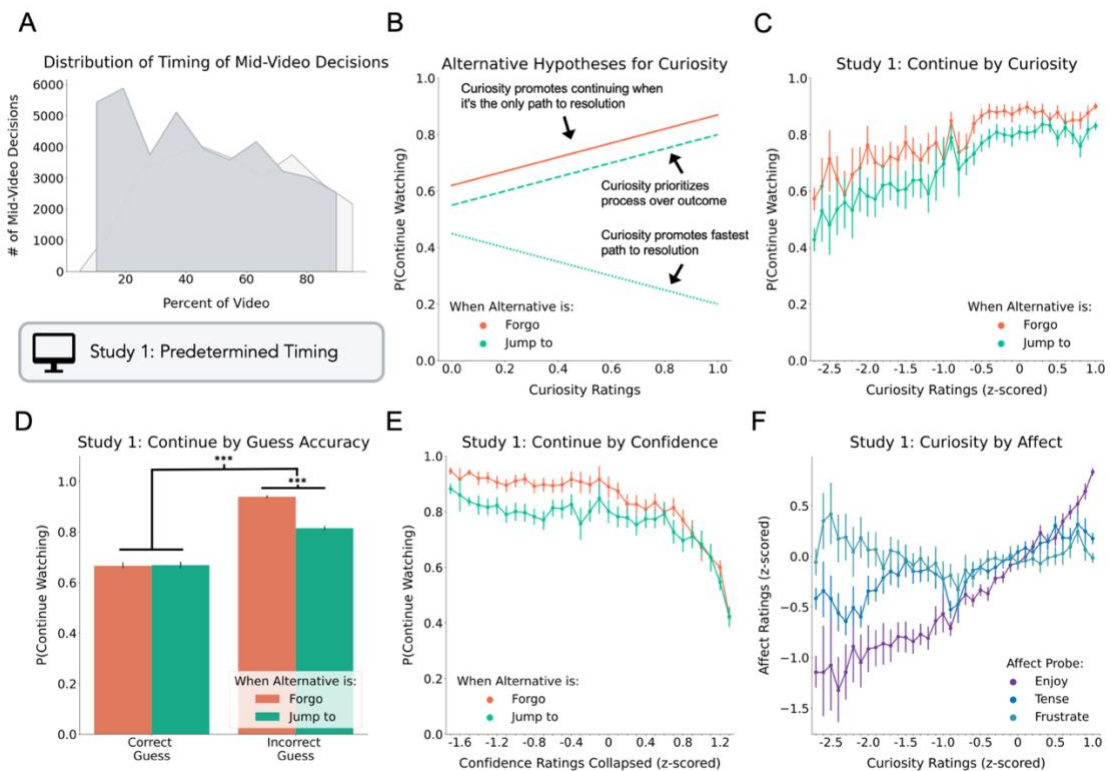


Figure 8: Results from Study 1. Curiosity predicts continuation across Forgo- and Jump-to Resolution conditions.

3.3.1.2 Uncertainty and Information Seeking

In addition to curiosity, across both the Forgo Resolution and Jump-to Resolution conditions, we found that correctly guessing the identity of the drawing (Forgo: $\beta = -2.20$, 95% CI [-2.37, -2.04], $z = -26.32$, $p < 0.001$; Jump-to: $\beta = -1.02$, 95% CI [-1.17, -0.88], $z = -14.07$, $p < 0.001$) and high states of confidence (Forgo: $\beta = -0.77$, 95% CI [-0.86, -0.68], $z = -17.05$, $p < 0.001$; Jump-to: $\beta = -0.80$, 95% CI [-0.88, -0.72], $z = -19.92$, $p < 0.001$) predicted stopping a video early, indicating that as participants moved towards resolution, they were less likely to continue information gathering (Fig. 2D, 2E). In the Forgo Resolution condition, the percent of time watching the video was not a significant

predictor of continuation ($\beta = -0.06$ 95% CI [-0.13, 0.02], $z = -1.43$, $p = 0.153$). In contrast, in the Jump-to Resolution condition, the longer a participant watched a video, the more likely they were to end that video early ($\beta = -0.15$, 95% CI [-0.21, -0.08], $z = -4.93$, $p < 0.001$). Taken together, as participants became more confident in their assessment of the drawing's identity, they were more likely to stop video watching, suggesting that reaching resolution decreased the value of collecting more information.

Lastly, to directly compare the two Stop-Outcome conditions, we adapted our mixed-effects logistic regression model to include participants from both Stop-Outcome conditions and included all two-way interactions between Stop-Outcome condition and our original predictors. Overall, we found that participants in the Jump-to Resolution condition had significantly lower rates of continuing than participants in the Forgo Resolution condition ($\beta = -1.59$, 95% CI [-2.00, -1.19], $z = -7.80$, $p < 0.001$), suggesting that having the option to resolve uncertainty immediately drove more choices to stop watching a video compared to the option to forgo receiving further information. Furthermore, we found that the Stop-Outcome manipulation significantly interacted with curiosity ($\beta = -0.21$, 95% CI [-0.32, -0.12], $z = 4.40$, $p < 0.001$) and guess accuracy ($\beta = 1.17$, 95% CI [0.96, 1.38], $z = 10.84$, $p < 0.001$). Post-hoc analyses indicated that participants in the Jump-to Resolution condition were less likely to continue watching a video when they reported an incorrect guess (OR = 3.23, $z = 10.84$, $p < 0.001$) compared to participants in the Forgo Resolution condition (Fig. 2D). Additionally, the slope for curiosity ratings was steeper in the Forgo Resolution condition ($z = 4.40$, $p < 0.001$), indicating that curiosity was a stronger predictor of choice in the Forgo Resolution condition compared to the Jump-to Resolution condition (Fig. 2C). To summarize, participants were more likely to continue watching videos when they were

uncertain compared to when they were certain. When given the option to receive immediate resolution (Jump-to Resolution), participants showed a decreased desire to continue watching videos even when uncertainty was high but nevertheless, curiosity remained positively associated with the willingness to continue watching videos.

3.3.1.3 Curiosity and Affect

Given that curiosity predicted delaying resolution over resolving uncertainty immediately, we hypothesized that the subjective experience of curiosity in this task was unlikely to be aversive and more likely to be positive. To test this, we collapsed data from both Stop-Outcome conditions and fit three mixed-effects linear regression models with both participant and video as random intercepts to test what affective states were associated with curiosity. Within the models, we included confidence ratings and guess accuracy, since these have previously been linked to the state of curiosity (Dubey & Griffiths, 2020a; Kang et al., 2009; Marvin & Shohamy, 2016). We also included the percent time elapse since the start of the video to be consistent with our choice model above and the trial number, again, as a regressor of no interest. We found that curiosity was positively related to the affective state of enjoyment ($\beta = 11.62$, 95% CI [11.14, 12.10], $t = 47.00$, $p < 0.001$) and tension ($\beta = 5.88$, 95% CI [5.33, 6.43], $t = 21.04$, $p < 0.001$) but was not related to frustration ($\beta = -0.14$, 95% CI [-0.66, 0.37], $t = -0.547$, $p = 0.584$) (Fig. 2F). Taken together, our results demonstrate that curiosity, as experienced in this task, was not an aversive state but instead a pleasant state of anticipation.

In addition to affect, consistent across all three mixed-effects linear regressions, curiosity was associated with low states of confidence (Enjoy: $\beta = -6.04$, 95% CI [-6.51, -5.57], $t = -25.22$, $p < 0.001$; Tense; $\beta = -4.47$, 95% CI [-5.00, -3.94], $t = -16.42$, $p < 0.001$; Frustrate: $\beta = -5.76$, 95% CI [-6.31, -5.21], $t = -20.64$, $p < 0.001$) and was lower

when guessing was correct (Enjoy: $\beta = -2.73$, 95% CI [-3.60, -1.86], $t = -6.14$, $p < 0.001$; Tense; $\beta = -2.34$, 95% CI [-3.33, -1.35], $t = -4.63$, $p < 0.001$; Frustrate: $\beta = -2.89$, 95% CI [-3.90, -1.87], $t = -5.57$, $p < 0.001$). Interestingly, we also found that across the Frustrate and Tension affective probe, curiosity increased the longer a participant stayed watching a video (Frustrate: $\beta = 1.70$, 95% CI [0.72, 1.62], $t = 5.10$, $p < 0.001$; Tense; $\beta = 0.45$, 95% CI [0.01, 0.89], $t = 2.03$, $p = 0.043$) but this relationship did not reach statistical significance in the Enjoy affective probe (Enjoy; $\beta = 0.38$, 95% CI [-0.01, 0.76], $t = 1.78$, $p = 0.06$). These findings indicate that high states of curiosity were associated with states of uncertainty, grew across video watching but as participants became more confident in the identity of the drawing, curiosity declined.

3.3.1.4 Study 1 Conclusions

In Study 1, we showed that under continuous information gathering, curiosity predicts choices to seek information when the alternative is to forgo information, extending previous findings using discrete information seeking (FitzGibbon et al., 2021; Kang et al., 2009; Lau et al., 2020; Marvin & Shohamy, 2016; Noordewier & van Dijk, 2017; van Lieshout et al., 2018). Critically, we find that curiosity increased the likelihood that participants continued the video even when they had the choice for immediate resolution, demonstrating that curiosity can encourage the intentional prolonging of uncertainty (Metcalfe et al., 2021). In contrast to curiosity, high subjective states of confidence and correctly identifying the drawing in the video led to early termination of video watching across both the Forgo Resolution and Jump-to Resolution conditions. Thus, the experience of reaching resolution promoted the cessation of information gathering.

Furthermore, the subjective experience of curiosity in this task was related to an affective state marked by enjoyment and tension but not frustration. We found that curiosity was also inversely related to confidence and guess accuracy, supporting prior work (Dubey & Griffiths, 2020a; Vogl, Pekrun, Murayama, Loderer, et al., 2019), as well as a positive association between curiosity and the time elapsed watching a video, potentially signaling a ramping of curiosity as the video progresses. Overall, seeking information when it continuously unfolds creates an opportunity where curiosity can support a preference for delayed resolution, often until the identity of the drawing is realized. However, in Study 1, participants were probed at set intervals beyond their control, and it is thus possible that this approach could have inflated participant's willingness to continue watching the video (particularly for earlier probes). To address this, we conducted a second experiment where participants had agency in determining when they were able to submit their guesses, and to stop the video.

3.3.2. Study 2

Here, we investigated whether curiosity would still predict prolonging uncertainty if participants could choose when to submit Mid-Video Decisions during their information gathering experience. We asked 1010 participants to complete the same video-watching task as in Study 1 (*See Methods & Materials for information on participant exclusions and demographics*), however this time instead of experiencing randomly placed pauses, a textbox was presented underneath the playing video, and participants could freely submit guesses and Mid-Video Decisions (Fig. 1B, Free Timing). When participants

began to type into the textbox, the video would pause until either a guess was submitted or the characters in the textbox were removed. Submitting a guess prompted the appearance of the same 3 ratings scales as in Study 1 (curiosity, confidence, and affect). After answering the Mid-Video ratings, participants could then choose to continue watching the video or move to the end of the trial. To match Study 1, participants were similarly divided into two Stop-Outcome conditions: *Forgo Resolution*, N = 485 and *Jump-to Resolution*, N = 525. We encouraged participants to guess often throughout video play. Study 2 was otherwise identical to Study 1 (SI Appendix B: Post-Video rating results).

3.3.2.1 Choice Behavior Results

Similar to Study 1, we first tested whether curiosity would predict continued video watching when the alternative was to forgo the rest of the video and move to the next trial (Forgo-Resolution). Using an identical mixed-effects logistic regression as in Study 1, we replicated the results and found that high states of curiosity predicted continuing to watch a video when the alternative was to forgo collecting more information about the drawing ($\beta = 0.60$, 95% CI [0.49, 0.69], $z = 11.80$, $p < 0.001$) (Fig. 3C). We next sought to investigate whether curiosity would still predict continued information seeking over choices to resolve uncertainty immediately (Jump-to Resolution) when participants could control how much information to collect. With the same mixed-effects logistic regression, we replicated the effect that even when participants had the option to receive immediate resolution, higher curiosity increased the likelihood of continuing to watch a video ($\beta = 0.27$, 95% CI [0.18, 0.35], $z = 6.14$, $p < 0.001$) (Fig. 3C). These results suggest that even when people have control over how much time to collect information, curiosity still

encourages choices to continue collecting information and prioritizes the process of experiencing unfolding information over simply obtaining the outcome.

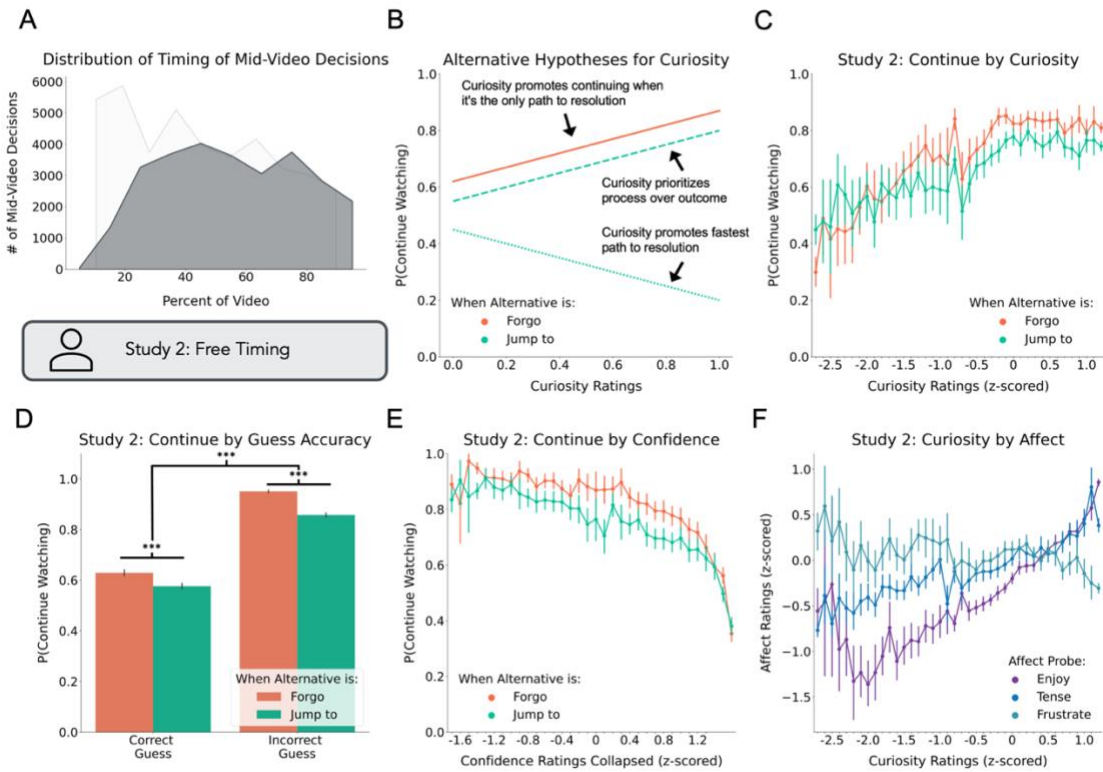


Figure 9: Results from Study 2. Curiosity predicted continuation even when participants were given agency over the timing of their choices.

Following the same analysis as Study 1, we found that in addition to curiosity, across both the Forgo Resolution and Jump-to Resolution conditions, correctly guessing the identity of the drawing (Forgo: $\beta = -2.20$, 95% CI [-2.37, -2.04], $z = -26.32$, $p < 0.001$; Jump-to: $\beta = -1.02$, 95% CI [-1.17, -0.88], $z = -14.07$, $p < 0.001$) and high states of confidence (Forgo: $\beta = -0.77$, 95% CI [-0.86, -0.68], $z = -17.05$, $p < 0.001$; Jump-to: $\beta = -0.80$, 95% CI [-0.88, -0.72], $z = -19.92$, $p < 0.001$) both predicted stopping a video early (Fig. 3D, 3E). In contrast to Study 1, across both Stop-Outcome conditions the more

time spent watching the video at the time of Mid-Video submission the less likely participants were to continue with the video (Forgo: $\beta = -0.62$, 95% CI [-0.73, -0.51], $z = -11.23$, $p < 0.001$; Jump-to: $\beta = -0.74$, 95% CI [-0.83, -0.64], $z = -15.20$, $p < 0.001$). To summarize, in Study 2 we replicate the results that as participants became more confident in their assessment of the drawing's identity, they were more likely to stop video watching. We also found that giving participants control over their time spent collecting information led to increased likelihood of stopping as the video progressed, consistent with the notion that participants were more likely to stop once they had acquired sufficient information.

Next, we directly compared the two Stop-Outcome conditions. We again found that participants in the Jump-to Resolution condition had significantly lower rates of continuing than participants in the Forgo Resolution condition ($\beta = -1.51$, 95% CI [-2.02, -0.99], $z = -5.78$, $p < 0.001$). Consistent with Study 1, we found that the Stop-Outcome manipulation significantly interacted with curiosity ($\beta = -0.32$, 95% CI [-0.45, -0.19], $z = 4.93$, $p < 0.001$) and guess accuracy ($\beta = 1.02$, 95% CI [0.75, 1.29], $z = 7.46$, $p < 0.001$). Post-hoc analyses revealed that participants in the Jump-to Resolution condition were less likely to continue watching a video when they reported an incorrect guess (OR = 2.78, $z = 7.46$, $p < 0.001$) (Fig. 3D) than participants in the Forgo Resolution condition. The slope for curiosity ratings was steeper in the Forgo Resolution condition ($z = 4.93$, $p < 0.001$), similar to Study 1. In Study 2, we also found that the Stop-Outcome condition significantly modulated the relationship between confidence ratings and continuation behavior such that the slope for confidence ratings was steeper in the Forgo compared to Jump-to Resolution condition ($z = 3.70$, $p = 0.0002$). Together, these results suggest that when participants had control over how long they spent gathering information, they

relied on similar markers as in Study 1 to guide their choice behavior but were more sensitive to the amount of time they spent collecting information within a trial.

3.3.2.2 Affect Results

We next aimed to characterize if the subjective state of curiosity was different in Study 2 compared to Study 1. To test this, we collapsed data from both Stop-Outcome conditions in Study 2 and fit three mixed-effects linear regression models with both participant and video as random intercepts to test what affective states were associated with curiosity. Curiosity was again, positively related to the affective state of enjoyment ($\beta = 8.20$, 95% CI [7.65, 8.75], $t = 29.61$, $p < 0.001$) and tension ($\beta = 4.35$, 95% CI [3.73, 4.98], $t = 13.70$, $p < 0.001$) and was not related to frustration ($\beta = 0.07$, 95% CI [-0.50, 0.65], $t = 0.248$, $p = 0.804$) (Fig. 3F). Our results match the findings from Study 1 and provide further support that here, curiosity was best characterized as a positive compared to negative state.

In line with Study 1, curiosity was also associated with low states of confidence (Enjoy: $\beta = -3.38$, 95% CI [-3.86, -2.91], $t = -14.02$, $p < 0.001$; Tense; $\beta = -3.17$, 95% CI [-3.73, -2.60], $t = -10.93$, $p < 0.001$; Frustrate: $\beta = -4.48$, 95% CI [-5.05, -3.90], $t = -15.29$, $p < 0.001$), and was lower when guessing was correct across the Enjoy and Tension affective probe (Enjoy: $\beta = -1.33$, 95% CI [-2.19, -0.46], $t = -3.00$, $p = 0.003$; Tense; $\beta = -1.43$, 95% CI [-2.48, -0.37], $t = -2.65$, $p = 0.008$) but not the Frustration affective probe (Frustrate: $\beta = 0.46$, 95% CI [-0.60, 1.52], $t = 0.85$, $p = 0.393$). In contrast to Study 1, we found that across the Enjoy and Tension affective probe, curiosity decreased the longer a participant stayed watching a video (Enjoy: $\beta = -1.06$, 95% CI [-1.48, -0.64], $t = -4.92$, $p < 0.001$; Tense; $\beta = -0.89$, 95% CI [-1.42, -0.36], $t = -3.30$, $p = 0.001$) but this relationship was not observed in the Frustrate affective probe (Frustrate; $\beta = 0.01$, 95% CI [-0.51,

0.53], $t = 0.04$, $p = 0.967$). Overall, we found that curiosity decreased as confidence and guess accuracy increased, but these results were not as consistent across all three affective probes. Additionally, having freedom over when to make information seeking choices changed the relationship between time elapsed watching the video and curiosity such that here, the longer a participant spent in a trial the lower their curiosity became.

3.3.2.3 Free vs. Predetermined Timing

We then directly compared Study 1 and Study 2 to assess how the addition of agency changed the information gathering experience. We first investigated differences in how long participants decided to engage in the information gathering process. Specifically, since participants had control over exactly when to stop seeking information, we were interested in testing if participants were indeed waiting to reach resolution before they stopped a video early. We first assessed across the two experiments whether there was an overall difference in guess accuracy, which provided a measure of whether a participant had reached resolution about the drawing's identity. Focusing on the last guess provided during video watching, which signified a participant's most updated idea about the identity of the drawing, we found the proportion of correct guesses in Study 2 was higher (more correct guesses reported) than in Study 1 ($\text{Chi-squared}(1, 40,754) = 2065.1$, $p < 0.001$, Study 1: 49.2% correct last guesses; Study 2: 70.6% correct last guesses) (Fig. 4A). Given that participants in Study 2 provided overall more correct last guesses, we explored whether participants in Study 2 were gathering more information prior to stopping compared to Study 1. Focusing only on trials where a participant stopped a video early, we found that in Study 2, participants watched more of a video before stopping ($F(1, 21,673) = 1841$, $p < 0.001$; Study 1: $M = 48.9\%$, $SD = 25.6\%$ video watched at stop; Study 2: $M = 65.1\%$, $SD = 21.7\%$ video

watched at stop) compared to in Study 1 (Fig. 4B). Next, we examined whether reaching resolution, as measured by a correct guess, led to different rates of stopping videos early across our two studies. Using a logistic regression that controlled for the timing of the guesses, we found a significant interaction between the Study and guess accuracy ($\beta = -0.55$, 95% CI [-0.63, -0.47], $z = -14.00$, $p < 0.001$), indicating that continuation rates were lower in Study 2 compared to Study 1 but only when there was a correct guess reported (OR = 1.40, $z = 16.91$, $p < 0.001$) (Fig. 4C). When an incorrect guess was reported, we found the opposite effect where continuation rates were higher in Study 2 compared to Study 1 (OR = 0.81, $z = -6.26$, $p < 0.001$). This interaction provides support towards our hypothesis that when participants had agency over their information seeking experience, they waited until they experienced resolution as information actively unfolded.

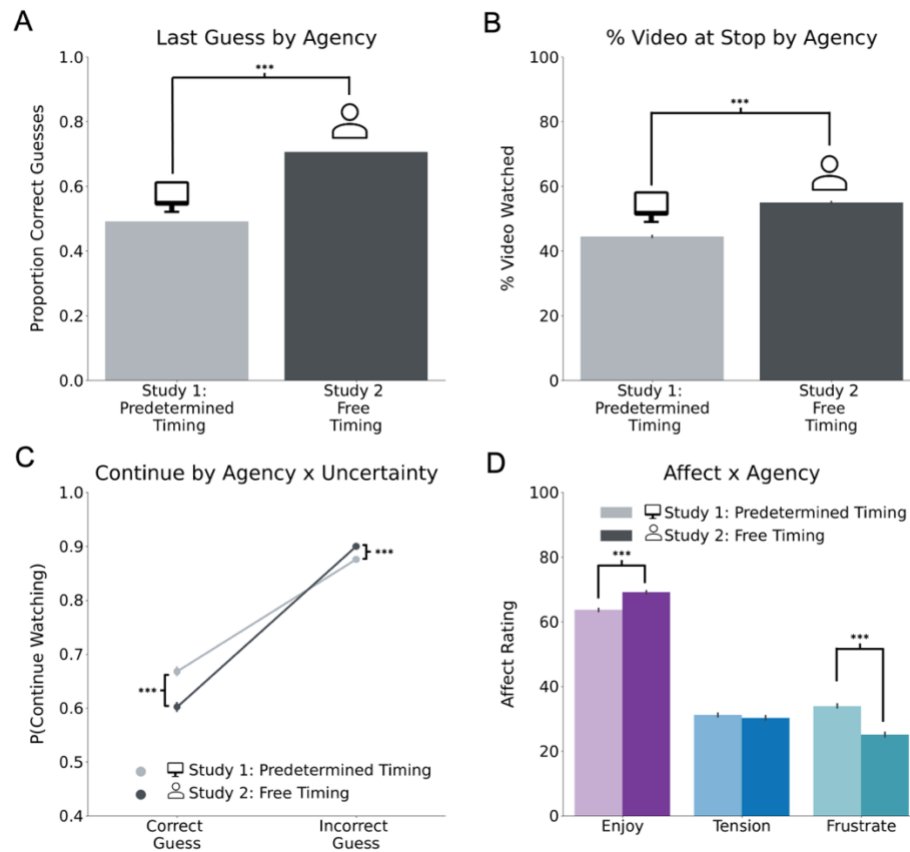


Figure 10: Comparing information gathering behavior between Study 1 and Study 2.

Lastly, using linear regression with Study as our predictor, we compared the overall affective experience between the two studies to assess if agency impacted how the information gathering process felt. We found that in Study 2, participants reported overall higher ratings of enjoyment ($\beta = 4.93$, 95% CI [4.33, 5.54], $t = 15.89$, $p < 0.001$), lower ratings of frustration ($\beta = -8.43$, 95% CI [-9.19, -7.67], $t = -21.71$, $p < 0.001$) and not a statistically significant difference in tension ratings ($\beta = -0.65$, 95% CI [-1.40, 0.09], $t = -1.72$, $p = 0.09$) (Fig. 4D). These findings suggest that allowing participants to freely decide when to make Mid-Video Decisions further heightened the enjoyment of watching these videos and reduced the experienced frustration, pointing to the additional reward

value inherent in exerting agency (Leotti & Delgado, 2011). Collectively, these results reveal that when participants could control how much information they gathered, they were more likely to wait until they reached resolution about the identity of the drawing, often collecting more information, before stopping to move to the next trial, and the ability to refine one's information gathering was linked to a more positive affective experience.

3.3.2.4 Study 2 Conclusions

In Study 2, we replicated our findings from Study 1, demonstrating that even when participants are given agency over exactly how much information to gather, curiosity still predicted delaying the receipt of the outcome to prioritize experiencing information unfold across time. We found that similar to Study 1 and across both the Forgo Resolution and Jump-to Resolution conditions, high subjective states of confidence and correctly guessing the drawing also led to early stopping, providing further support that realizing the identity of the drawing prompted terminating the information seeking process. Again, we found that curiosity was positively related to the affective states of enjoyment and tension and was not related to frustration. Moreover, curiosity was inversely related to confidence and guess accuracy, but in contrast to Study 1, we found that curiosity was also inversely related to the time elapsed watching a video, suggesting a decline in curiosity as information gathering moves on that may have been obscured by the Predetermined Timing of decisions in Study 1. When we assessed the differences between Study 1 and Study 2, we discovered that giving participants agency over exactly when they provided Mid-Video Decisions augmented their sensitivity to experiencing resolution, such that participants waited longer during video watching, were more likely to stop videos early when they had reached resolution and were more likely

to continue videos when they were uncertain. When we examined the affective experience between Study 1 and Study 2, we found that in Study 2, participants reported greater enjoyment, lower frustration and no difference in tension. These findings suggest that allowing participants to freely decide when to make Mid-Video Decisions increased the overall positive affect experienced during video watching, pointing to the additional reward value inherent in exerting agency (Leotti & Delgado, 2011). Overall, Study 2 provided further support that high states of curiosity promote continuing to engage in video watching compared to seeking immediate outcomes and showed how the experience of resolution can signal an end to the value of information gathering and thus promote cessation.

3.4 General Discussion

Across two experiments using a novel dynamically evolving line drawings paradigm, we showed that higher curiosity predicted the choice to delay resolution not only when the alternative was to forgo resolution, but even when the alternative was to immediately resolve uncertainty and view the final drawing. These findings support the notion that under continuously evolving information, curiosity can prioritize the process of information gathering above simply obtaining the outcome. In both experiments, the subjective state of curiosity was characterized by enjoyment and tension, but not by frustration. While curiosity encouraged continuing to watch videos, confidence and guess accuracy drove choices to stop videos early, indicating that reaching resolution changed the value of continued information gathering. Moreover, giving participants freedom over when they could choose to continue or stop videos (Study 2) strengthened the relationship between resolution and choice, such that participants were more likely to

continue watching when uncertain and more likely to stop once resolution was reached. Overall, these results demonstrate that curiosity promotes withholding early reception of an outcome in contexts that afford the opportunity for discovery.

Amplifying and extending prior work exploring curiosity using discrete-resolution tasks (e.g., trivia Q&A paradigms), we demonstrate that curiosity also arbitrates between seeking versus forgoing information when it is continuously evolving (Kang et al., 2009; Lang et al., 2021; Marvin & Shohamy, 2016; van Lieshout et al., 2018). High levels of uncertainty, as measured through low levels of confidence and incorrect guessing, also increased choices to continue information seeking, supporting past research (Bennett et al., 2016; Blanchard et al., 2015; Brydevall et al., 2018; Marvin & Shohamy, 2016).

Critically, we also extend prior work by demonstrating that when information evolves across time, curiosity can discourage immediate resolution and instead encourage prolonging the state of uncertainty, similar to how people prefer to experience movies, sports, or narratives that unfold across time (Boyd-Graber et al., 2013; Ely et al., 2015; Maxwell, 2021; Su-lin et al., 1997). These findings demonstrate that the process of information gathering also holds value beyond the outcome of an event and that curiosity is sensitive to this value of experiencing information unfold. Expanding on recent work that shows people prefer receiving hints compared to immediately seeking the answer when they are curious (Metcalf et al., 2021), here we find the intentional prolonging of information gathering happens most often when people are uncertain about that outcome; once resolution is reached there is a distinct decrease in the desire for more information. Moreover, we demonstrate that curiosity also changes across information gathering, decreasing as participants become more confident in the drawing's identity,

evincing that curiosity is specifically tuned to the dynamic informational state of an individual.

In both of our experiments, curiosity was related to enjoyment and tension but not with frustration. These findings illustrate that the experience of curiosity can be associated with positive anticipation, which stands in contrast with previous work characterizing the state of curiosity or the prolonging of uncertainty as aversive (Bennett et al., 2016; Hsee & Ruan, 2016; Jepma et al., 2012; Loewenstein, 1994) and linked to unhappiness (van Lieshout et al., 2021). This signals that one's state experience of curiosity can likely take on a variety of affective flavors. Although the link between curiosity and enjoyment has been explored in personality trait differences in information seeking (Jach et al., 2021; Kashdan et al., 2018; Litman, 2010), understanding within-person changes in the affective experience of curiosity are underexplored but can provide deeper insight into how different information seeking behaviors are prioritized (Noordewier & van Dijk, 2020). For example, curiosity can prioritize the process of gathering when the affective experience is positive, as in our task, but it is likely that when curiosity is paired with more negative affective states, an outcome-focused approach would prevail, similar to the preference for immediate vs. delayed information seen in other tasks (Blanchard et al., 2015; Charpentier et al., 2018; Rodriguez Cabrero et al., 2019; van Lieshout et al., 2021). Moreover, we found that increasing a person's agency over the timing of their choices (Study 2) further increased the positive affect experienced, similar to previous work emphasizing the value of exerting control (Jiwa et al., 2021; Leotti & Delgado, 2011). Consequently, curiosity and agency can guide choices towards prolonging or resolving uncertainty and characterizing the affective experience of curiosity is essential to understanding its impact on behavior.

Participants in Study 2 could control when they rated their curiosity and when to stop or continue videos. The addition of this control highlighted the importance of reaching resolution during the information gathering process. Specifically, compared to participants in Study 1, participants in Study 2 were more likely to continue watching when they were uncertain but also more likely to stop watching when they resolved the drawing's identity. Stopping videos early in Study 2 also came at later time points of the video, suggesting that participants waited until after resolution was reached before they stopped watching. These findings suggest that people use the subjective experience of resolution as a marker of when the process of information gathering can be concluded, providing key insight into how people self-organize their information seeking (Fastrich & Murayama, 2018; Murayama et al., 2019). Notably, between Study 1 and Study 2, we saw contrasting relationships between curiosity and elapsed time watching a video; where in Study 1, curiosity was positively predicted by time elapsed and in Study 2, it was negatively predicted by time elapsed. Such differences could be due to the interaction between timing information and agency. Timing information appeared more salient in the choice process in Study 2, whereas in Study 1, timing information was only a significant determinant of choice in the Jump-to Resolution condition. The role of time in curiosity-driven choices has remained an underexplored topic (but see (Noordewier & van Dijk, 2017) and future work is needed to unpack how agency and time spent information seeking interact with the subjective state of curiosity.

The divergent effects of curiosity on information seeking generates new questions about boundary conditions: *when* does curiosity prioritize the process of information gathering over a rapid resolution of uncertainty? Future studies will need to elucidate what factors shape the experience of uncertainty and choices to relieve or

continue it. For example, a person's expectations about the process and outcome of anticipated information have previously been shown to influence choice. Prior work has demonstrated that the expected valence of the outcome impacts information seeking, where information expected to be positive is often sought after, while information expected to be negative is often avoided (Charpentier et al., 2018; Marvin & Shohamy, 2016), (but see (Oosterwijk, 2017; Oosterwijk et al., 2020)). Expected valence could also influence preferences on *when* that information is revealed, such that when an outcome is expected to be negative, curiosity could favor immediate reduction of uncomfortable uncertainty, while when an outcome is expected to be positive or neutral, curiosity could prioritize the process of information gathering (Dan et al., 2020). Furthermore, expectations around the structure of the environment can also influence information seeking choices. An individual might deem prolonging uncertainty to be worthwhile when its resolution is definite, but not when resolution may not occur (Lang et al., 2021; Ligneul et al., 2018). Lastly, another important factor to consider is what process is happening while the resolution of uncertainty is being delayed. In our task, information continued to unfold across time, allowing for active hypothesis testing and updating. This stands in contrast with delay periods in which the participant passively waits for resolution, without the opportunity to gather more information (Blanchard et al., 2015; Iigaya et al., 2020, 2016). Under conditions where no further information is obtained, receiving immediate information, indeed, seems favored and is linked to anticipatory utility (Iigaya et al., 2020). Future work may delineate under what contexts curiosity promotes immediate uncertainty reduction or engaging in a prolonged gathering process.

What does it mean if curiosity sometimes directs choices towards and other times away from immediate resolution of uncertainty? One implication is that curiosity can prioritize understanding *how* information fits together. Curiosity has been linked to the process of sense-making (Chater & Loewenstein, 2016), explanation-seeking (Brockbank et al., 2022; Gopnik, 2000; E. G. Liquin & Lombrozo, 2020a; Schwartenbeck et al., 2019), and in the creations of knowledge networks (Lydon-Staley et al., 2019; Zhou et al., 2020). All of these conceptualizations emphasize the process of experiencing an evolving informational state in building representations of the world. Thus, the choice to suspend resolution in this study could reflect curiosity's preference for and valuation of understanding the process of how each drawing comes to fruition. Delaying resolution to experience the process of information gathering can be a rational solution to maximize information gain (Dubey & Griffiths, 2020) or help account for the slow and steady learning as documented in the Learning-Progress Hypothesis (Gottlieb & Oudeyer, 2018; Kaplan & Oudeyer, 2007). Thus, curiosity as prioritizing the process of information gathering offers a new avenue to extend current theoretical and computational models of non-instrumental information seeking.

Overall, we developed a novel dynamic line drawing task that provides a unique way to examine the evolution of curiosity over continuous information gathering as opposed to isolated, discrete episodes. Our results demonstrate that the experience of curiosity does not necessitate a need for immediate resolution but can instead promote choices to prolong a state of uncertainty, encouraging patience over instant gratification. This experience of curiosity, characterized not by frustration, but rather by tension and enjoyment, could reflect a state where uncertainty presents opportunities for exploration,

hypothesis generation, and testing. Our work thus demonstrates the reward value inherent in reaching a desired resolution, in due time, via a process of discovery.

4. The process of discovery engages distinct neural correlates across unfolding information

4.1 Introduction

Discovery, the act of uncovering new information or the resolution of uncertainty, is an essential tool for learning. Previous work has demonstrated the importance of discovery in driving exploration and building knowledge across development (Brockbank et al., 2022; Gopnik, 2000; Gottlieb & Oudeyer, 2018; Ripollés et al., 2016; L. E. Schulz & Bonawitz, 2007). Specifically, the motivation to experience discovery directs choices to seek out explanations (E. G. Liquin & Lombrozo, 2020b, 2021; Wojtowicz et al., 2022), to investigate the causal structure of our environment (Gopnik et al., 2007; Schulz & Bonawitz, 2007; Ten et al., 2021), and even to delay premature resolution of uncertainty (Hsiung et al., 2022). Beyond shaping choices, experiencing discovery also benefits encoding (Jepma et al., 2012; Ripollés et al., 2016; Van de Cruys et al., 2021), and promotes positive affect (Danek et al., 2014; Hsiung et al., 2022; Shen et al., 2016). Thus, understanding how the process of discovery arises and shapes information seeking and memory formation, reflects a core goal in the study of psychology and neuroscience.

The process of discovery, however, can be challenging to measure. One obstacle is characterizing how people transition from uncertainty to resolution. In causal reasoning tasks, discovery happens through trial and error as a learner constructs an understanding of the bounds and relationships with variables in the environment (Ruggeri et al., 2019; Schulz & Bonawitz, 2007; Wang et al., 2021). Conversely, insight type problems often describe discovery as sudden and unpredictable (Danek et al., 2020; Metcalfe & Wiebe, 1987). Another form of discovery involves disambiguating

perceptual information into its semantic form (e.g., seeing the leaves of a bush rustling and discovering there's a bird making a nest). Research into perceptual discovery has found that perceptual ambiguity elicits curiosity (Cohanpour et al., 2022) and the relief of this ambiguity elicits the experience discovery, the “Aha!” moment, and can enhance memory (Jepma et al., 2012; Van de Cruys et al., 2021). Yet, these studies treat the experience of disambiguation as instantaneous, overlooking the dynamic interplay between a viewer's informational state and the unfolding information. In contrast, disambiguation is often an active and multi-step process that involves learners probing their environment and updating their beliefs as information continues to unfold (Voss et al., 2011; Yang, Lengyel, et al., 2016; Yang, Wolpert, et al., 2016). Consequently, it is not just the relief from uncertainty that is important to understand, but also the process that leads to such relief. However, the evolution and dynamics of the process of discovery for perceptual information remains unexplored.

Another barrier toward understanding the process of discovery is characterizing how the brain supports the dynamic updating of a learner's belief state as information is accumulated. Prior work on perceptual information accumulation posits that the building of evidence is represented in parietal cortex (Gold & Shadlen, 2007; Kiani et al., 2008), but these studies have only examined evidence accumulation as it relates to a subsequent choice and not toward the experience of discovery. In contrast, work examining the neural correlates of both the experience of insight and perceptual discovery have documented the role of dopaminergic reward systems, hippocampal regions, and prefrontal cortices in signaling the experience of reaching resolution or clarity (Aziz-Zadeh et al., 2009; Cohanpour et al., 2022; Jepma et al., 2012; Oh et al., 2020; Ripollés et al., 2016; Shen et al., 2018; Tik et al., 2018). Despite these advances,

this body of work has isolated their analyses to the *moment* of realization or discovery, leaving open investigation into the neural correlates that support the preceding process of gathering information until discovery.

Critically, it remains unclear whether the process toward perceptual discovery is a singular process or actually composed of several distinct cognitive processes that require unique neural circuitry support. For example, during the early phases of perceptual ambiguity, a learner may have no predictions about the semantic identity of a hidden object (e.g., seeing the leaves of a bush rustling without strong ideas of what could be causing it) but after further collection, candidate characters could be identified (e.g., a squirrel, or a bird) and tested until an answer is reached (e.g., it IS a bird). The exploration of potential identities is likely more related to semantic memory retrieval as learners select potential identifiers given the current perceptual information, and would rely more on memory retrieval processes including the posterior cingulate cortex, hippocampus, precuneus, and medial frontal cortex (Binder & Desai, 2011; Flanagin et al., 2023; Moscovitch et al., 2016; Sestieri et al., 2017). Conversely, the active testing and anticipation of discovery is more akin to the moments leading up to realization documented in prior insight and perceptual discovery studies, which report corresponding activation in regions associated with object recognition (Cohanpour et al., 2022; Dolan et al., 1997; Hsieh et al., 2010) as well as reward and memory encoding systems (Aziz-Zadeh et al., 2009; Jepma et al., 2012; Tik et al., 2018). Alternatively, if this process is homogenous it might involve the same select regions shown to be active in perceptual information accumulation (Ding & Gold, 2012; Gold & Shadlen, 2007; Kiani et al., 2008; Shadlen & Newsome, 2001). Thus, whether the processes of exploring for

potential identities and the testing of those candidate identities represent the same or unique components toward perceptual discovery is currently unexplored.

Taken together, there are several key questions regarding how the brain supports the cognitive journey towards discovery: characterizing how information unfolds across time, and whether the path to discovery represents a singular process or multiple interacting cognitive processes of collecting and building evidence. We propose that the process for perceptual discovery under continuous information gathering is characterized by an ongoing transition from evaluating perceptual evidence to a semantic experience of recognizing the identity of a stimulus. We posit that this transition from perceptual ambiguity to semantic realization is comprised of two distinct phases: an exploration phase where participants generated potential identities for the drawing, and a testing phase where evidence toward potential identities accumulated until realization. Critically, we hypothesize that these phases are also supported by distinct neural systems. Yet, little work has investigated the potential for different cognitive or neural components of the discovery process.

In the present study, we investigated the neural correlates supporting the process of discovery by leveraging a perceptual task in which participants watched videos of complex line drawings as they slowly resolved into objects. As participants watched these videos in the MRI scanner, they indicated via button press when they first had an idea about the eventual identity of the drawing and again when they reached resolution about the object in the drawing. Using these individual behavioral responses, we divided the video watching experience into two distinct processes: an Explore Phase occurring prior to their first idea, where participants explored potential identities for the drawing, and a Hypothesis Testing Phase following the first idea to resolution, where

evidence for potential identities accumulated until realization. To isolate the neural correlates specific to the experience of discovery, we included a subset of videos that never resolved into objects but were otherwise perceptually identical to the other videos. We hypothesized that the process to discovery for resolvable objects, compared to unresolvable drawings, would elicit greater activation in regions associated with exploration, semantic memory, and anticipation, and that this differentiation would increase as participants moved from the Explore Phase to the Hypothesis Testing Phase. Moreover, we hypothesized that each phase would also be supported by distinct neural correlates. Lastly, we examined how the process toward realization influenced memory formation by having participants complete a surprise memory test 24-hours later, where they rewatched the same videos and indicated whether they remembered seeing the video the previous day.

4.2 Methods and Materials

4.2.1 Participants

Forty-seven participants were recruited from the Duke Brain Imaging and Analysis Center subject pool. All participants completed written informed consent prior to participation in the study and received monetary compensation for their time (\$25.00/hr for the scanning session, \$20.00/hr for the online behavior session). The study was approved by the Duke University Health System Institutional Review Board. Participants were right-handed and had normal or corrected-to-normal vision. Two participants were excluded for not providing responses for more than 50% of the trials, two people were excluded for technical errors during encoding, and five participants did not return for the Session 2 online session. Our final sample used in all behavioral and neuroimaging

analyses consisted of 38 participants ($M_{\text{age}} = 23.59$ years, $SD = 3.57$ years; 24 females, 14 males).

4.2.2 Experimental Procedure

The study consisted of two sessions separated by 24 hours (Figure 1). Session 1 took place in an fMRI scanner, where participants watched videos of single-line drawings unfold over time. Each drawing consisted of a single item that could be named using a single word. The items belonged to a variety of categories, including objects (e.g., guitar, headphones), food (e.g., taco, pepper), people (e.g., ballerina, golfer), and animals (e.g., tiger, eel). Each trial started with a fixation cross (2s) and a brief blank screen (1s) before the start of each video, which played automatically and lasted 20 seconds. Videos could fall into one of two categories (Figure 1A). Two-thirds of videos were in the “Real” condition, as they eventually resolved into interpretable objects (e.g., a brain). The remaining third of videos fell into the “Scribble” condition, where the image never resolved into a recognizable object. Participants were told that they would experience both types of videos throughout the task but were not told what trials would contain real or scribble videos. As participants watched each video, they indicated via button press when they first had an idea of what the drawing could become (First-Guess) and a second time when they realized what the drawing was (Realized). These responses were later used to divide the video watching experience into two phases, Explore and Hypothesis Testing. A blank screen was presented for 1.25s at the end of each video, followed by a rating screen (3s) where participants rated how surprised they were by the outcome of the video using a 4-point scale. Each trial was separated by an ITI between 4 to 8 seconds (uniform distribution). During the ITI, participants completed an active baseline task where they were presented arrows (1s) and indicated which direction the

arrow was facing (left or right) (Stark & Squire, 2001). Participants completed three runs inside the scanner with 15 videos in each run. Each run contained 10 resolving videos and 5 scribble videos. Videos were randomly shuffled for each participant. Before the experiment began, participants received detailed instructions on how to complete the task and experienced several practice videos that were not included in the task, followed by a quiz to ensure their comprehension.

The following day, participants completed the second, online session, which consisted of a self-paced surprise incidental memory test. The session began with a free-recall test where participants were asked to report the name of all the objects that they remembered seeing in Session 1. Participants made their responses using a textbox, which populated a growing list used to remind participants of the names they had already reported. Participants then moved to the next portion of the memory test. Here, participants watched another set of single-line drawing videos and were asked to report whether they remembered seeing it the day before (Figure 1B). The set of videos included all 45 old videos presented in Session 1 (30 Real, 15 Scribble) as well as 10 new videos (5 Real, 5 Scribble). Videos played automatically at the start of each trial and participants indicated via a button press when they remembered the video. After their response, participants were prompted to report the identity of the drawing before the video finished playing. After each video, participants reported a final identity judgment for the object in the video, an old/new judgment on a 4-point Likert scale from “Definitely Old” to “Definitely New”, and a source memory judgment for which scan run of the task in Session 1 they remembered watching the video in.

The single-line drawings used in the tasks were collected through Google Image searches and transformed into standardized video stimuli using iMovie. Videos were

grayscale and adjusted to fit the center of the screen. All videos were standardized to 20 seconds in length. Scribble videos were constructed by combining together other real line drawing videos, separate from the Real videos used in this task. To program the task used in session 1, we used the Python framework, PsychoPy (Peirce et al., 2019). The task in session 2 was built in the JavaScript framework Vue.js and was hosted on a private virtual web server using psiTurk (Gureckis et al., 2016), which participants accessed through a browser.

4.2.3 Statistical analysis of memory data

To investigate predictors of memory, we focused our analysis on whether participants provided a correct mid-video “I remember” judgment during the video rewatch in Session 2 as this measure offered the most precision about memory performance. To score the mid-video judgments, we first removed all new videos and then confirmed that participants were correctly identifying the objects in the drawings by examining the final identity judgments they provided after each video was complete. Mid-video memory was then computed by comparing each participant’s mid-video answer to their final identity. Correct judgments were coded as 1 and incorrect judgments were coded as 0. Mid-video memory across all videos seen in Session 1 was 73%.

Behavioral analyses were conducted with R v4.0.3 and RStudio v1.3.1093. To investigate the effects of video type (Real versus Scribble videos), response timing (First Guess response and Realized response), and whether participants reported a Realized response, on mid-video memory, we ran a mixed-effects logistic regression using the lme4 package (Bates et al., 2015) and bound optimization by quadratic approximation (BOBYQA controller). The timing for both the First Guess and Realized response were treated as continuous predictors and centered and scaled. We included all two-way

interactions between video type and our other predictors to investigate whether the encoding process differed as a function of video type. We obtained p -values with the lmerTest package (Kuznetsova & Brockhoff, 2017). The model included random effects with both participant and video as random intercepts. For all statistical analyses, we report standardized β values (and 95% confidence intervals (Cis) around the slope estimate). Figures were produced using the raw data in Python with Seaborn and Matplotlib.

4.2.4 fMRI data acquisition and preprocessing

Neuroimaging data was acquired at the Duke University Brain Imaging and Analysis Center with a 3T General Electric MR750 scanner and 48-channel head coil. An MP-RAGE pulse sequence was used to acquire high-resolution T1-weighted structural scans (TR = 2.25 s; TE = 30 ms; image matrix = 256^2 ; voxel size = $0.94 \times 0.94 \times 1$ mm). Whole-brain functional scans were collected using an EPI-RT sequence parallel to the AC-PC line (TR = 2 s; TE = 30 ms; image matrix = 128^2 ; FOV = 192×192 mm²; voxel size = $2 \times 2 \times 4$ mm, 34 slices with interleaved acquisition, flip angle = 80°). The first four timepoints in each functional scan were discarded to ensure scanner stabilization. Participants provided responses and ratings using a 4-button MR-compatible response boxes in their right hand.

Preprocessing of fMRI data was performed using *fMRIPrep* version v22.0.1 (Esteban et al., 2020, 2019), which is based on *NiPype* 1.8.4 (Esteban et al., 2022; Gorgolewski et al., 2011). The following description of the anatomical and functional preprocessing steps is provided by *fMRIPrep*.

4.2.4.1 Anatomical data preprocessing

T1-weighted (T1w) images were corrected for intensity non-uniformity (INU) (Avants et al., 2008; Tustison et al., 2010) and were used as T1w-reference throughout the workflow. For each participant, the T1w-reference was then skull-stripped and brain tissue segmentation was performed on the brain-extracted T1w to segment cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) (Zhang et al., 2001). Volume-based spatial normalization to the standard ICBM 152 Nonlinear Asymmetrical template version space was performed through nonlinear registration using brain-extracted versions of both T1w reference and the T1w template (Evans et al., 2012; Fonov et al., 2009).

4.2.4.2 Functional data preprocessing

The following preprocessing was performed for each of subject's three BOLD runs. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) were estimated and then a spatiotemporal filtering was applied using FSL *mcfliirt* (Mark Jenkinson et al., 2002). BOLD runs were slice-time corrected using AFNI's 3dTshift (Cox & Hyde, 1997). The BOLD time-series (including slice-timing correction when applied) were resampled onto their original, native space by applying the transforms to correct for head-motion. The BOLD reference was then co-registered to the T1w reference using *mri_coreg* (FreeSurfer) followed by FSL's *flirt* (M. Jenkinson & Smith, 2001) with the boundary-based registration cost-function (Greve & Fischl, 2009). Co-registration was configured with six degrees of freedom. Four confounding time-series were calculated based on the preprocessed BOLD: framewise displacement (FD),

DVARs and two region-wise global signals. FD and DVARs were calculated for each functional run, both using their implementations in *Nipype* (Power et al., 2014). The two global signals were extracted within the CSF and WM masks. All resamplings can be performed with a single interpolation step by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using *antsApplyTransforms* (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels (Lanczos 1964). Non-gridded (surface) resamplings were performed using *mri_vol2surf* (FreeSurfer). Prior to analysis, the outputs from fMRIPrep were skull-stripped using FSL's brain extraction tool (BET), spatial smoothing with an isotropic, Gaussian kernel of 5mm FWHM (full-width half-maximum), and high-pass filtered at 100s. For more details of the pipeline, see the workflows in *fMRIPrep's* documentation.

4.2.5 Univariate fMRI analyses

Univariate general linear model (GLM) analyses were conducted using FSL's FEAT (v6.0.0). We built two GLMs, one to examine the process of resolution (Resolution Model) and one to investigate subsequent memory differences for only Real videos (Memory Model). First-level GLMs for each model were created for each participant and for each run. For the Memory Model, we included boxcar functions to model the onset and durations for remembered and forgotten Real videos (2 EVs). For the Resolution Model, we split our trial into two periods based on the responses of our participants. We defined a boxcar function based on the time preceding the "First Guess" response as our "Explore" period, where participants gather information toward the first idea about what the drawing could be. We defined a boxcar function that started at the offset of the "First

Guess” and continued until the “Resolution” response as the “Hypothesis Testing” phase where participants were updating and refining their guess toward finally identifying the drawing (Figure 1A). Consequently, the timing and duration of each phase was set by our participants on each trial. For scribble videos that did not have a resolution response, we estimated the response time by taking the average time point of resolution across all real videos based on the first guess timing provided in the scribble video. Videos without any button presses were excluded from analysis. Both first-level GLMs also included confound regressors for CSF, WM, 24 motion parameters (3 translational, 3 rotational, motion parameters squared, all derivatives, and squared derivatives), and the top 6 components of the anatomical PCA, CompCor (Behzadi et al., 2007). All regressors were convolved with a double-gamma hemodynamic response function. For our Resolution Model, our primary contrasts investigated the difference between Real and Scribble videos during the Explore and Hypothesis Testing phase and the difference between the Explore and Hypothesis Testing phase for Real videos. For the Memory Model, our primary contrasts examined the difference between Real videos that were subsequently remembered against Real videos subsequently forgotten.

Each participant’s three runs were then combined at second-level analyses using a fixed-effects approach and this was then used as input into the group-level group analysis, using FMRIB’s mixed effects approach (FLAME 1). The group results were assessed at a cluster-forming threshold of $z = 2.58$ and a cluster significance threshold of $p = .05$. This combination of FLAME 1 with cluster forming and significance thresholds provides an effective control against inflated false positive findings in fMRI studies (Eklund et al., 2016).

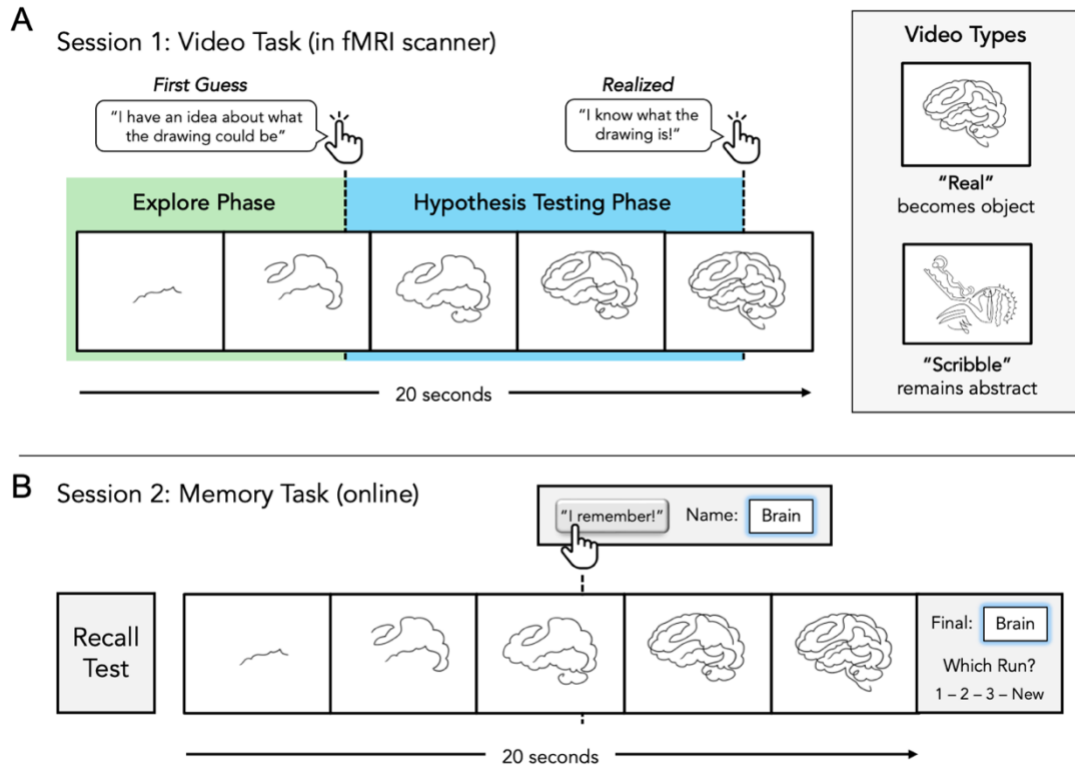


Figure 11: Two session task design.

4.3 Results

4.3.1 Behavioral Response Distribution

We first investigated the distribution of responses throughout video watching. In the time preceding the First Guess when there was very little information available, we hypothesized that participants would be more uncertain about whether the current video was a Real or Scribble drawing. Consequently, we expected a little difference in the likelihood of making a First Guess and the timing of this guess. First Guess responses were made on 86% of videos (Real: 88%, Scribble = 83%). On average, participants made their First Guess responses for Real videos at 8.1 seconds (SD = 3.4 seconds) and their First Guess responses for Scribble videos at 8.0 seconds (SD = 3.8 seconds).

A 2-sample test for equality of proportions revealed that First Guess responses were provided on significantly less trials during Scribble compared to Real videos (Chi-squared = 6.31, $p = 0.012$), however when they were provided, there was no difference in the timing of when responses were made ($t = 1.03$, $p = 0.31$).

In contrast as the video continued, information discriminating whether the drawing was Real or Scribble also increased. Here, we anticipated behavioral differences in the likelihood and timing of Realized responses, specifically that there would be fewer Realized responses for Scribble videos and they would occur at later points in the video. Realized responses were made on 71% of videos (Real: 83%, Scribble = 48%). On average, participants made their Realized responses for Real videos at 14.6 seconds (SD = 3.4 seconds) and their Realized responses for Scribble videos at 15.7 seconds (SD = 2.9 seconds). A 2-sample test for equality of proportions also revealed a significant difference (Chi-squared = 235.65, $p < 0.001$) in the proportion of Realized responses, again showing that there were less Realized responses made during Scribble compared to Real videos, although this difference was greater for Realized compared to First Guess responses. Additionally, there was a significant difference in the timing of Realized responses between Real and Scribble videos ($t = 6.47$, $p < 0.001$), supporting our hypothesis that when participants made Realized responses for Scribble videos, they made them slower than those for Real videos. Although this effect is robust, it is important to note that the estimated timing difference between Real and Scribble Realized responses was 1.4 seconds. Collectively, these descriptive statistics revealed that during the first part of video watching, there were smaller differences between the experience of watching Real and Scribble videos but as

information accumulated, the difference between Real and Scribble videos exerted changes in behavioral responses.

4.3.2 Behavioral Correlates of Memory

We next sought to determine how the experience of watching these drawings unfold in Session 1 related to memory for the same videos in Session 2. Memory was assessed through the accuracy of mid-video “I remember” judgments that participants made during video rewatch of Session 2. We found that Real videos produced more correct mid-video memory judgments compared to Scribble videos ($\beta = 1.52$, 95% CI [2.30, 0.75], $z = 3.84$, $p < 0.001$, Figure 2), as did whether participants reported a realization response during Session 1 ($\beta = 1.36$, 95% CI [0.80, 1.92], $z = 4.78$, $p < 0.001$, Figure 2). Both the timing of First Guesses and Realized responses were not significant predictors of later memory (all $p > 0.05$). Moreover, there was a significant interaction between video type (Real versus Scribble) and realization ($\beta = 0.96$, 95% CI [0.11, 1.81], $z = 2.20$, $p = 0.027$, Figure 2). Post-hoc analyses indicated that the memory enhancement of realization was only significant for Real videos ($z = 4.78$, $p < 0.001$) and not for Scribble videos ($z = 1.16$, $p = 0.74$). Taken together, when videos resolved into interpretable objects, participants had improved memory for them. Similarly, the act of reaching resolution during video watching enhanced memory but this effect was specific to when participants were watching Real videos.

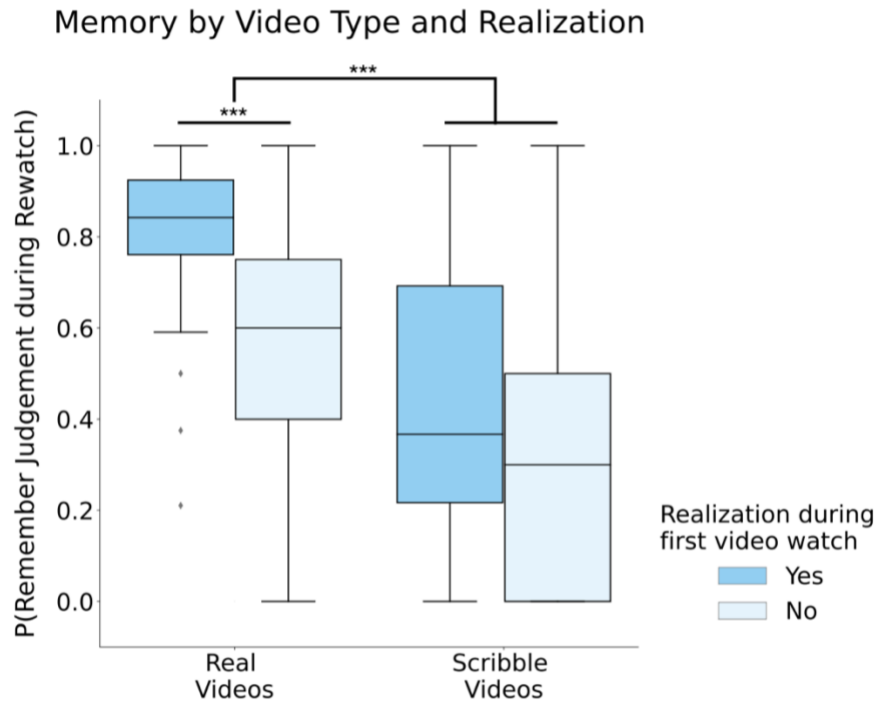


Figure 12: Memory results for session 2. Rates of remembering a video during re-watch (Session 2).

4.3.3 Univariate Whole Brain Activation

A key feature of experiencing information unfolding across time is that it allows for the information collected to build upon itself. Within this task as line drawings unfolded, participants slowly transitioned from the perceptual experience of tracking the line to a semantic experience of recognizing the identity of the drawing. We posited that this transition from perceptual ambiguity to semantic realization could be divided into two distinct components: an exploration and information gathering phase where participants generated potential identities for the drawing (Explore Phase) and a testing and updating phase where evidence toward potential identities accumulated until realization (Hypothesis Testing Phase). Trials were thus divided into these two phases using the participant-specific indications with the Explore Phase starting at the onset of the video

and continuing until their First Guess and the Hypothesis Testing Phase starting after their First Guess and continuing on until they reported realizing the identity of the drawing (Realized) (Figure 1).

Additionally following our behavioral data, we hypothesized that these phases would be differentially sensitive to whether videos resolved into objects (Real) or never resolved (Scribble). We expected that during the Explore Phase there would be fewer differences in BOLD signal activation between Real and Scribble videos since there was more uncertainty about whether the drawing was going to resolve. Conversely, across the Hypothesis Testing Phase as participants moved toward realization, we anticipated larger differences and more diversity in the neural activation patterns between Scribble and Real videos. To organize the results below, we first investigated the differences between Real and Scribble videos across the two information gathering phases and then examined the differences between the Explore and Hypothesis Testing Phases specifically for Real videos. Finally, we assessed the brain areas that were preferentially active during encoding for Real videos that were later remembered compared to those that were forgotten.

4.3.3.1 Explore Phase

We first tested for differences in brain activation between Real and Scribble videos during the Explore Phase. We examined both the contrast for Real videos greater than Scribble videos as well as the reverse contrast. Since we anticipated this period of the trial to function similarly between Scribble and Real videos, we did not have strong hypotheses about any specific differences between the two video types. However, for Real videos we found significantly increased activation in the posterior cingulate cortex, precuneus, medial prefrontal cortex, and right hippocampus (Figure 3). These regions

have been shown to be active during retrieval of information (Flanagin et al., 2023; Gilmore et al., 2015; Moscovitch et al., 2016; Sestieri et al., 2017) and also in regulating moving between internally and externally directed focus (Leech et al., 2012), suggesting that during exploration for Real videos, participants were more actively surveying potential identity candidates. Conversely, for Scribble videos, there was increased activation compared to Real videos in early visual cortex, supplementary motor area, and right anterior insula, potentially reflecting increases in externally focused attention (Gandhi et al., 1999; Offen et al., 2009) and monitoring (Bonini et al., 2014) (Table 1). Together, these findings highlight that even leading up to the First Guess, participants were starting to show different patterns of activation for Real compared to Scribble videos.

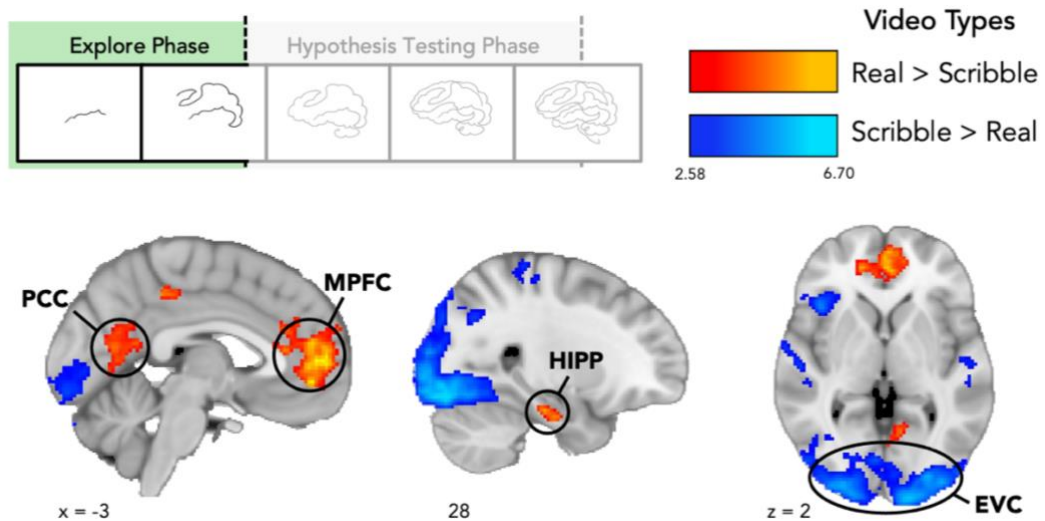


Figure 13: Whole-brain during Exploration Phase (for a full list of regions, see Table 1).

Table 1: Effects of video type during the Explore Phase on whole-brain activity.

Brain region	Hemi	Cluster	Z-statistic	MNI peak coordinates (mm)			Voxels
				X	Y	Z	
Explore Phase	--	--	--	--	--	--	--
Real > Scribble	--	--	--	--	--	--	--
Medial prefrontal cortex / paracingulate gyrus	L	6	4.91	-2	54	6	1461
Lateral occipital complex	L	5	4.85	-42	-76	32	597
Precuneus / intracalcarine cortex	L	4	4.69	-12	-62	8	433
Supramarginal gyrus	L	3	4.38	-56	-30	46	361
Posterior cingulate cortex	L	2	3.8	-2	-38	44	183
Right Hippocampus	R	1	3.99	26	-20	-22	153
Scribble > Real	--	--	--	--	--	--	--
Early visual cortex, inferior lateral occipital cortex	R	6	7.57	22	-84	-8	12174
Middle frontal gyrus / insula	R	5	5.08	42	22	2	1735
Planum temporale	R	4	4.34	62	-6	6	517
Planum temporale	L	3	4.3	-46	-32	6	402
Supplementary Motor Area	R	2	3.82	12	2	58	298
Post-central gyrus	R	1	3.47	28	-34	68	164

4.3.3.2 Hypothesis Testing Phase

To investigate the difference between experiencing resolution for resolvable compared to unresolvable objects, we contrasted the Hypothesis Testing Phase between Real and Scribble videos. Overall, we anticipated that this contrast would show more distinct and varied activation patterns as participants came to realize what type of video they were watching than in the early Explore Phase. Specifically, we hypothesized

that the resolution of a drawing into an object would preferentially activate both reward and memory related regions, consistent with previous work on insight (Becker et al., 2020; Tik et al., 2018) and perceptual disambiguation (Cohanpour et al., 2022; González-García et al., 2018; Jepma et al., 2012).

Indeed, we found that during the Hypothesis Testing Phase, as participants gathered evidence about the identity of the drawing, there was increased activation in reward-processing regions including the VTA, bilateral NACC, and caudate as well as medial temporal lobe regions including the left PHC (Figure 4). We also found widespread increased activation in bilateral SMG, LOC, and OTC, regions commonly active during the recognition of the objects (Cohanpour et al., 2022; Dolan et al., 1997), and in the inferior frontal gyrus leading into DLPFC as well as DACC, regions show to reflect the evaluation and updating of information (Kaanders, Juechems, et al., 2021; Kobayashi et al., 2021; Shenhav et al., 2018) (Table 2). Taken together, these regions suggest that reaching realization is a rewarding experience and involves widespread prefrontal and temporal engagement, likely a reflection of the active testing and evaluation of different identities until the object is finally recognized.

When we examined the reverse contrast, we found that Scribble videos during the Hypothesis Testing phase elicited increased activation in the PCC and EVC extending into superior LOC (Table 2). Collectively, these regions suggest that when resolution is not offered, participants might still continue attending toward the perceptual line drawing searching for structure, in other words to stall in the Explore Phase.

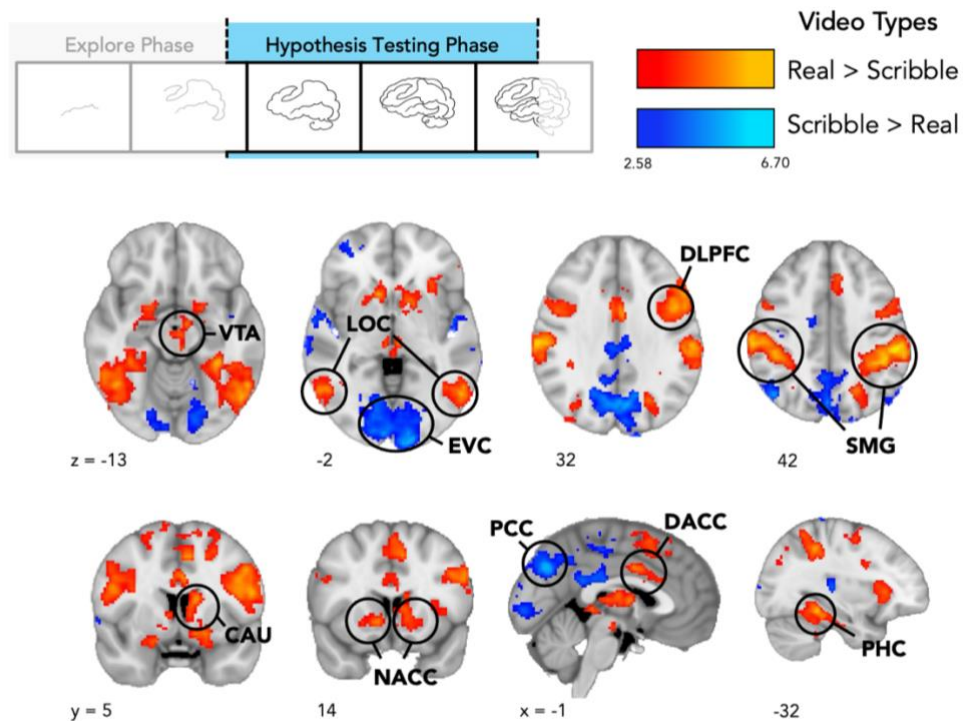


Figure 14: Reaching realization preferentially activated reward-related regions and object recognition regions for Real compared to Scribble videos.

Table 2: Effects of video type during the Hypothesis Testing Phase on whole-brain activity.

Brain region	Hemi	Cluster	Z-statistic	MNI peak coordinates (mm)			Voxels
				X	Y	Z	
Hypothesis Testing Phase							
Real > Scribble							
Precentral gyrus / inferior frontal gyrus	L	11	5.68	-54	6	22	3750
Middle frontal gyrus	L	11	5.38	-50	6	38	
Paracingulate gyrus	L	11	5.35	-6	22	46	
Superior frontal gyrus	L	11	5.3	-4	28	48	
Occipital temporal cortex / lateral occipital cortex	R	10	5.32	48	-58	-16	742
Temporal occipital fusiform cortex	R	10	5.15	38	-48	-26	
Caudate	L	9	5.6	-10	0	16	621

	Caudate / nucleus accumbens,	R	9	5.42	12	16	-2	
	Anterior cingulate gyrus	L	9	5	-4	0	32	
	Thalamus	L	9	4.87	-4	-22	-2	
	Inferior lateral occipital cortex	L	8	5.91	-54	-64	-4	417
	Inferior temporal gyrus,	L	8	5.77	-54	-58	-16	
	Posterior parahippocampal gyrus,	L	8	5.62	-32	-40	-14	
	Middle temporal gyrus	L	8	5.4	-50	-66	-16	
	Supramarginal gyrus	L	7	6.57	-54	-30	40	361
	Supramarginal gyrus	R	6	6.64	60	-26	38	2148
	Inferior frontal gyrus	R	5	4.84	50	4	24	762
	Superior parietal lobe / precuneus	L	4	4.79	-20	-64	48	616
	Middle frontal gyrus	R	3	4.19	30	-6	52	368
	Precentral Gyrus / Middle Frontal Gyrus	R	3	4.19	30	-6	52	368
	Lateral occipital cortex	R	2	4	44	-78	22	236
	Amygdala	R	1	4.22	20	2	-16	155
	Nucleus accumbens	R	1	2.98	10	8	8	
Scribble > Real		--	--	--	-	-	-	-
	Early visual cortex / precuneus / posterior cingulate cortex	L	7	6.59	-14	-102	10	8402
	Post central gyrus	L	6	4.88	-14	-34	74	2711
	Planum temporale	L	5	5.02	-42	-32	6	1216
	Planum temporale	R	4	4.49	52	-26	8	1181
	Superior lateral occipital cortex	R	3	5.03	46	-66	42	398
	Superior lateral occipital cortex	L	2	4.4	-44	-72	46	171
	Frontal pole	R	1	4.1	38	56	4	155

4.3.3.3 Comparing Explore and Hypothesis Testing Phases

Next, we compared the experience of the Explore Phase to the experience of Hypothesis Testing. Because of the widespread differences in activation patterns between Real and Scribble videos during Hypothesis Testing, we focused our comparison on only Real videos. We hypothesized that during the Explore Phase, participants would toggle between examining the perceptual qualities of the line drawing and internally searching for potential identities that could fit the evolving drawing. After

candidate identities had been collected, participants would shift to increase their attention on the line drawing as they anticipated resolution.

Consistent with our hypothesis, we found that Explore over Hypothesis Testing preferentially increased activation in areas relevant for memory and information retrieval, semantic processing, and exploration (Figure 5). During the Explore Phase, we identified significant clusters in the bilateral precuneus, angular gyrus, PCC, and hippocampus, core regions involved in memory retrieval (Flanagin et al., 2023; Gilmore et al., 2015; Leech et al., 2012; Moscovitch et al., 2016) and also in semantic processing (Binder & Desai, 2011; Tik et al., 2018). Additionally, there was increased activation in bilateral insula and in DACC into MPFC, regions previously shown to be active when people make exploratory choices and collecting information (Blanchard & Gershman, 2018; Kaanders, Juechems, et al., 2021) (Table 3). These regions suggest that during Exploration compared to Hypothesis Testing, there is increased semantic exploration as participants search and retrieve potential identities for the drawing.

In contrast, the Hypothesis Testing Phase was characterized by increased activation in the visual cortex with extensions into the OTC and LOC, suggesting that participants were attending more to the identity of the line drawing as they moved towards resolution (Gandhi et al., 1999; Kravitz et al., 2013). Additionally, we found increased activation in bilateral caudate (Figure 5), potentially reflecting the reward value in discovering the identity of the drawing (Tik et al., 2018) and relief from perceptual ambiguity (Jepma et al., 2012). Finally, there was also increased activation in the inferior frontal gyrus, orbitofrontal cortex, and the right amygdala, regions that have previously been shown in experiencing insight and across problem solving tasks (Aziz-Zadeh et al., 2009; Shen et al., 2018; Tik et al., 2018) (Table 3).

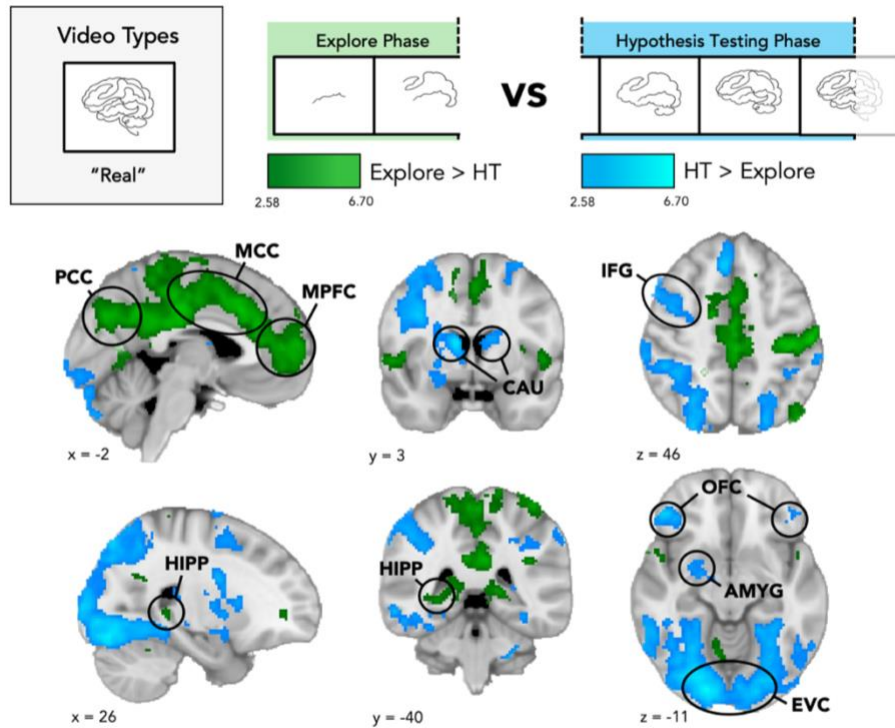


Figure 15: Whole-brain activation comparing Exploration versus Hypothesis Testing Phase (for a full list of regions, see Table 3).

Table 3: Effect of trial phase on whole-brain activity for Real videos only.

Brain region	Hemi	Cluster	Z-statistic	MNI peak coordinates (mm)			Voxels
				X	Y	Z	
Explore > Hypothesis Testing	--	--	--	--	--	--	--
Post central gyrus	L	6	6.58	-56	-20	50	14022
Precentral gyrus	L	6	5.97	-40	-20	58	
Precuneus	R	6	5.66	12	-66	34	
Precuneus	L	6	5.64	-12	-60	32	
Posterior cingulate gyrus	L	6	5.6	-4	-22	34	
Lateral occipital cortex	L	5	4.59	-50	-74	40	705
Angular gyrus	L	5	3.88	-36	-62	32	
Central opercular cortex	L	4	4.86	-50	-4	2	478
Insular cortex	L	4	3.42	-34	12	-8	
Right hippocampus / thalamus	R	3	4.2	16	-36	8	475
Parietal Operculum Cortex	L	2	4.82	-38	-36	26	401

	Central Opercular Cortex	R	1	4.3	56	2	0	289
	Insular cortex	R	1	3.85	38	8	0	
Hypothesis Testing > Explore		--	--	--	--	--	--	--
	Early visual cortex, inferior lateral occipital cortex / occipital temporal cortex	R	11	7.33	18	-88	-12	21238
	Middle frontal gyrus	R	10	6.65	36	8	34	5605
	Inferior frontal gyrus	R	10	6.38	54	18	30	
	Caudate	R	9	7.23	14	-4	20	1066
	Caudate	L	8	5.29	-18	-8	26	753
	Superior frontal gyrus	R	7	5.49	2	28	52	420
	Frontal pole / orbitofrontal cortex	L	6	4.73	-48	40	-4	393
	Precentral frontal gyrus / middle frontal gyrus	L	5	4.39	-56	14	34	
	Putamen	R	4	4.78	32	-22	0	320
	Amygdala	R	4	3.73	22	-6	-12	
	Supramarginal gyrus	L	3	4.89	-66	-32	28	262
	Superior parietal lobe	L	2	4.24	-38	-42	42	183
	Superior frontal gyrus	L	1	3.73	-22	-4	62	380

4.3.3.4 Subsequent Memory Effects

Lastly, we used participant's mid-video memory judgements to assess differences in activation patterns for videos that were later remembered compared to forgotten. We focused our analyses only on Real videos to examine the subsequent memory effects while controlling for the resolvability of drawings. For videos later remembered, there was increased activation in the left hippocampus and bilateral parahippocampal cortices, consistent with prior work underscoring the importance of the hippocampus and greater MTL structures in driving memory formation (Barense et al., 2005; Ergorul & Eichenbaum, 2004; Moscovitch et al., 2016) (Figure 6). Moreover, there was a significant cluster found in the brainstem, overlapping with midbrain nuclei, SN/VTA, suggesting videos that elicited greater dopaminergic activity were also better remembered (Adcock et al., 2006; Poh et al., 2022; Shohamy & Adcock, 2010). When examining the reverse contrast, we found increased activation in the EVC and bilateral

LOC for videos that were later forgotten (Table 4). Together, these findings highlight key involvement of both memory systems and reward structures in driving successful encoding of these videos.

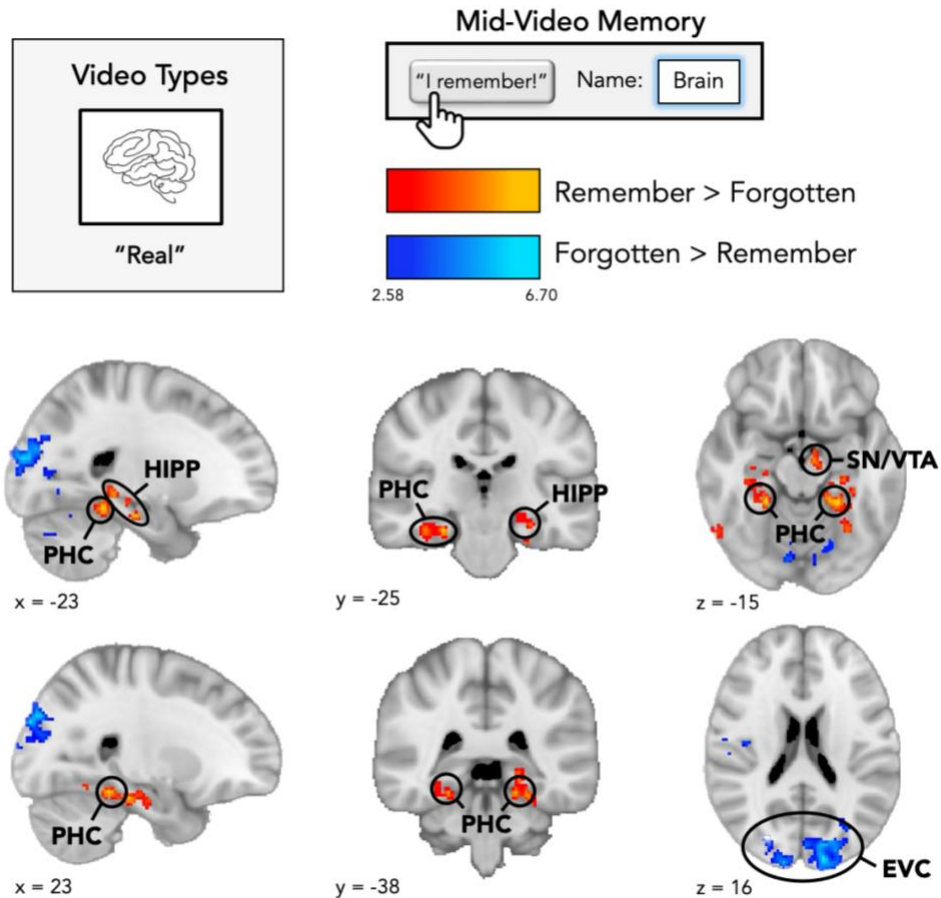


Figure 16: Real videos that were later remembered show increased activation in the hippocampus, PHC, and midbrain at encoding.

Table 4: Subsequent memory effect on whole-brain activity for Real videos only.

Brain region	Hemi	Cluster	Z-statistic	MNI peak coordinates (mm)	Voxels
--------------	------	---------	-------------	---------------------------	--------

				X	Y	Z	
Remembered > Forgotten							
	Hippocampus	L	4	4.55	-22	-20	368
	Temporal fusiform cortex / Parahippocampal gyrus	L	4	3.43	-34	-32	-18
	Parahippocampal gyrus / temporal fusiform cortex	R	3	4.5	26	-26	311
	Inferior temporal gyrus	R	2	4.28	60	-52	211
	Brainstem	L	1	4.86	-4	-12	160
Forgotten > Remembered							
	Early visual cortex	L	4	4.95	-16	-94	18
	Cerebellum	L	3	4.54	-16	-80	714
	Lateral occipital cortex	R	2	4.9	26	-84	32
	Lateral occipital cortex	L	1	3.87	-34	-72	10

4.4 Discussion

In the present study, we used a dynamically evolving line drawings paradigm to examine the neural correlates supporting the process toward discovery. Throughout the unfolding of these line drawings, we identified two distinct cognitive processes that characterized the path to discovery: an Explore Phase where participants generated potential identities for the drawing and a Hypothesis Testing Phase where evidence for potential identities accumulated until realization. Our neural analyses indicated that each phase was differentially sensitive to videos that resolved into real objects (Real videos) compared to videos that remained abstract (Scribble videos). Specifically, we demonstrate that during the initial Explore Phase, videos that resolved engage more midline brain regions including the PCC, precuneus, hippocampus, and cingulate cortex, suggesting greater recruitment from semantic memory and exploration areas. In the Hypothesis Testing Phase, we found greater activation in the prefrontal cortex, striatum, VTA, and parahippocampal cortex as the true identity of the drawing was revealed,

indicating that the experience of resolution elicited reward and memory regions. Moreover, even when isolating the process of discovery to only resolving videos, the Explore and Hypothesis Testing Phase showed distinct activation when compared to each other, suggestive of two unique components characterizing the discovery process. Finally, for resolvable videos, we show that videos that elicited increased activation in hippocampal and midbrain dopaminergic regions were also preferentially remembered. Collectively, these findings demonstrate that during ongoing evolution of information, people engage in a complex and multi-step process from searching, to testing, and finally discovery. Consequently, the neural correlates of discovery also change as a function of stage in the process as well as whether videos resolve or not.

Our results extend prior research on both insight problems and perceptual curiosity. Corroborating reports on insight problems, we demonstrated that reaching discovery compared to when it is not offered engages brain regions canonically associated with reward including areas of the striatum (caudate, NACC) and midbrain structures (SN/VTA) (Aziz-Zadeh et al., 2009; Shen et al., 2018; Tik et al., 2018). Our findings help extend this body of work by positing that the process of problem solving can actually be disentangled into distinct stages as people explore and test out different hypotheses toward solutions. Here, the contrast between exploration and hypothesis revealed distinct neural patterns differences, where increased activation in the semantic memory network dominated during exploration and increased activation in visual cortex, PFC, and reward anticipatory regions dominated during hypothesis testing. We posit that this difference in activation patterns marks a shift as participants move from sampling semantic space to gathering evidence toward various hypotheses. Particularly for insight problems where discovery is more sudden and unpredictable, delineating the paths that

people sample and how people decide when to stick with an idea are important directions to understand and better predict whether and when a solution will be reached (Bilalić et al., 2021). One remaining question is understanding whether these different types of discovery experiences (i.e., moving from perceptual ambiguity to a semantic understanding versus insight problems) elicit similar affective feelings. Future work should aim to directly compare the two experiences to further investigate how different types of discovery relate to different affective states and potential different supporting neural correlates.

The results from this study provide key extensions to understanding perceptual curiosity (Cohanpour et al., 2022; Jepma et al., 2012; Van de Cruys et al., 2021). We found that the relief of perceptual ambiguity increased activation in the caudate and ventral striatum, providing further support that reaching resolution was a rewarding experience (Jepma et al., 2012). We also find that regions important for the recognition of objects, including the lateral occipital cortex and the occipital temporal cortex (Dolan et al., 1997; Goodale & Milner, 1992), were more active when participants discovered what the object in the line drawing resolved into (Cohanpour et al., 2022; Jepma et al., 2012). Critically, we extend these findings by emphasizing the importance of considering the internal informational state of a person. By allowing information to unfold across time, creating a slow shift from perceptual ambiguity to resolution, we find that the experience of perceptual relief is tied to the internal understanding of a participant. In other words, when participants experienced information across time, they were able to titrate their behavioral responses to their own informational state, precisely signaling when they realized the identity of the drawing. This behavioral change helps link the experience of perceptual relief to that of insight problems, where participants play a

more active role in evaluating and searching for solutions. However, since we did not explicitly measure curiosity in their task, understanding how curiosity guides active search under perceptual ambiguity reflects an important next step for future research (Gottlieb & Oudeyer, 2018).

Similar to previous work, we also found that experiencing discovery enhanced memory for the videos a day later (Jepma et al., 2012; Van de Cruys et al., 2021), and such enhancement was also tied to increased activation in hippocampal regions (Jepma et al., 2012). Beyond the hippocampus, we identified a cluster in the midbrain that was more active in later remembered videos, suggesting that dopaminergic contributions facilitated encoding (Adcock et al., 2006; Poh et al., 2022; Shohamy & Adcock, 2010). Interestingly, we did not find any relationship between when participants indicated their first guess or their resolution and later memory. Given prior work investigating the benefits of delay on answers on later memory (Mullaney et al., 2014), we anticipated that the longer participants anticipated resolution, the better their later memory would be. In subsequent analyses of the current work, I plan to further investigate how the dynamics of the videos and participant's engagement with them drive subsequent memory. One potential analysis would be to treat the "I remember" mid-video memory judgment as a continuous variable based on the difference between when participants first indicated a resolution response during Session 1 versus when they reported remembering the video during their rewatch in Session 2. Understanding the neural patterns during encoding that could predict how far in advance the outcome of video can be subsequently identified remains an open question for this dataset. Such findings could provide a deeper look into how the transition from perceptual ambiguity to discovery shapes memory formation.

Within this study, we identified and differentiated two unique phases of the information gathering process toward discovery. Using participant responses, we divided trials into a first phase of exploration, in which participants were searching for possible identities for the drawing, and a hypothesis testing phase, in which participants gathered information toward potential identities until resolution was reached. These phases showed differential activation with each other as well as differential activation between resolvable and unresolvable videos that grew more distinct across information gathering. In the Explore Phase, we saw increased activation across the midline, including the PCC, ACC, and MPFC, as well as the precuneus and hippocampus. Collectively, these regions have been shown to play a role in memory retrieval (Flanagin et al., 2023; Gilmore et al., 2015; Hutchinson et al., 2014; Sestieri et al., 2017; Summerfield et al., 2009). During the first phase of the videos, it is highly likely that participants rely on semantic memory as they are searching for potential identities. Furthermore, comparing resolvable and unresolvable videos for the Explore Phase revealed increased activation in the PCC, MPFC, and hippocampus for videos that later resolved and increased activation in the EVC for videos that never resolved. This suggests that during the early part of the videos, participants were exploring both potential identities for the drawing and gathering evidence about whether they thought the drawing was likely to resolve or not. Indeed, the increased activation in the EVC could signal more externally directed attention toward the perceptual qualities of the video (Gandhi et al., 1999; Silver et al., 2007).

In the Hypothesis Phase the difference between resolving and unresolved videos grew more stark. For videos that resolved, we found greater activation in the caudate, NACC, and VTA, all regions involved in reward-processing and successful problem

solving (Daw & Doya, 2006; Knutson et al., 2001; Tik et al., 2018). In contrast, for videos that never resolved, we found increased activation the PCC, EVC, and bilateral superior LOC, areas involved in external attention (Gandhi et al., 1999), object recognition (Dolan et al., 1997), and semantic retrieval (Leech et al., 2012), and areas that were also active during earlier stages of information gathering. We interpret this as signifying that when resolution is not offered, participants may still continue searching for structure in line drawing. Lastly, when comparing the Explore and Hypothesis Testing phases for resolvable videos, we see increased activation in both frontal and early visual cortex areas, suggesting that as people identify potential guesses, they turn their attention back toward the video to gather evidence and anticipate the solution (Wager et al., 2005). One potential limitation of these phases is that participants were instructed to make these responses to indicate their process of information gathering, and thus are not organic parts of the process but were simply experimentally imposed. Although the instructions were explicit, guessing and updating guesses has been shown to be a natural component of watching these videos (Hsiung et al., 2022). Furthermore, we found distinct patterns of neural activation present in both phases, suggesting these phases represented different cognitive processes occurring over information gathering. Still, future work will need to confirm the time course and fidelity of these distinct phases of information gathering when people are unprompted.

Taken together, these findings identify unique neural patterns across information gathering as participants move from perceptual ambiguity to a discovery. Despite these initial findings, there is still much to explore within this dataset. To further investigate the process of discovery, I plan to conduct a series of inter-subject connectivity analyses to track how both the time course and the patterns of activation within key brain regions

change as a function of moving toward resolution. Specifically, I hypothesize that participant's brain states will be more diverse and dissimilar at the start of each video, but as information accumulates and evidence for a particular identity heightens, there will be increased similarity across participants. This can be corroborated by comparing the change in similarity for Real and Scribble videos. Furthermore, I will examine the patterns of functional connectivity between regions of interest identified in the whole brain contrasts to understand how people move from exploration to hypothesis testing and whether this shift represents a monotonic function or a strict step function. Finally, using the moment-by-moment time courses from an earlier behavioral study with these videos, I will study what brain areas are sensitive to the uncertainty and entropy of these videos as they unfold. Ultimately, these analyses will further advance our understanding of the process that unfolds across information gathering and the experience of discovery.

5. Complementary roles of curiosity and prediction error in information seeking and learning

5.1 Introduction

Seeking information is fundamental to navigating and flourishing in an uncertain and ever-changing world. Curiosity is important for directing learning, both in how information is sought and how information is later remembered (Gottlieb & Oudeyer, 2018; Gruber & Ranganath, 2019; Kidd & Hayden, 2015). When individuals report feeling curious towards the outcome of an event, whether for the answer to a trivia question or the outcome of a gamble, subsequent behavior is directed at obtaining revealing information, even at the expense of limited resources (Kang et al., 2009; Lau et al., 2020; Marvin & Shohamy, 2016; Rodriguez Cabrero et al., 2019; van Lieshout et al., 2018). Information obtained under high states of curiosity compared to low states is also preferentially remembered (Fastrich et al., 2018; Gruber et al., 2014; Kang et al., 2009; Marvin & Shohamy, 2016; Poh et al., 2022), and this enhancement has been demonstrated in both immediate memory (Lang et al., 2021; Poh et al., 2022) and also after a delay (Kang et al., 2009; Marvin & Shohamy, 2016; Stare et al., 2018). Curiosity, therefore, guides us to seek answers for questions of interest and improves the long-term memory for that information. Yet, few studies have investigated whether curiosity perpetuates search once a question has been answered.

Prior work has limited curiosity paradigms to focus on singular choices of information search, mainly relying on the assumption that the relief of curiosity is complete when an answer is revealed. This leaves open the question of whether and how the initial relief of curiosity may potentially perpetuate the desire to seek further information. In particular, the experience of surprise, such as that induced by a

prediction error about the answer could contribute to subsequent information search and how information is later remembered. Experiencing surprise has previously been shown to alter information seeking towards reconciling differences between the expected and the realized outcomes (Vogl, Pekrun, Murayama, & Loderer, 2019; Vogl, Pekrun, Murayama, Loderer, et al., 2019) and has also been shown to improve memory for encountered information (Antony et al., 2021; Kang et al., 2009; Pine et al., 2018; Rouhani et al., 2018). Consequently, one potential reason for when curiosity could induce further information seeking after an event has been resolved is if the outcome was unexpected or surprising. The current study, therefore, aims to understand how curiosity and surprise interact to drive sequential information seeking and memory formation, providing initial insight into the self-perpetuating properties of curiosity-driven search. Furthermore, previous work has also demonstrated that curiosity-guided information seeking enhances memory for that information (Kang et al., 2009; Lang et al., 2021; Marvin & Shohamy, 2016), but it remains unclear how the choice to seek information related to an original statement impacts memory for the original statement.

In the current study, we aimed to characterize the role of curiosity and surprise, via prediction error, in driving information seeking and memory formation. To capture the impact of both processes, we developed a task where participants learned about the veracity of a variety of trivia facts and urban myths. We specifically selected a balanced range of statements that were unlikely to be known as well as statements designed to elicit large negative prediction errors. After learning about each statement, participants had the opportunity to request for additional related information. This allowed us to examine the relationship between curiosity and surprise on information seeking that was not dependent on participants choosing to reveal the answer to the original statement of

interest. We hypothesized that participants would seek additional information about the original statement if they expressed curiosity in the statement or if the statement elicited a large negative prediction error (high-confident wrong answer). After reading through the trivia statements, participants then underwent a memory test either immediately (Immediate Memory Group) or 24 hours later (Delayed Memory Group). We hypothesized that successful learning would be predicted by ratings of curiosity during encoding as well as the size of the prediction error experienced when the answer was revealed. Additionally, we predicted that the choice to seek information would also improve memory for the original statement, in addition to the factors driving the search for extra information.

5.2 Methods and Materials

5.2.1 Participants

We recruited a total of 222 US adults (Immediate Memory Group: $N = 100$, Delayed Memory Group: $N = 122$) through the online labor marketplace, Prolific, in January of 2023. We collected a slightly higher number of participants in our Delayed Memory group to account for attrition rates. 6 participants did not return to complete the memory test for our Delayed Memory Group, leaving a final sample of 116 for that condition. Within our sample, 52% of our participants identified as women, 45% identified as men, and 3% identified as gender non-conforming or preferred not to say. The mean age of our participants was 33.5 years ($SD = 8.2$ years; range, 18 – 50 years). Our racial demographic breakdown showed that 64% of our participants identified as White, 8% as Asian, 8.5% as Black, and 20% as either mixed or other. There were no differences in demographics between our two memory groups. Our study in total was

estimated to take 55 minutes, either in one session (Immediate Memory Group) or across two sessions separated by 24-hours (Delayed Memory Group). Participants were paid a flat \$12.00 for completion of the entire study. All participants received informed consent under the guidelines of Duke University's Institutional Review Board.

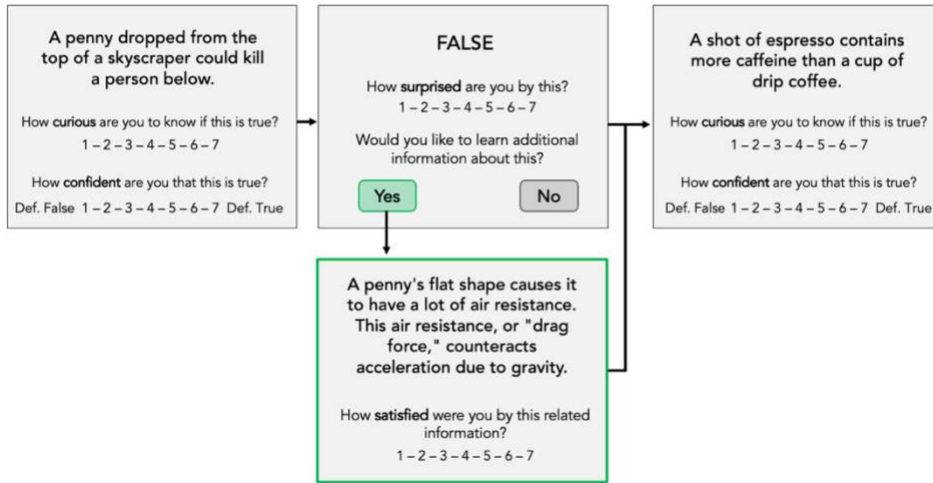
5.2.2 Task Design and Procedure

To explore the influences of curiosity and prediction error on information seeking and memory updating, we designed a task where participants learned about the veracity of different trivial facts and urban myths. On each trial, participants were presented with a statement, and they reported on how curious they were to know if the statement was true as well as how confident they were that the statement was either true or false on a Likert scale of 1 – 7 (Figure 1a). For curiosity, only the lower and upper bounds of the Likert scale were labeled (e.g., Not at All Curious, Very Curious). For the confidence rating, the scale included end labels for “Highly Confident False” and “Highly Confident True”, as well as a center point label for “Unsure”. This was done to give participants an option to report no expectations about the veracity of the current statement. After providing these ratings, participants advanced to the next screen where they saw the answer to the trivia question, either true or false and rated their surprise at finding out the truth of the statement on a Likert scale of 1 – 7. Participants were then given the option to read more information about the statement. The choice to seek additional information led participants to a new screen where they viewed a short (one-to-two sentence) piece of related information. Participants rated their satisfaction (Likert scale of 1 – 7) with the additional information before moving on to the next trial. If participants did not choose to seek additional information, they simply advanced to the next trial. In the Learning Phase of the study, participants experienced 75 statements and had 75

chances to seek additional information. After completion of the Learning Phase, we asked participants to provide demographic ratings and complete a few questionnaires as filler tasks.

Participants then completed the Memory Test Phase either immediately or after a 24-hour delay. During the Memory Test Phase, participants saw the same 75 statements as in the Learning Phase and reported on whether the statement was True or False on a Likert scale of 1 – 5, which included end labels for “Highly Confident False” and “Highly Confident True”, as well as a center point label “Unsure”. (Figure 1b). They then recalled if they sought additional information on that statement, and if so, they were asked to select which of three statements they remembered seeing in the Learning Phase. All three statements used in the Memory Test Phase were factual and related to the original statement. The entire task was self-paced, but participants could not return to previous trial questions or answers. For our Immediate Memory group, the entire task, both Learning and Memory Test Phases, took on average $M = 51.8$ minutes ($SD = 26.1$ minutes) to complete. For our 24-Hour Delay Memory group, the Learning Phase took an average of $M = 34.1$ minutes ($SD = 17.9$ minutes) to complete, and the Memory Test Phase took an average of $M = 15.7$ minutes ($SD = 12.4$ minutes) to complete.

a. Session 1: Learning Phase



b. Session 2: Memory Test Phase

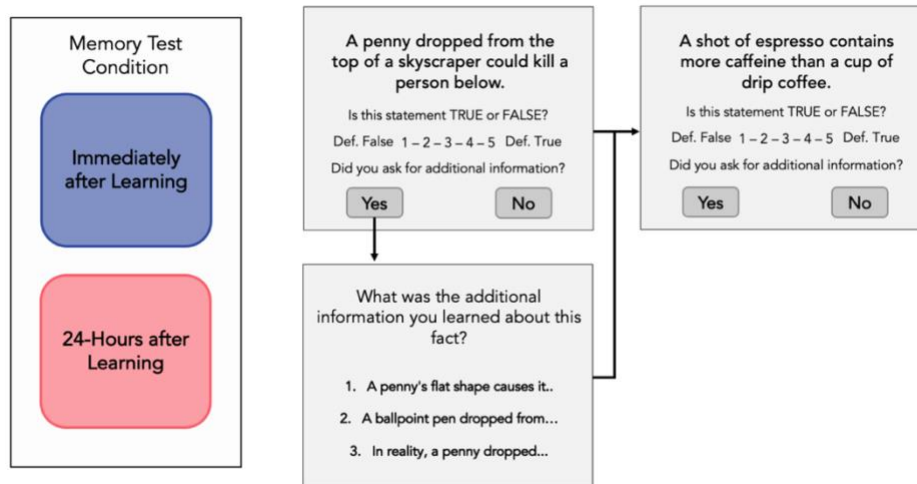


Figure 17: Task design for testing the influence of curiosity and prediction error on information seeking and memory updating.

5.2.3 Data Analysis

All analyses were conducted with R v4.0.3 and RStudio v1.3.1093. Confidence ratings during the Learning Phase were first collapsed into a 4-point Likert scale such that “Highly Confident False” and “Highly Confident True” now both represented a “High Confidence” rating of three. The lower bound, providing an “Unsure” response, was

coded as a zero. We then used these collapsed ratings to compute a prediction error term for each trivia statement for each participant using their confidence ratings and their accuracy in assessing the veracity of the trivia statement. For statements that were correctly judged, the prediction error term was computed as $4 - \text{confidence rating}$, such that high confident correct judgements were coded as low positive prediction errors and low confident correct judgements were coded as high positive prediction errors (i.e., I am surprised that I was correct!). We treated confidence ratings of 0 (“Unsure”) as incorrect judgements with no prediction error ($PE = 0$) since no expectation about the veracity of the statement was provided. For statements that participants judged incorrectly, prediction error was computed as the additive inverse of confidence. To assess successful memory updating or formation, we compared the judged veracity of the statement from Session 2 to the answer provided in Session 1. Responses of “Unsure” during Session 2 were marked as unsuccessful memory formation or updating.

We created mixed effects logistic regression models with the lme4 package (Bates et al., 2015) and bound optimization by quadratic approximation (BOBYQA controller). To explore the determinants of seeking additional information, we ran two models, one for correctly judged statements and another for incorrectly judged statements. We split our models by accuracy to understand whether curiosity and prediction error changed information seeking behavior as a function of receiving confirmatory or contrary information. Each model included curiosity and prediction error and their interaction as well as Memory Test condition (Immediate vs. 24-hour Delay). To understand memory updating, we ran a single model that contained only statements that participants judged incorrectly during the Learning Phase. Within this model, we again included curiosity, prediction error, and their interaction as well as the choice to

learn additional information and Memory Test condition. We added all two-way interactions between each predictor and Memory Test condition to test for delay-dependent effects of uncertainty and prediction error, which may be suggestive of effects on memory consolidation. Memory accuracy for the decision to seek additional information as well as recognition memory for the extra information seen were close to ceiling (both 94%) and were not analyzed further. Similarly, the extra information participants could examine was overall rated as highly satisfying ($M = 5.8$ out of 7, $SD = 1.58$) and thus not included in the models. In all cases curiosity and prediction error were treated as continuous variables and centered and scaled.

We obtained p-values with the lmerTest package (Kuznetsova & Brockhoff, 2017). All mixed effects models included random effects with both participant and statement as random intercepts. Each of our models converged successfully. For all statistical analyses, we report standardized β values (and 95% confidence intervals (CIs) around the slope estimate). Figures were produced using the raw data in Python with Seaborn and Matplotlib.

5.3 Results

5.3.1 Seeking Additional Information

We first examined the choice to seek additional information when participants had incorrectly judged the veracity of the trivia statement. We found that higher ratings of curiosity predicted a higher probability of seeking related information ($\beta = 1.36$, 95% CI [1.22, 1.51], $z = 18.08$, $p < 0.001$, Figure 2a), suggesting that curiosity can extend beyond the answer to a question to motivate the search for extra information as well. Additionally, the stronger the negative prediction error, the higher the probability of

seeking more information ($\beta = -1.37$, 95% CI [-1.49, -1.25], $z = -21.74$, $p < 0.001$, Figure 2a), confirming prior reports of prediction error driving information seeking (Vogl, Pekrun, Murayama, Loderer, et al., 2019). Curiosity and prediction error also showed a significant interaction ($\beta = 0.53$, 95% CI [0.42, 0.64], $z = 9.64$, $p < 0.001$, Figure 2a), such that higher ratings of curiosity uniformly predict seeking additional information whereas prediction error drove seeking choices when curiosity was low.

We then investigated the choice to seek additional information when participants correctly judged the veracity of the trivia statements. Here, we found that, again, curiosity was a significant positive predictor of seeking information ($\beta = 1.00$, 95% CI [0.77, 1.24], $z = 8.29$, $p < 0.001$, Figure 3b). Prediction error was also a significant predictor ($\beta = 0.65$, 95% CI [0.40, 0.90], $z = 5.16$, $p < 0.001$, Figure 3b), indicating that information seeking was increased for statements where a correct judgement was made with low confidence. We also found an interaction between curiosity and prediction error ($\beta = 0.41$, 95% CI [0.17, 0.66], $z = 3.32$, $p < 0.001$, Figure 3b). When statements were correctly judged, curiosity and prediction error appear to work synergistically such that experiencing a correct judgement with low confidence and being curious about the answer at the time of first seeing the statement create an even greater inclination toward seeking related information.

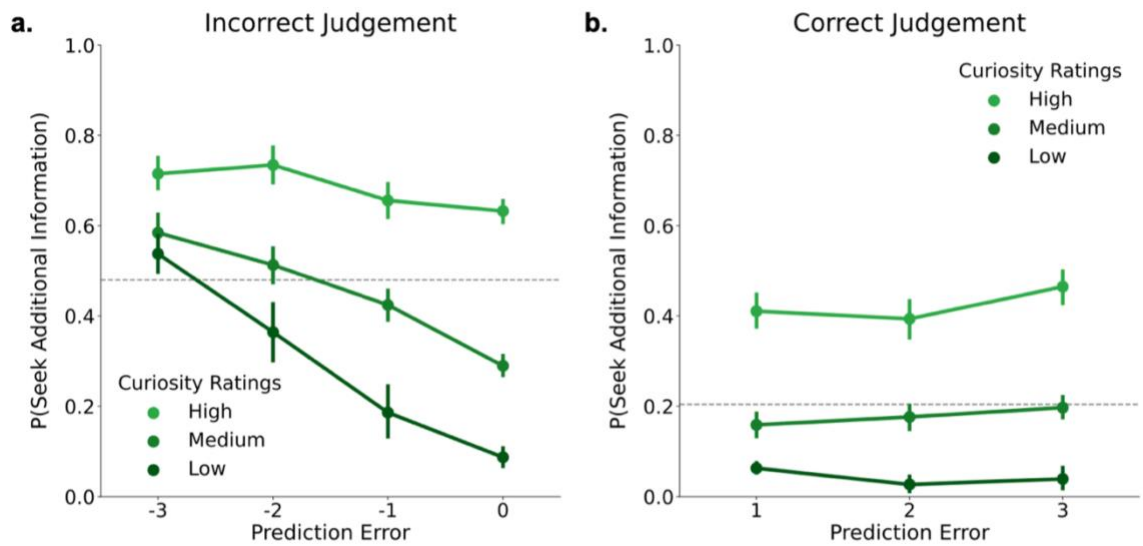


Figure 18: Curiosity and prediction error offer complementary motivations to seek additional information.

5.3.2 Memory Updating

We next examined how curiosity, prediction error, and the choice to seek additional information are related to learning. Specifically, we were interested in memory updating, whether participants would update a belief about a statement after they wrongly judged the veracity in the first phase. To this end, we focused our analyses on statements that were incorrectly judged during the Learning Phase across our two Memory Test conditions (Immediate versus Delayed).

Consistent with prior work, we showed that curiosity and prediction error were both significant predictors of memory updating (Kang et al., 2009; Marvin & Shohamy, 2016; Poh et al., 2022; Rouhani et al., 2018; Vogl, Pekrun, Murayama, Loderer, et al., 2019). Higher curiosity ($\beta = 0.45$, 95% CI [0.30, 0.61], $z = 5.69$, $p < 0.001$, Figure 3) and larger prediction error ($\beta = -0.20$, 95% CI [-0.35, -0.05], $z = -2.54$, $p = 0.011$, Figure 3) were both associated with a greater likelihood of correcting previously erroneous

answers. There was also a significant interaction between curiosity and prediction error whereby high states of curiosity enhanced memory across the magnitude of experienced prediction error while the effect of prediction error was greater when initial curiosity was low ($\beta = 0.20$, 95% CI [0.10, 0.31], $z = 3.80$, $p < 0.001$, Figure 3). Lastly, there were no differences in the effects for curiosity and prediction error between the Immediate Memory Group and the Delayed Memory Group. These findings suggest that curiosity and prediction error provide complementary paths towards driving memory updating and these paths may be relatively stable across a delay.

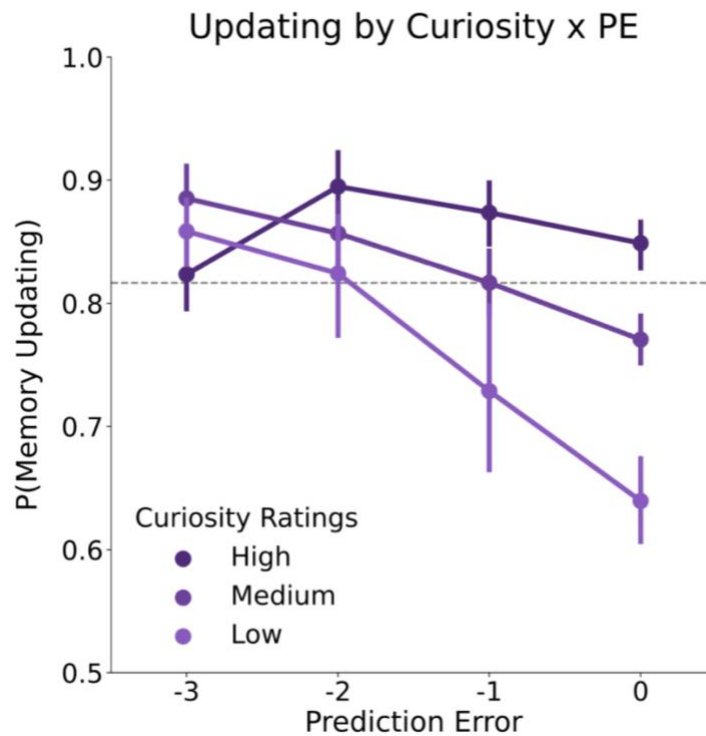


Figure 19: Curiosity and Prediction error independently enhance memory updating.

Central to our hypothesis, the choice to seek additional information was also a significant predictor of subsequent memory updating ($\beta = 1.75$, 95% CI [1.48, 2.02], $z =$

12.82, $p < 0.001$, Figure 4), whereby participants were more likely to update erroneous memory following the choice to seek additional information. Contrasting to the effects of curiosity and prediction error, we also observed an interaction with the Memory Test condition ($\beta = -0.51$, 95% CI [-0.91, -0.12], $z = -2.55$, $p = 0.011$, Figure 4), such that the effects of information seeking was more pronounced following a 24h-delay (OR = 1.67, $z = 2.55$, $p = 0.033$). In addition, we found that memory performance overall was better in the Immediate Memory Group than in the 24-Hour Delay group ($\beta = 0.61$, 95% CI [0.15, 1.08], $z = 2.59$, $p = 0.010$, Figure 4). These findings suggest that seeking additional information can have protective effects against memory degradation, particularly for longer delays.

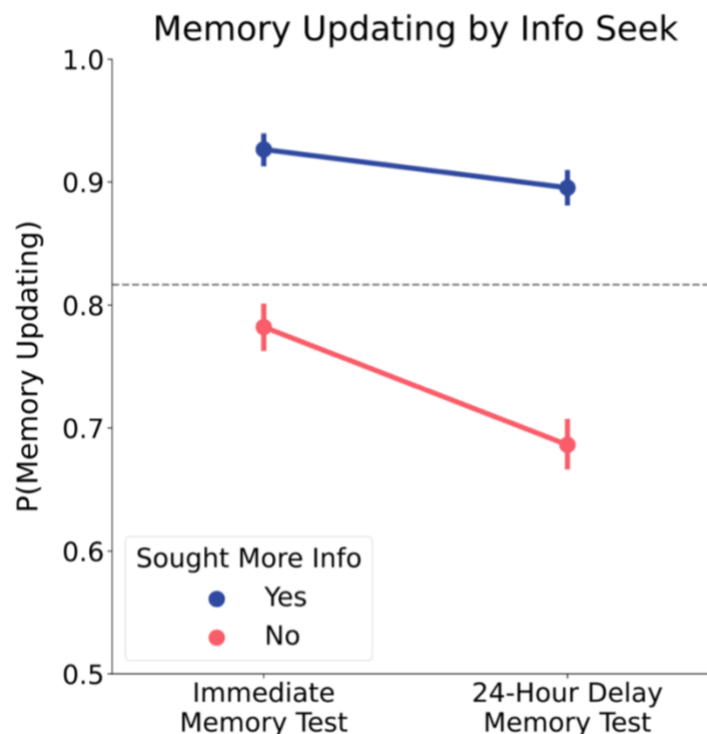


Figure 20: Seeking additional information boosts memory updating, particularly for information tested after a delay.

5.4 Discussion

We investigated the roles of curiosity and prediction error in a trivia paradigm that explored both information seeking for related information and memory updating. We found that curiosity and prediction error both contribute to continued information seeking through distinct mechanisms. High curious states drive subsequent information seeking across the range of experienced prediction errors while prediction error drives subsequent information seeking most strongly when a statement elicits low curiosity and when an original assessment is incorrect. When we examined learning, we found that choosing to seek more information benefits memory performance for the original statement. This effect is largest when memory was tested 24 hours later as opposed to immediately, suggesting that seeking additional information can support memory via delay-dependent mechanisms. Lastly, we saw that curiosity and prediction error provide complementary paths towards driving memory updating, where curiosity provides a general improvement to memory, while prediction error specifically drives memory updating for statements that elicited low curiosity when first encountered. Collectively, these findings demonstrate that information seeking, and memory formation can be uniquely driven by curiosity and prediction error, signaling a balance between new learning, and correcting previously misinformed beliefs.

Our findings support and extend the extant literature on curiosity, prediction error and information seeking. Similar to prior work, we found that curiosity drives information seeking (Baranes et al., 2015; Kang et al., 2009; Lang et al., 2021; Marvin & Shohamy, 2016; Metcalfe et al., 2021; Vogl, Pekrun, Murayama, & Loderer, 2019; Vogl, Pekrun, Murayama, Loderer, et al., 2019). We extend these reports here by demonstrating that curiosity experienced at the first encounter of a statement can predict subsequent

seeking for related information even after the resolution to the statement has been received. The continuation of information seeking after resolution offers evidence for the self-perpetuating nature of curiosity (Gruber & Ranganath, 2019; Murayama, 2022). Diverging from traditional reward motivation, curiosity-guided information-seeking can sustain itself as new information offer new questions without the need for other external reinforcers to drive behavior. We also found evidence for prediction error in shaping information seeking, particularly when high-confidence errors are experienced. Surprise has previously been shown to increase related information seeking through the mediating factor of curiosity (Vogl, Pekrun, Murayama, & Loderer, 2019) but here we show that curiosity and prediction error can independently drive information seeking.

We also demonstrate that curiosity can enhance memory formation, corroborating past work (Gruber et al., 2014; Kang et al., 2009; Lang et al., 2021; Marvin & Shohamy, 2016; Janet Metcalfe et al., 2021; Mullaney et al., 2014; Poh et al., 2022). One possibility is that curiosity drove learning through participants having substantial prior knowledge that aided in successful encoding (Wade & Kidd, 2019). We think this is unlikely, given curiosity was highest for statements that participants were least confident about, reflecting less prior knowledge. We also found evidence for prediction error in shaping information seeking and memory updating, supporting previous research (Kang et al., 2009; Pine et al., 2018; Potts & Shanks, 2014; Rouhani et al., 2018; Sinclair & Barense, 2018). One important distinction to contextualize our findings is understanding how prediction error is defined. Here, we computed prediction error as the signed difference between a learner's assessment of their knowledge and the actual veracity of a statement, reflecting how correct or incorrect their prior knowledge state was. In other work, "Informational Prediction Errors" (IPEs) are computed as the difference between

the anticipated satisfaction of an answer (measured through curiosity) minus the realized satisfaction of that answer (Abir et al., 2022; Lang et al., 2021; Marvin & Shohamy, 2016). IPEs also show memory enhancements and suggest that throughout the information seeking process there are multiple opportunities where one can experience a prediction error (e.g., in the assessment of whether or not a fact is known or in how satisfying an answer might be). Teasing apart these different prediction errors and their role in information seeking and memory updating reflect an important next step for future research.

Similar to the effects of information seeking, we found that curiosity benefitted memory across prediction errors while prediction error selectively enhanced memory for statements that elicited low curiosity. Again, these distinct roles highlight unique motivational conditions for learning: one that prioritizes learning new information and another to correct misinformed beliefs. Having two systems sensitive to different types of information is likely adaptive and rational to balance our information seeking and later learning (Dubey & Griffiths, 2020). Although we did not collect imaging data for this study, one open question points to whether learning from curiosity and prediction error arise from different underlying neural mechanisms. Both prediction error and curiosity have been shown to relate to activation in the hippocampus (Gruber et al., 2014; Jepma et al., 2012; Poh et al., 2022; Sinclair et al., 2021) and reward circuitry including the VTA/SN, striatum, and basal forebrain (Charpentier et al., 2018; Gruber et al., 2014; Iglesias et al., 2013; Kang et al., 2009; Pine et al., 2018; Poh et al., 2022; Sinclair et al., 2021). One possibility is that although these constructs rely on similar neural mechanisms, the timing and activation patterns could signal meaningful differences that characterize how memory is enhanced. For example, curiosity can engender an

anticipatory state through increasing dopaminergic midbrain activation and modulating the stability of the hippocampus, readying it for upcoming information (Poh et al., 2022). In contrast, prediction errors might disrupt the stability of hippocampal patterns through cholinergic modulation, creating an opportunity for prior beliefs to update (Sinclair et al., 2021). Future work is needed to directly compare roles of curiosity and prediction error in enhancing learning.

Lastly, we found that seeking related information helps with successful memory encoding and updating for the original statement. One potential mechanism underlying this enhancement could be an increase in attentional processing and time spent thinking about and integrating the learned information (Craik & Tulving, 1975). Another potential avenue could come from having the control to choose or forgo related information. Prior work has documented that agency over choices can increase memory for chosen information (Murty et al., 2015, 2019; Rotem-Turchinski et al., 2018), however in our task we did not show information to individuals that did not explicitly choose to see it. Additional work has highlighted the role of agency in titrating a learner's expectations and assessment of their own knowledge state with better memory (Markant et al., 2016; Voss et al., 2011), providing evidence for self-regulated search as an important mechanism for successful learning. Interestingly, the effects of additional information seeking were particularly prominent following a 24-hour delay, suggestive of potential consolidation-dependent mechanisms. This observation is reminiscent of prior work showing that memory can be strengthened via post-learning exposure to novelty (Moncada & Viola, 2007; Salvetti et al., 2014). As such, while the act of seeking additional information may itself be sufficient for enhancing memory for previously learned information, the novelty of the additional information may also be a contributing

factor. Dissociating the roles of depth of processing, agency, and novelty as mechanisms driving memory enhancement from information seeking will help to uncover how and why self-guided engagement shapes our memory.

A critical question moving forward is to understand what the time course of curiosity is across information gathering. Prior work investigating the shape of curiosity over extended information gathering has demonstrated that curiosity initially grows with information seeking but then declines as the collected information reaches a critical point (Fastrich & Murayama, 2020). These dynamics appear to be sensitive to how much perceived information has been traversed as well as perceptions of how much more information can be learned (Wade & Kidd, 2020). Yet, future work is needed to characterize what defines changepoint between curiosity and information gathered and the role of prediction error in sustaining search. Additionally, our findings here also stand in contrast with theories of curiosity as concluding once an informational gap has been closed (Loewenstein, 1994). One factor that might account for whether curiosity leads to the extension or termination of information seeking, is how complete the closure of an informational gap feels to a learner. For example, closing one gap may open another or lead to a surprising discovery, leading to more questions and sustaining information seeking (Murayama, 2022; Vogl, Pekrun, Murayama, & Loderer, 2019; Vogl, Pekrun, Murayama, Loderer, et al., 2019). Conversely, if an answer feels particularly satisfying (Gruber & Ranganath, 2019), it might signal the end of a learning experience and thus end information gathering (E. G. Liquin & Lombrozo, 2021).

Despite the advances this study provides, there are limitations to how well it can be generalized. First, we tested these relationships using trivia questions that might inherently be more engaging and elicit more curiosity than other topical information.

Future work can expand this work by including informational mediums of varying interest and prior knowledge. Second, we limited the amount of extra information people could seek out to one additional fact. In order to more deeply distinguish the impact of prediction error and curiosity on information seeking, the time course of how long information seeking is perpetuated should be considered. Moreover, the extra information participants could examine was overall highly satisfying. This lack of variability dampened our ability to understand whether the memory effects from seeking information were influenced by how satisfying the related information was judged to be. Future work should aim to manipulate the satisfaction of related information to investigate the role of closure in aiding memory formation and updating.

Collectively, we demonstrate that information seeking and learning can be shaped by two distinct processes: the subjective state of curiosity and the experience of a prediction error. Additionally, we show that the act of seeking additional information can improve memory for the original statement, suggesting that our motivational states as well as our actions within our environment have consequences for how information is remembered. Understanding how information seeking and memory formation are driven through these distinct processes has implications for how educators encourage learning in students as well as for shaping interventions to target misinformation. Ultimately through our work here, we propose that different motivational states can prioritize unique goals during information seeking, producing multiple avenues to grow our knowledge.

6. Conclusion

Understanding how we organize our information gathering processes is a challenge that we constantly face as we navigate a world with infinite information. Across the collection of experiments presented, I find that the internal motivations of an individual help to shape sequential information seeking and have consequences for memory formation.

6.1 Study Summaries

In Study 1, I examined how people implement strategies to guide their information sampling. To characterize the computations that underlay information sampling, I studied the choice processes in people deployed as they sampled sequences of images (e.g. indoor and outdoor scenes) in an attempt to infer the majority category (e.g. indoor or outdoor) under two reward conditions. I discovered that choices were best described by a flexible function that lay between optimal and heuristic strategies; one that integrated the magnitude of evidence favoring each category and the number of samples collected thus far. Integration of these criteria resulted in a trade-off between evidence and samples collected, in which the strength of evidence needed to stop sampling decreased linearly as the number of samples accumulated over the course of a trial. This non-optimal trade-off best accounted for choice behavior even under high reward contexts. These results demonstrated that unlike the optimal strategy, people perform simple accumulations instead of computing expected values, and that unlike a simple heuristic strategy, people dynamically integrate multiple sources of information in lieu of using only one source. This study highlighted an evidence-by-costs tradeoff that illustrates a computationally

efficient strategy that balances competing motivations for accuracy and cost minimization.

In Study 2, I explored whether states of high curiosity favored delaying resolution when information unfolds across time. I developed a novel video task in which line drawings slowly resolved into objects, and participants made choices to keep watching as the video progressed. I examined the role of curiosity in driving information seeking choices across two studies that manipulated the amount of agency people were given over the timing of their choices. In Study 1, the timing of these choices was predetermined and externally imposed while in Study 2, participants had agency over exactly when to make these choices. Across both studies, I showed that when curious, participants often chose to prolong uncertainty as information is gradually revealed, rather than alleviate uncertainty immediately (i.e., viewing a spoiler). The affective experience of curiosity was positive, associated not with frustration but rather tension and enjoyment. Additionally, I found that when participants were given control over the timing of their choices (Study 2), they specifically waited until they knew the identity of the drawing before stopping. This finding indicated that the value of seeking information was tied to reaching resolution, rather than another value, such as aesthetic experience. Collectively, these findings demonstrate that curiosity not only confers value to information itself, but also influences preferences for how one arrives at that information, highlighting the reward value inherent in the process of discovery.

Study 3 provided insight into the neural mechanisms that support the experience of resolution. Using the same dynamic stimuli as in Study 2, I explored the difference phases people experience across unfolding information and how such experiences shaped later memory. Within the fMRI scanner, people watched videos of single-line

drawings evolve while indicating key components of their thought process along the way. I manipulated the videos people saw such that some videos resolved into interpretable objects while other videos never resolved, remaining scribbles the entire way through. 24-hours later, people watched the same videos again and indicated throughout video play whether and when they remembered the video from the day before. Through the examination of brain activity, I found that the information gathering process could be divided into two distinct phases: (1) an exploration phase where participants generated ideas about the identity of the line drawing that was associated with regions involved in semantic memory retrieval, including the posterior cingulate cortex, the precuneus, and the hippocampus, and (2) a hypothesis testing phase where evidence towards various identities accumulated until realization occurred, supported by regions involved in reward, memory, and object recognition, including the VTA, striatum, parahippocampal cortex, and occipital temporal cortex. Importantly, the neural correlates supporting the hypothesis testing phase were dependent on videos resolving into objects and were not present when line drawings remained abstract. Finally, during the process of discovery, increased activation in the hippocampus and midbrain predicted successful memory formation. Together, these findings illustrate the dynamic interplay between perceptual information and internal representations as information unfolds towards discovery, providing insight into the mechanisms of why discovery is a rewarding and adaptive experience.

In Study 4, I investigated the role of curiosity and surprise in driving information seeking choices as well as memory updating. This work sought to understand if and how curiosity would self-perpetuate information seeking after resolution was received. To accomplish this, I had participants navigate through a trivia task where they judged the

veracity of trivia facts and urban legends. After uncovering the answers, they had the option to request additional relevant information. I found that while both curiosity and prediction error contributed to information seeking and subsequent memory, they exhibited dissociable influences depending on the initial state of uncertainty. Curiosity enhanced information seeking and memory when participants experienced high uncertainty. In contrast, prediction error enhanced information seeking and memory when high certainty was found to be erroneous. Additionally, I uncovered that seeking additional information also improved memory for the original trivia and this benefit was most pronounced following a delay. Together, these results point to instances where information seeking and memory formation are differentially driven by curiosity and prediction error, and that the choice to seek related information can provide an additional boost to learning. Lastly, this work demonstrated that curiosity and prediction error both work to self-perpetuate information seeking without the need for other external reinforcers to drive behavior, providing key insight into how we self-direct our learning.

These studies, together, highlight how the process of information gathering impacts the formation of judgements, the valuation of uncertainty and discovery, and how we learn and update our knowledge base. Central to this process is the motivational context. By examining these processes through a variety of paradigms and targeting different motivational states, we demonstrate that motivational states direct how information gathering is enacted and can shape what we are able to glean from these experiences.

6.2 Future Directions

Across Study 1 and Study 2, I explored how much information people gathered before they terminated their search across two different motivational contexts: in pursuit of reward and in pursuit of resolution. Both findings pointed to individualized thresholds that set bounds how much information was enough to satisfy the needs of the individual. In Study 1, this threshold was determined by both the magnitude of evidence collected and the total amount of information gathered, reflecting a balance between accuracy and accumulation costs. In Study 2, this threshold was determined by resolution, the moment a person understood the identity of the line drawing. Critically, both qualities rely on the metacognitive perception of confidence. Although past research has thoroughly investigate the role of confidence in choice behavior across both instrumental (Desender et al., 2018; Hausmann-Thürig & Läge, 2008; Kiani et al., 2014; Pouget et al., 2016; Schulz et al., 2023) and non-instrumental information seeking (Baranes et al., 2015; Cohanpour et al., 2022; Dubey & Griffiths, 2020a; Fastrich et al., 2018; Käckenmester et al., 2018; Kang et al., 2009), there remain gaps in our understanding of the computation and stability of confidence. For example, how do people establish a level of confidence for their search process, how malleable is this threshold across time, and how does this threshold change as a function of the search context (i.e., instrumental versus non-instrumental).

When information gathering extends over time it is often rare that each piece of information is weighted the same or contributes termination decisions in the same manner. Prior work has documented that decision-makers are sensitive to the reliability of individual pieces of information (Kira et al., 2015) but how individuals learn such reliability as well as how this translates to non-instrumental information seeking contexts

remains unknown. Interestingly, the most revealing pieces of information might not always be sought after. In Study 2, when information about drawings unfolded over time people often avoided the option to obtain the most illuminating information, instead preferring to remain in suspense. These findings suggest that preference for reliable pieces of information may be different depending on the motivational context. Research into this challenge is also critical for understanding how humans fall into misinformation and how best to signal whether information encountered should be trusted or not (Majerczak & Strzelecki, 2022; Muhammed T & Mathew, 2022; van der Linden, 2022).

In the studies above, participants never needed to search for where to find information, but knowing where to look is fundamental to adaptive and continued search (Gureckis & Markant, 2012; Voss et al., 2011). Prior work has reported that deciding where to look for relevant information is taxing and effortful, even in instances where the search space is limited (Petitet et al., 2021; Schwartenbeck et al., 2019; Vidal et al., 2022; Wu et al., 2018). Expanding research into how individuals select and change amongst different “sources” of information as well as how this is impacted by motivational states is critical for the understanding information seeking behaviors in the real-world (Wu et al., 2017).

In Studies 3 and 4, I examined how the process of seeking information over time shapes memory formation. In Study 3, I found enhanced encoding came from experiencing discovery as well as increased activation from hippocampal and midbrain areas. Experiencing resolution has previously been shown to enhance memory and elicit dopaminergic and hippocampal areas, particularly for information that evokes curiosity (Gruber et al., 2014; Jepma et al., 2012; Kang et al., 2009). However, the current study provided novel contributions to understanding how information presented dynamically

over time can also shape memory formation. Expanding on this work, future research should aim to uncover how different paths to resolution impact memory formation and the neural correlates that support encoding. One example of this could be to test whether information revealed over time is differentially remembered compared to information revealed all at once. Another advancement would be to test if the duration of suspense shapes memory formation in a linear or non-linear pattern. For example, it is possible that the longer a person stays in suspense the better their memory for the events will be (Bezdek et al., 2017). Conversely, suspense might influence memory in an inverted U-shaped pattern where too little time in suspense might degrade memory due to boredom (Goldberg & Todman, 2018; van Hooft & van Hooft, 2018), and too long in suspense could produce selective memory for the outcome at the expense of the process (Mather & Sutherland, 2011). These findings could be used to help tailor educational practices to enhance learning.

In Study 4, I showed that seeking additional information after resolution was also predictive of later successful memory updating. Here, we gave participants just a single choice to seek information, but future work will need to expand upon this to understand how long curiosity can perpetuate search and how it relates to the qualities of traversed information. Prior work has begun to explore the rise and fall of curiosity (Fastrich & Murayama, 2020) but more work is needed to explore how curiosity might inspire new topic exploration and what precedes these topic shifts. Other work has shown that giving people the freedom to explore topics was influenced by a person's trait curiosity (Lydon-Staley et al., 2021), but this study did not explore the interplay between how encountered information directs subsequent search behaviors. Overall, understanding the motivational state a person enters information seeking with as well as the nuanced

dynamic between information, affect, and curiosity remain key areas of interest for the field.

Broadly, understanding the role of motivation in directing sequential search behaviors has relevance for both education and computer science. The findings here and future research along this path can inform our understanding of educational practices and help build interventions toward engendering adaptive search strategies for knowledge acquisition (Halamish et al., 2019; Oudeyer et al., 2016; Wimpenny & Clarke, 2017). Additionally, future research can work to establish how the relationship between curiosity and sequential information seeking can guide and take guidance from computational models of curiosity and autonomous learning (Cangelosi & Schlesinger, 2015; Friston et al., 2017; Kaplan & Oudeyer, 2007). Ultimately, future work will advance our understanding of adaptive information search and self-directed learning, fundamental qualities of the human experience.

6.3 Conclusion

From adaptive decision-making to helping guide scientific progress, information seeking is a critical component to how we engage and learn from the world around us. The work in this dissertation aimed to advance the study of information seeking through examining how different motivational states guide this process. I organized my work into two broad aims. In the first aim, I investigated how motivational states shape the development of criteria to terminate information seeking. Towards this end, I found that during instrumental information seeking, people implement simple yet flexible strategies to balance the competing motivations of the benefits information provides with the costs of obtaining it (**Study 1**). Additionally, within non-instrumental information seeking,

curiosity increases the value of delaying resolution, and the self-perceived experience of discovery decreases this value (**Study 2**). My second aim explored how the process of information seeking impacted learning and memory. Within this arm of work, I found that the experience of resolution enhances memory, and this corresponded to increased activation in the hippocampus and midbrain areas (**Study 3**). I also demonstrated that both curiosity, prediction error, and seeking related information can all improve memory updating through complementary mechanisms (**Study 4**). Collectively, this work provides insights into how sequential information seeking is shaped by different motivational states, while also highlighting key research areas to continue advancing our understanding of why and how we direct our own learning.

References

- Abir, Y., Marvin, C. B., van Geen, C., Leshkowitz, M., Hassin, R. R., & Shohamy, D. (2022). An energizing role for motivation in information-seeking during the early phase of the COVID-19 pandemic. *Nature Communications*, *13*(1), 2310.
- Achtziger, A., Alós-Ferrer, C., Hügelschäfer, S., & Steinhauser, M. (2015). Higher incentives can impair performance: neural evidence on reinforcement and rationality. *Social Cognitive and Affective Neuroscience*, *10*(11), 1477–1483.
- Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: mesolimbic activation precedes memory formation. *Neuron*, *50*(3), 507–517.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*(6), 716–723.
- Antony, J. W., Hartshorne, T. H., Pomeroy, K., Gureckis, T. M., Hasson, U., McDougle, S. D., & Norman, K. A. (2021). Behavioral, Physiological, and Neural Signatures of Surprise during Naturalistic Sports Viewing. *Neuron*, *109*(2), 377-390.e7.
- Arrow, K. J., Blackwell, D., & Girshick, M. A. (1949). Bayes and Minimax Solutions of Sequential Decision Problems. *Econometrica: Journal of the Econometric Society*, *17*(3/4), 213–244.
- Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*, 403–450.
- Avants, B. B., Epstein, C. L., Grossman, M., & Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, *12*(1), 26–41.
- Averbeck, B. B. (2015). Theory of choice in bandit, information sampling and foraging tasks. *PLoS Computational Biology*, *11*(3), e1004164.
- Aziz-Zadeh, L., Kaplan, J. T., & Iacoboni, M. (2009). “Aha!”: The neural correlates of verbal insight solutions. *Human Brain Mapping*, *30*(3), 908–916.
- Bach, D. R., & Dolan, R. J. (2012). Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature Reviews. Neuroscience*, *13*(8), 572–586.
- Baranes, A., Oudeyer, P.-Y., & Gottlieb, J. (2015). Eye movements reveal epistemic curiosity in human observers. *Vision Research*, *117*, 81–90.

- Barense, M. D., Bussey, T. J., Lee, A. C. H., Rogers, T. T., Davies, R. R., Saksida, L. M., Murray, E. A., & Graham, K. S. (2005). Functional specialization in the human medial temporal lobe. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(44), 10239–10246.
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, *76*, 412–427.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting Linear Mixed-Effects Models using lme4. In *arXiv [stat.CO]*. arXiv. <http://arxiv.org/abs/1406.5823>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using **lme4**. In *Journal of Statistical Software* (Vol. 67, Issue 1). <https://doi.org/10.18637/jss.v067.i01>
- Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(23), 12750–12755.
- Becker, M., Sommer, T., & Kühn, S. (2020). Verbal insight revisited: fMRI evidence for early processing in bilateral insulae for solutions with AHA! experience shortly after trial onset. *Human Brain Mapping*, *41*(1), 30–45.
- Behzadi, Y., Restom, K., Liou, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, *37*(1), 90–101.
- Bellman, R. (1957). *Dynamic programming*. Princeton University Press.
- Bennett, D., Bode, S., Brydevall, M., Warren, H., & Murawski, C. (2016). Intrinsic Valuation of Information in Decision Making under Uncertainty. *PLoS Computational Biology*, *12*(7), e1005020.
- Bennett, D., Sasmita, K., Maloney, R. T., Murawski, C., & Bode, S. (2019). Monetary feedback modulates performance and electrophysiological indices of belief updating in reward learning. *Psychophysiology*, *56*(10), e13431.
- Berlyne, D. (08/1954). A Theory of Human Curiosity. *British Journal of Psychology*, *45*(3). <https://doi.org/10.1111/j.2044-8295.1954.tb01243.x>
- Berlyne, D. E. (1962). Uncertainty and epistemic curiosity. *British Journal of Psychology*, *53*, 27–34.
- Berlyne, D. E. (1966). Curiosity and exploration. *Science*, *153*(3731), 25–33.

- Bezdek, M. A., Wenzel, W. G., & Schumacher, E. H. (2017). The effect of visual and musical suspense on brain activation and memory during naturalistic viewing. *Biological Psychology*, *129*, 73–81.
- Bilalić, M., Graf, M., Vaci, N., & Danek, A. H. (2021). The temporal dynamics of insight problem solving – restructuring might not always be sudden. *Thinking & Reasoning*, *27*(1), 1–37.
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, *15*(11), 527–536.
- Blanchard, T. C., & Gershman, S. J. (2018). Pure correlates of exploration and exploitation in the human brain. *Cognitive, Affective & Behavioral Neuroscience*, *18*(1), 117–126.
- Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, *85*(3), 602–614.
- Bonini, F., Burle, B., Liégeois-Chauvel, C., Régis, J., Chauvel, P., & Vidal, F. (2014). Action monitoring and medial frontal cortex: leading role of supplementary motor area. *Science*, *343*(6173), 888–891.
- Bossaerts, P., & Murawski, C. (2017). Computational Complexity and Human Decision-Making. *Trends in Cognitive Sciences*, *21*(12), 917–929.
- Bossaerts, P., Yadav, N., & Murawski, C. (2019). Uncertainty and computational complexity. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *374*(1766), 20180138.
- Boyd-Graber, J., Glasgow, K., & Zajac, J. S. (2013). Spoiler alert: Machine learning approaches to detect social media posts with revelatory information. *Proceedings of the American Society for Information Science and Technology*, *50*(1), 1–9.
- Braver, T. S., Krug, M. K., Chiew, K. S., Kool, W., Westbrook, J. A., Clement, N. J., Adcock, R. A., Barch, D. M., Botvinick, M. M., Carver, C. S., Cools, R., Custers, R., Dickinson, A., Dweck, C. S., Fishbach, A., Gollwitzer, P. M., Hess, T. M., Isaacowitz, D. M., Mather, M., ... MOMCAI group. (2014). Mechanisms of motivation-cognition interaction: challenges and opportunities. *Cognitive, Affective & Behavioral Neuroscience*, *14*(2), 443–472.
- Brockbank, E., Lombrozo, T., Gopnik, A., & Walker, C. M. (2022). Ask me why, don't tell me why: Asking children for explanations facilitates relational thinking. *Developmental Science*, e13274.
- Brydevall, M., Bennett, D., Murawski, C., & Bode, S. (2018). The neural encoding of information prediction errors during non-instrumental information seeking. In *Sci. Rep.* (Issue 1, p. 179507). <https://doi.org/10.1038/s41598-018-24566-x>

- Cangelosi, A., & Schlesinger, M. (2015). *Developmental Robotics: From Babies to Robots*. The MIT Press.
- Cassey, P., Hawkins, G. E., Donkin, C., & Brown, S. D. (2016). Using alien coins to test whether simple inference is Bayesian. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 42(3), 497–503.
- Charness, G., & Levin, D. (2005). When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect. *The American Economic Review*, 95(4), 1300–1309.
- Charpentier, C. J., Bromberg-Martin, E. S., & Sharot, T. (2018). Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proceedings of the National Academy of Sciences of the United States of America*.
<https://doi.org/10.1073/pnas.1800547115>
- Chater, N., & Loewenstein, G. (2016). The under-appreciated drive for sense-making. *Journal of Economic Behavior & Organization*, 126, 137–154.
- Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 29(37), 11560–11571.
- Clark, L., Robbins, T. W., Ersche, K. D., & Sahakian, B. J. (2006). Reflection impulsivity in current and former substance users. *Biological Psychiatry*, 60(5), 515–522.
- Coenen, A., & Gureckis, T. M. (2016). The distorting effect of deciding to stop sampling. *Pdfs.Semanticscholar.Org*.
<https://pdfs.semanticscholar.org/fbcf/783935679b02c6b66d3def9e5bdcaaa2626b.pdf>
- Cohanpour, M., Aly, M., & Gottlieb, J. (2022). Leveraging vision to understand curiosity. In *bioRxiv* (p. 2022.09.23.509220). <https://doi.org/10.1101/2022.09.23.509220>
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362(1481), 933–942.
- Conlisk, J. (1996). Why Bounded Rationality? *Journal of Economic Literature*, 34(2), 669–700.
- Cox, R. W., & Hyde, J. S. (1997). Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine*, 10(4–5), 171–178.
- Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology. General*, 104(3), 268–294.

- Dan, O., Leshkowitz, M., & Hassin, R. R. (2020). On clickbaits and evolution: curiosity from urge and interest. *Current Opinion in Behavioral Sciences*, 35, 150–156.
- Danek, A. H., Fraps, T., von Müller, A., Grothe, B., & Öllinger, M. (2014). It's a kind of magic-what self-reports can reveal about the phenomenology of insight problem solving. *Frontiers in Psychology*, 5, 1408.
- Danek, A. H., Williams, J., & Wiley, J. (2020). Closing the gap: connecting sudden representational change to the subjective Aha! experience in insightful problem solving. *Psychological Research*, 84(1), 111–119.
- Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, 16(2), 199–204.
- Desender, K., Boldt, A., & Yeung, N. (2018). Subjective Confidence Predicts Information Seeking in Decision Making. *Psychological Science*, 29(5), 761–778.
- Desender, K., Murphy, P., Boldt, A., Verguts, T., & Yeung, N. (2019). A post-decisional neural marker of confidence predicts information-seeking in decision-making. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.2620-18.2019>
- Ding, L., & Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cerebral Cortex*, 22(5), 1052–1067.
- Ditterich, J. (2006). Evidence for time-variant decision making. *The European Journal of Neuroscience*, 24(12), 3628–3641.
- Dolan, R. J., Fink, G. R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R. S., & Friston, K. J. (1997). How the brain learns to see objects and faces in an impoverished context. *Nature*, 389(6651), 596–599.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 32(11), 3612–3628.
- Dubey, R., & Griffiths, T. L. (2020a). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, 127(3), 455–476.
- Dubey, R., & Griffiths, T. L. (2020b). Understanding exploration in humans and machines by formalizing the function of curiosity. *Current Opinion in Behavioral Sciences*, 35, 118–124.
- Dubey, R., Griffiths, T. L., & Lombrozo, T. (2022). If it's important, then I'm curious: Increasing perceived usefulness stimulates curiosity. *Cognition*, 226, 105193.

- Edwards, W., & Slovic, P. (1965). SEEKING INFORMATION TO REDUCE THE RISK OF DECISIONS. *The American Journal of Psychology*, 78, 188–197.
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences of the United States of America*, 113(28), 7900–7905.
- Ely, J., Frankel, A., & Kamenica, E. (2015). Suspense and Surprise. *The Journal of Political Economy*, 123(1), 215–260.
- Ergorul, C., & Eichenbaum, H. (2004). The hippocampus and memory for “what,” “where,” and “when.” *Learning & Memory*, 11(4), 397–405.
- Esslinger, C., Braun, U., Schirmbeck, F., Santos, A., Meyer-Lindenberg, A., Zink, M., & Kirsch, P. (2013). Activation of midbrain and ventral striatal regions implicates salience processing during a modified beads task. *PloS One*, 8(3), e58536.
- Esteban, O., Ciric, R., Finc, K., Blair, R. W., Markiewicz, C. J., Moodie, C. A., Kent, J. D., Goncalves, M., DuPre, E., Gomez, D. E. P., Ye, Z., Salo, T., Valabregue, R., Amlien, I. K., Liem, F., Jacoby, N., Stojić, H., Cieslak, M., Urchs, S., ... Gorgolewski, K. J. (2020). Analysis of task-based functional MRI data preprocessed with fMRIPrep. *Nature Protocols*, 15(7), 2186–2202.
- Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S. S., Wright, J., Durnez, J., Poldrack, R. A., & Gorgolewski, K. J. (2019). fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature Methods*, 16(1), 111–116.
- Esteban, O., Markiewicz, C. J., Burns, C., Goncalves, M., Jarecka, D., Ziegler, E., Berleant, S., Ellis, D. G., Pinsard, B., Madison, C., Waskom, M., Notter, M. P., Clark, D., Manhães-Savio, A., Clark, D., Jordan, K., Dayan, M., Halchenko, Y. O., Loney, F., ... Ghosh, S. (2022). *nipy/nipype: 1.8.3*. <https://doi.org/10.5281/zenodo.6834519>
- Evans, A. C., Janke, A. L., Collins, D. L., & Baillet, S. (2012). Brain templates and atlases. *NeuroImage*, 62(2), 911–922.
- Fastrich, G. M., Kerr, T., Castel, A. D., & Murayama, K. (2018). The role of interest in memory for trivia questions: An investigation with a large-scale database. *Motivation Science*, 4(3), 227–250.
- Fastrich, G. M., & Murayama, K. (2018). *Development of Interest and Role of Choice During Sequential Knowledge Acquisition*. <https://doi.org/10.31219/osf.io/rhykp>
- Fastrich, G. M., & Murayama, K. (2020). Development of Interest and Role of Choice During Sequential Knowledge Acquisition. *AERA Open*, 6(2), 2332858420929981.

- FitzGerald, T. H. B., Schwartenbeck, P., Moutoussis, M., Dolan, R. J., & Friston, K. (2015). Active inference, evidence accumulation, and the urn task. *Neural Computation*, 27(2), 306–328.
- FitzGibbon, L., Komiya, A., & Murayama, K. (2021). The Lure of Counterfactual Curiosity: People Incur a Cost to Experience Regret. *Psychological Science*, 32(2), 241–255.
- Flanagin, V. L., Klinkowski, S., Brodt, S., Graetsch, M., Roselli, C., Glasauer, S., & Gais, S. (2023). The precuneus as a central node in declarative memory retrieval. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhac476>
- Fonov, V. S., Evans, A. C., McKinstry, R. C., Almlí, C. R., & Collins, D. L. (2009). Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, 47, S102.
- Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., & Ondobaka, S. (2017). Active Inference, Curiosity and Insight. *Neural Computation*, 29(10), 2633–2683.
- Furl, N., & Averbeck, B. B. (2011). Parietal cortex and insula relate to evidence seeking relevant to reward-related decisions. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 31(48), 17572–17582.
- Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 96(6), 3314–3319.
- Garner, W. R. (1962). *Uncertainty and structure as psychological concepts* (Vol. 369). Wiley Uncertainty and structure as psychological concepts.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103(4), 650–669.
- Gigerenzer, Gerd, & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 451–482.
- Gilmore, A. W., Nelson, S. M., & McDermott, K. B. (2015). A parietal memory network revealed by multiple MRI methods. *Trends in Cognitive Sciences*, 19(9), 534–543.
- Goh, A. X.-A., Bennett, D., Bode, S., & Chong, T. T.-J. (2021). Neurocomputational mechanisms underlying the subjective value of information. *Communications Biology*, 4(1), 1346.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574.

- Goldberg, R., & Todman, M. (2018). Induced boredom suppresses the recall of positively valenced information: A preliminary study. *Psychological Thought*, 11(1), 18–32.
- Golman, R., & Loewenstein, G. (2018). Information gaps: A theory of preferences regarding the presence and absence of information. *Decisions*, 5(3), 143–164.
- González-García, C., Flounders, M. W., Chang, R., Baria, A. T., & He, B. J. (2018). Content-specific activity in frontoparietal and default-mode networks during prior-guided visual perception. *ELife*, 7. <https://doi.org/10.7554/eLife.36068>
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15(1), 20–25.
- Gopnik, A. (2000). Explanation as orgasm and the drive for causal knowledge: The function, evolution, and phenomenology of the theory formation system. In F. C. Keil (Ed.), *Explanation and cognition*, (pp (Vol. 396, pp. 299–323). The MIT Press, x.
- Gopnik, A., Schulz, L., & Schulz, L. E. (2007). *Causal Learning: Psychology, Philosophy, and Computation*. Oxford University Press, USA.
- Gorgolewski, K., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L., & Ghosh, S. S. (2011). Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Frontiers in Neuroinformatics*, 5, 13.
- Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, reward, and information seeking. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 34(46), 15497–15504.
- Gottlieb, J., & Oudeyer, P.-Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews. Neuroscience*, 19(12), 758–770.
- Greve, D. N., & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48(1), 63–72.
- Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, 84(2), 486–496.
- Gruber, M. J., & Ranganath, C. (2019). How Curiosity Enhances Hippocampus-Dependent Memory: The Prediction, Appraisal, Curiosity, and Exploration (PACE) Framework. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2019.10.003>
- Gureckis, T. M., & Markant, D. B. (2012). Self-Directed Learning: A Cognitive and Computational Perspective. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 7(5), 464–481.

- Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., Halpern, D., Hamrick, J. B., & Chan, P. (2016). psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior Research Methods*, *48*(3), 829–842.
- Halamish, V., Madmon, I., & Moed, A. (2019). Motivation to Learn. *Experimental Psychology*, 1–12.
- Hauser, T. U., Moutoussis, M., NSPN Consortium, Dayan, P., & Dolan, R. J. (2017). Increased decision thresholds trigger extended information gathering across the compulsivity spectrum. *Translational Psychiatry*, *7*(12), 1296.
- Hauser, T. U., Moutoussis, M., Purg, N., Dayan, P., & Dolan, R. J. (2018). Beta-blocker propranolol modulates decision urgency during sequential information gathering. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0192-18.2018>
- Hausmann-Thürig, D., & Läge, D. (2008). Sequential evidence accumulation in decision making: the individual desired level of confidence can explain the extent of information acquisition. *Judgment and Decision Making*, *3*(3), 229–243.
- Howard, R. A. (1966). Information Value Theory. *IEEE Transactions on Systems Science and Cybernetics*, *2*(1), 22–26.
- Hsee, C. K., & Ruan, B. (2016). The Pandora Effect: The Power and Peril of Curiosity. *Psychological Science*, *27*(5), 659–666.
- Hsieh, P.-J., Vul, E., & Kanwisher, N. (2010). Recognition alters the spatial pattern of fMRI activation in early retinotopic cortex. *Journal of Neurophysiology*, *103*(3), 1501–1507.
- Hsiung, A., Poh, J.-H., Huettel, S. A., & Adcock, R. A. (2022). *Spoiler Alert! Curiosity encourages patience and joy in the presence of uncertainty*. <https://doi.org/10.31234/osf.io/x5hgc>
- Huettel, S. A., Song, A. W., & McCarthy, G. (2005). Decisions under uncertainty: probabilistic context influences activation of prefrontal and parietal cortices. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(13), 3304–3311.
- Hutchinson, J. B., Uncapher, M. R., Weiner, K. S., Bressler, D. W., Silver, M. A., Preston, A. R., & Wagner, A. D. (2014). Functional heterogeneity in posterior parietal cortex across attention and episodic memory retrieval. *Cerebral Cortex*, *24*(1), 49–66.
- Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E. M., & Stephan, K. E. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, *80*(2), 519–530.

- ligaya, K., Hauser, T. U., Kurth-Nelson, Z., O'Doherty, J. P., Dayan, P., & Dolan, R. J. (2020). The value of what's to come: Neural mechanisms coupling prediction error and the utility of anticipation. *Science Advances*, 6(25), eaba3828.
- ligaya, K., Story, G. W., Kurth-Nelson, Z., Dolan, R. J., & Dayan, P. (2016). The modulation of savouring by prediction error and its effects on choice. *ELife*, 5. <https://doi.org/10.7554/eLife.13747>
- Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The Effort Paradox: Effort Is Both Costly and Valued. *Trends in Cognitive Sciences*, 22(4), 337–349.
- Jach, H. K., DeYoung, C. G., & Smillie, L. D. (2021). Why do people seek information? The role of personality traits and situation perception. *Journal of Experimental Psychology. General*. <https://doi.org/10.1037/xge0001109>
- Jenkinson, M., & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2), 143–156.
- Jenkinson, Mark, Bannister, P., Brady, M., & Smith, S. (2002). Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *NeuroImage*, 17(2), 825–841.
- Jepma, M., Verdonschot, R. G., van Steenbergen, H., Rombouts, S. A. R. B., & Nieuwenhuis, S. (2012). Neural mechanisms underlying the induction and relief of perceptual curiosity. *Frontiers in Behavioral Neuroscience*, 6, 5.
- Jiwa, M., Cooper, P. S., Chong, T. T.-J., & Bode, S. (2021). Choosing increases the value of non-instrumental information. *Scientific Reports*, 11(1), 8780.
- Juni, M. Z., Gureckis, T. M., & Maloney, L. T. (2016). Information sampling behavior with explicit sampling costs. *Decision (Washington, D.C.)*, 3(3), 147–168.
- Kaanders, P., Juechems, K., O'Reilly, J., & Hunt, L. (2021). Dissociable mechanisms of information sampling in prefrontal cortex and the dopaminergic system. *Current Opinion in Behavioral Sciences*, 41, 63–70.
- Kaanders, P., Nili, H., O'Reilly, J. X., & Hunt, L. T. (2021). Medial frontal cortex activity predicts information sampling in economic choice. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0392-21.2021>
- Käckenmester, W., Kroencke, L., & Wacker, J. (2018). Frontal asymmetry predicts the incentive value of perceptual information. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, 134, 22–29.
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T.-Y., & Camerer, C. F. (2009). The wick in the candle of learning: epistemic curiosity

- activates reward circuitry and enhances memory. *Psychological Science*, 20(8), 963–973.
- Kaplan, F., & Oudeyer, P.-Y. (2007). In search of the neural circuits of intrinsic motivation. *Frontiers in Neuroscience*, 1(1), 225–236.
- Kashdan, T. B., Stikma, M. C., Disabato, D. J., McKnight, P. E., Bekier, J., Kaji, J., & Lazarus, R. (2018). The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *Journal of Research in Personality*, 73, 130–149.
- Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron*, 84(6), 1329–1342.
- Kiani, R., Hanks, T. D., & Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 28(12), 3017–3029.
- Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron*, 88(3), 449–460.
- Kira, S., Yang, T., & Shadlen, M. N. (2015). A neural implementation of Wald's sequential probability ratio test. *Neuron*, 85(4), 861–873.
- Klimmt, C., Rizzo, A., Vorderer, P., Koch, J., & Fischer, T. (2009). Experimental evidence for suspense as determinant of video game enjoyment. *Cyberpsychology & Behavior: The Impact of the Internet, Multimedia and Virtual Reality on Behavior and Society*, 12(1), 29–31.
- Knutson, B., Fong, G. W., Adams, C. M., Varner, J. L., & Hommer, D. (2001). Dissociation of reward anticipation and outcome with event-related fMRI. *Neuroreport*, 12(17), 3683.
- Kobayashi, K., & Hsu, M. (2019). Common neural code for reward and information value. In *Proc. Natl. Acad. Sci. U. S. A.* (p. 324665). <https://doi.org/10.1073/pnas.1820145116>
- Kobayashi, K., Lee, S., Filipowicz, A. L. S., McGaughey, K. D., Kable, J. W., & Nassar, M. R. (2021). Dynamic Representation of the Subjective Value of Information. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0423-21.2021>
- Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-019-0589-3>

- Kolling, N., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2012). Neural mechanisms of foraging. *Science*, 336(6077), 95–98.
- Korn, C. W., & Bach, D. R. (2018). Heuristic and optimal policy computations in the human brain during sequential decision-making. *Nature Communications*, 9(1), 325.
- Kravitz, D. J., Saleem, K. S., Baker, C. I., Ungerleider, L. G., & Mishkin, M. (2013). The ventral visual pathway: an expanded neural framework for the processing of object quality. *Trends in Cognitive Sciences*, 17(1), 26–49.
- Kuznetsova, A., & Brockhoff, P. B. (2017). lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software*. <http://statistik-jstat.uibk.ac.at/index.php/jss/article/view/v082i13>
- Lang, E. A., van Geen, C., Tedeschi, E., Marvin, C. B., & Shohamy, D. (2021). Learned temporal statistics guide information seeking and shape memory. *Journal of Experimental Psychology: General*. <https://doi.org/10.1037/xge0001122>
- Lau, J. K. L., Ozono, H., Kuratomi, K., Komiya, A., & Murayama, K. (2020). Shared striatal activity in decisions to satisfy curiosity and hunger at the risk of electric shocks. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-020-0848-3>
- Leech, R., Braga, R., & Sharp, D. J. (2012). Echoes of the brain within the posterior cingulate cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 32(1), 215–222.
- Leotti, L. A., & Delgado, M. R. (2011). The inherent reward of choice. *Psychological Science*, 22(10), 1310–1318.
- Li, Y., Daddaoua, N., Horan, M., Foley, N. C., & Gottlieb, J. (2022). Uncertainty modulates visual maps during noninstrumental information demand. *Nature Communications*, 13(1), 5911.
- Li, Z.-W., Bramley, N. R., & Gureckis, T. M. (2021). Expectations about future learning influence moment-to-moment feelings of suspense. *Cognition & Emotion*, 35(6), 1099–1120.
- Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, 124(6), 762–794.
- Ligneul, R., Mermillod, M., & Morisseau, T. (2018). From relief to surprise: Dual control of epistemic curiosity in the human brain. *NeuroImage*. <https://doi.org/10.1016/j.neuroimage.2018.07.038>
- Liquin, E. G., & Lombrozo, T. (2020a). A functional approach to explanation-seeking curiosity. *Cognitive Psychology*, 119, 101276.

- Liquin, E. G., & Lombrozo, T. (2020b). Explanation-seeking curiosity in childhood. *Current Opinion in Behavioral Sciences*, 35, 14–20.
- Liquin, E. G., & Lombrozo, T. (2021). Motivated to learn: An account of explanatory satisfaction. *Cognitive Psychology*, 132, 101453.
- Liquin, E., & Lombrozo, T. (2017). Explain, Explore, Exploit: Effects of Explanation on Information Search. *CogSci*.
<https://pdfs.semanticscholar.org/0da5/3f6ec0476d9cc1eb1ad7dc2f72f097d1d096.pdf>
- Litman, J. A. (2010). Relationships between measures of I- and D-type curiosity, ambiguity tolerance, and need for closure: An initial test of the wanting-liking model of information-seeking. *Personality and Individual Differences*, 48(4), 397–402.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*. <http://psycnet.apa.org/record/1994-41058-001>
- Luwel, K., Schillemans, V., Onghena, P., & Verschaffel, L. (2009). Does switching between strategies within the same task involve a cost? *British Journal of Psychology*, 100(Pt 4), 753–771.
- Lydon-Staley, D. M., Zhou, D., Blevins, A. S., Zurn, P., & Bassett, D. S. (2019). *Hunters, busybodies, and the knowledge network building associated with curiosity*.
<https://doi.org/10.31234/osf.io/undy4>
- Lydon-Staley, D. M., Zhou, D., Blevins, A. S., Zurn, P., & Bassett, D. S. (2021). Hunters, busybodies and the knowledge network building associated with deprivation curiosity. *Nature Human Behaviour*, 5(3), 327–336.
- Majerczak, P., & Strzelecki, A. (2022). Trust, Media Credibility, Social Ties, and the Intention to Share towards Information Verification in an Age of Fake News. *Behavioral Sciences*, 12(2). <https://doi.org/10.3390/bs12020051>
- Malhotra, G., Leslie, D. S., Ludwig, C. J. H., & Bogacz, R. (2017). Overcoming indecision by changing the decision boundary. *Journal of Experimental Psychology. General*, 146(6), 776–805.
- Manohar, S. G., Chong, T. T.-J., Apps, M. A. J., Batla, A., Stamelou, M., Jarman, P. R., Bhatia, K. P., & Husain, M. (2015). Reward Pays the Cost of Noise Reduction in Motor and Cognitive Control. *Current Biology: CB*, 25(13), 1707–1716.
- Markant, D. B., Ruggeri, A., Gureckis, T. M., & Xu, F. (2016). Enhanced Memory as a Common Effect of Active Learning. *Mind, Brain and Education: The Official Journal of the International Mind, Brain, and Education Society*, 10(3), 142–152.

- Marvin, C. B., & Shohamy, D. (2016). Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *Journal of Experimental Psychology. General*, 145(3), 266–272.
- Marvin, C. B., Tedeschi, E., & Shohamy, D. (2020). Curiosity as the impulse to know: common behavioral and neural mechanisms underlying curiosity and impulsivity. *Current Opinion in Behavioral Sciences*, 35, 92–98.
- Mather, M., & Sutherland, M. R. (2011). Arousal-Biased Competition in Perception and Memory. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 6(2), 114–133.
- Maxwell, L. C. (2021). Spoilers ahead, proceed with caution: How engagement, enjoyment, and FoMO predict avoidance of spoilers. *Psychology of Popular Media*. <https://doi.org/10.1037/ppm0000362>
- Metcalfe, J., & Wiebe, D. (1987). Intuition in insight and noninsight problem solving. *Memory & Cognition*, 15(3), 238–246.
- Metcalfe, Janet. (2017). Learning from Errors. *Annual Review of Psychology*, 68, 465–489.
- Metcalfe, Janet, Kennedy-Pyers, T., & Vuorre, M. (2021). Curiosity and the desire for agency: wait, wait ... don't tell me! *Cognitive Research: Principles and Implications*, 6(1), 69.
- Moncada, D., & Viola, H. (2007). Induction of long-term memory by exposure to novelty requires protein synthesis: evidence for a behavioral tagging. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(28), 7476–7481.
- Moscovitch, M., Cabeza, R., Winocur, G., & Nadel, L. (2016). Episodic Memory and Beyond: The Hippocampus and Neocortex in Transformation. *Annual Review of Psychology*, 67, 105–134.
- Moutoussis, M., Bentall, R. P., El-Deredy, W., & Dayan, P. (2011). Bayesian modelling of Jumping-to-Conclusions bias in delusional patients. *Cognitive Neuropsychiatry*, 16(5), 422–447.
- Muhammed T, S., & Mathew, S. K. (2022). The disaster of misinformation: a review of research in social media. *International Journal of Data Science and Analytics*, 13(4), 271–285.
- Mullaney, K. M., Carpenter, S. K., Grotenhuis, C., & Burianek, S. (2014). Waiting for feedback helps if you want to know the answer: the role of curiosity in the delay-of-feedback benefit. *Memory & Cognition*, 42(8), 1273–1284.

- Murayama, K. (2022). A reward-learning framework of knowledge acquisition: An integrated account of curiosity, interest, and intrinsic-extrinsic rewards. *Psychological Review*, 129(1), 175–198.
- Murayama, K., FitzGibbon, L., & Sakaki, M. (2019). Process Account of Curiosity and Interest: A Reward-Learning Perspective. *Educational Psychology Review*. <https://doi.org/10.1007/s10648-019-09499-9>
- Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature Communications*, 7, 13526.
- Murty, V. P., DuBrow, S., & Davachi, L. (2015). The simple act of choosing influences declarative memory. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 35(16), 6255–6264.
- Murty, V. P., DuBrow, S., & Davachi, L. (2019). Decision-making Increases Episodic Memory via Postencoding Consolidation. *Journal of Cognitive Neuroscience*, 31(9), 1308–1317.
- Noordewier, M. K., & van Dijk, E. (2017). Curiosity and time: from not knowing to almost knowing. *Cognition & Emotion*, 31(3), 411–421.
- Noordewier, M. K., & van Dijk, E. (2020). Deprivation and discovery motives determine how it feels to be curious. *Current Opinion in Behavioral Sciences*, 35, 71–76.
- Offen, S., Schluppeck, D., & Heeger, D. J. (2009). The role of early visual cortex in visual short-term memory and visual attention. *Vision Research*, 49(10), 1352–1362.
- Oh, Y., Chesebrough, C., Erickson, B., Zhang, F., & Kounios, J. (2020). An insight-related neural reward signal. *NeuroImage*, 214, 116757.
- Oosterwijk, S. (2017). Choosing the negative: A behavioral demonstration of morbid curiosity. *PloS One*, 12(7), e0178399.
- Oosterwijk, S., Snoek, L., Tekoppele, J., Engelbert, L. H., & Scholte, H. S. (2020). Choosing to view morbid information involves reward circuitry. *Scientific Reports*, 10(1), 15291.
- Otto, A. R., & Daw, N. D. (2019). The opportunity cost of time modulates cognitive effort. *Neuropsychologia*, 123, 92–105.
- Oudeyer, P.-Y., Gottlieb, J., & Lopes, M. (2016). Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. *Progress in Brain Research*, 229, 257–284.

- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *14*(3), 534–552.
- Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Computational Biology*, *7*(1), e1001048.
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203.
- Petit, P., Attaallah, B., Manohar, S. G., & Husain, M. (2021). The computational cost of active information sampling before decision-making under uncertainty. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-021-01116-6>
- Pine, A., Sadeh, N., Ben-Yakov, A., Dudai, Y., & Mendelsohn, A. (2018). Knowledge acquisition is governed by striatal prediction errors. *Nature Communications*, *9*(1), 1673.
- Poh, J.-H., Vu, M.-A. T., Stanek, J. K., Hsiung, A., Egner, T., & Adcock, R. A. (2022). Hippocampal convergence during anticipatory midbrain activation promotes subsequent memory formation. *Nature Communications*, *13*(1), 6729.
- Potts, R., & Shanks, D. R. (2014). The benefit of generating errors during learning. *Journal of Experimental Psychology. General*, *143*(2), 644–667.
- Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience*, *19*(3), 366–374.
- Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage*, *84*, 320–341.
- Preuschoff, K., Quartz, S. R., & Bossaerts, P. (2008). Human insula activation reflects risk prediction errors as well as risk. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *28*(11), 2745–2752.
- Rakow, T., Newell, B. R., & Zougkou, K. (2010). The role of working memory in information acquisition and decision making: lessons from the binary prediction task. *Quarterly Journal of Experimental Psychology*, *63*(7), 1335–1360.
- Rieskamp, J., & Otto, P. E. (2006). SSL: a theory of how people learn to select strategies. *Journal of Experimental Psychology. General*, *135*(2), 207–236.
- Ripollés, P., Marco-Pallarés, J., Alicart, H., Tempelmann, C., Rodríguez-Fornells, A., & Noesselt, T. (2016). Intrinsic monitoring of learning success facilitates memory

encoding via the activation of the SN/VTa-Hippocampal loop. *ELife*, 5.
<https://doi.org/10.7554/eLife.17441>

- Rodriguez Cabrero, J. A. M., Zhu, J.-Q., & Ludvig, E. A. (2019). Costly curiosity: People pay a price to resolve an uncertain gamble early. *Behavioural Processes*, 160, 20–25.
- Rotem-Turchinski, N., Ramaty, A., & Mendelsohn, A. (2018). The opportunity to choose enhances long-term episodic memory. *Memory*, 1–10.
- Rouhani, N., Norman, K. A., & Niv, Y. (2018). Dissociable effects of surprising rewards on learning and memory. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 44(9), 1430–1443.
- Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition*, 193, 104013.
- Ryakhovskaya, Y., Jach, H. K., & Smillie, L. D. (2021). Curiosity as feelings of interest versus deprivation: Relations between curiosity traits and affective states when anticipating information. *Journal of Research in Personality*, 104164.
- Salvetti, B., Morris, R. G. M., & Wang, S.-H. (2014). The role of rewarding and novel events in facilitating memory persistence in a separate spatial memory task. *Learning & Memory*, 21(2), 61–72.
- Sang, K., Todd, P. M., Goldstone, R. L., & Hills, T. T. (2020). Simple Threshold Rules Solve Explore/Exploit Trade-offs in a Resource Accumulation Search Task. *Cognitive Science*, 44(2), e12817.
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, 43(4), 1045–1050.
- Schulz, L., Fleming, S. M., & Dayan, P. (2023). Metacognitive computations for information search: Confidence in control. *Psychological Review*.
<https://doi.org/10.1037/rev0000401>
- Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H. B., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. In *Elife* (p. e41703). <https://doi.org/10.7554/eLife.41703>
- Schwarz, G. (1978). Estimating the Dimension of a Model. *Annals of Statistics*, 6(2), 461–464.
- Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., Reiss, A. L., & Greicius, M. D. (2007). Dissociable intrinsic connectivity networks for

- salience processing and executive control. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(9), 2349–2356.
- Sestieri, C., Shulman, G. L., & Corbetta, M. (2017). The contribution of the human posterior parietal cortex to episodic memory. *Nature Reviews. Neuroscience*, 18(3), 183–192.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4), 1916–1936.
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: an effort-reduction framework. *Psychological Bulletin*, 134(2), 207–222.
- Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-019-0793-1>
- Shen, W., Tong, Y., Li, F., Yuan, Y., Hommel, B., Liu, C., & Luo, J. (2018). Tracking the neurodynamics of insight: A meta-analysis of neuroimaging studies. *Biological Psychology*, 138, 189–198.
- Shen, W., Yuan, Y., Liu, C., & Luo, J. (2016). In search of the ‘Aha!’ experience: Elucidating the emotionality of insight problem-solving. *British Journal of Psychology*, 107(2), 281–298.
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron*, 79(2), 217–240.
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a Rational and Mechanistic Account of Mental Effort. *Annual Review of Neuroscience*, 40, 99–124.
- Shenhav, A., Straccia, M. A., Musslick, S., Cohen, J. D., & Botvinick, M. M. (2018). Dissociable neural mechanisms track evidence accumulation for selection of attention versus action. *Nature Communications*, 9(1), 2485.
- Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, 14(10), 464–472.
- Silver, M. A., Ress, D., & Heeger, D. J. (2007). Neural correlates of sustained spatial attention in human early visual cortex. *Journal of Neurophysiology*, 97(1), 229–237.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.

- Simon, Herbert A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Sinclair, A. H., & Barense, M. D. (2018). Surprise and destabilize: prediction error influences episodic memory reconsolidation. *Learning & Memory*, 25(8), 369–381.
- Sinclair, A. H., Manalili, G. M., Brunec, I. K., Adcock, R. A., & Barense, M. D. (2021). Prediction errors disrupt hippocampal representations and update episodic memories. *Proceedings of the National Academy of Sciences of the United States of America*, 118(51). <https://doi.org/10.1073/pnas.2117625118>
- Stare, C. J., Gruber, M. J., Nadel, L., Ranganath, C., & Gómez, R. L. (2018). Curiosity-driven memory enhancement persists over time but does not benefit from post-learning sleep. *Cognitive Neuroscience*, 1–16.
- Stark, C. E., & Squire, L. R. (2001). When zero is not zero: the problem of ambiguous baseline conditions in fMRI. *Proceedings of the National Academy of Sciences of the United States of America*, 98(22), 12760–12766.
- Steyvers, M., Lee, M. D., & Wagenmakers, E.-J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3), 168–179.
- Su-lin, G., Tuggle, C. A., Mitrook, M. A., Coussement, S. H., & Zillmann, D. (1997). THE THRILL OF A CLOSE GAME: Who Enjoys It and Who Doesn't? *Journal Of Sport & Social Issues*, 21(1), 53–64.
- Summerfield, J. J., Hassabis, D., & Maguire, E. A. (2009). Cortical midline involvement in autobiographical memory. *NeuroImage*, 44(3), 1188–1200.
- Ten, A., Kaushik, P., Oudeyer, P.-Y., & Gottlieb, J. (2021). Humans monitor learning progress in curiosity-driven exploration. *Nature Communications*, 12(1), 5972.
- Tik, M., Sladky, R., Luft, C. D. B., Willinger, D., Hoffmann, A., Banissy, M. J., Bhattacharya, J., & Windischberger, C. (2018). Ultra-high-field fMRI insights on insight: Neural correlates of the Aha!-moment. *Human Brain Mapping*, 39(8), 3241–3252.
- Todd, P. M., & Gigerenzer, G. (2000). Précis of Simple heuristics that make us smart. *The Behavioral and Brain Sciences*, 23(5), 727–741; discussion 742-80.
- Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., & Gee, J. C. (2010). N4ITK: improved N3 bias correction. *IEEE Transactions on Medical Imaging*, 29(6), 1310–1320.
- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, 71(5), 680–683.

- Van de Cruys, S., Damiano, C., Boddez, Y., Król, M., Goetschalckx, L., & Wagemans, J. (2021). Visual affects: Linking curiosity, Aha-Erlebnis, and memory through information gain. *Cognition*, *212*, 104698.
- van der Linden, S. (2022). Misinformation: susceptibility, spread, and interventions to immunize the public. *Nature Medicine*, *28*(3), 460–467.
- van Hooft, E. A. J., & van Hooff, M. L. M. (2018). The state of boredom: Frustrating or depressing? *Motivation and Emotion*, *42*(6), 931–946.
- van Lieshout, L. L. F., de Lange, F. P., & Cools, R. (2020). Why so curious? Quantifying mechanisms of information seeking. *Current Opinion in Behavioral Sciences*, *35*, 112–117.
- van Lieshout, L. L. F., de Lange, F. P., & Cools, R. (2021). Uncertainty increases curiosity, but decreases happiness. *Scientific Reports*, *11*(1), 14014.
- van Lieshout, L. L. F., Vandenbroucke, A. R. E., Müller, N. C. J., Cools, R., & de Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal activity. In *J. Neurosci.* (p. 195461). <https://doi.org/10.1523/JNEUROSCI.2816-17.2018>
- Vidal, A., Soto-Faraco, S., & Moreno Bote, R. (2022). Balance between breadth and depth in human many-alternative decisions. *ELife*, *11*. <https://doi.org/10.7554/eLife.76985>
- Vogl, E., Pekrun, R., Murayama, K., & Loderer, K. (2019). Surprised-curious-confused: Epistemic emotions and knowledge exploration. *Emotion*. <https://doi.org/10.1037/emo0000578>
- Vogl, E., Pekrun, R., Murayama, K., Loderer, K., & Schubert, S. (2019). Surprise, Curiosity, and Confusion Promote Knowledge Exploration: Evidence for Robust Effects of Epistemic Emotions. *Frontiers in Psychology*, *10*, 2474.
- Voss, J. L., Gonsalves, B. D., Federmeier, K. D., Tranel, D., & Cohen, N. J. (2011). Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nature Neuroscience*, *14*(1), 115–120.
- Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-019-01598-6>
- Wade, S., & Kidd, C. (2020). Influence of Topic Knowledge on Curiosity. *CogSci*.
- Wager, T. D., Jonides, J., Smith, E. E., & Nichols, T. E. (2005). Toward a taxonomy of attention shifting: individual differences in fMRI during multiple shift types. *Cognitive, Affective & Behavioral Neuroscience*, *5*(2), 127–143.

- Wang, J. J., Yang, Y., Macias, C., & Bonawitz, E. (2021). Children With More Uncertainty in Their Intuitive Theories Seek Domain-Relevant Information. *Psychological Science*, 956797621994230.
- Wang, M. Z., & Hayden, B. Y. (2019). Monkeys are curious about counterfactual outcomes. *Cognition*, 189, 1–10.
- Wimpenny, K., & Clarke, S. (2017). Curiosity, Mystery and Risk: Exploring the Challenges of Adopting Creative and Playful Pedagogies in Higher Education. *Playful Learning*. <https://pureportal.coventry.ac.uk/en/publications/curiosity-mystery-and-risk-exploring-the-challenges-of-adopting-c>
- Wittmann, B. C., Bunzeck, N., Dolan, R. J., & Düzel, E. (2007). Anticipation of novelty recruits reward system and hippocampus while promoting recollection. *NeuroImage*, 38(1), 194–202.
- Wojtowicz, Z., Chater, N., Loewenstein, G., Dezza, I. C., Schulz, E., & Wu, C. M. (2022). The Motivational Processes of Sense-Making. In *The Drive for Knowledge: The Science of Human Information Seeking* (pp. 3–30). Cambridge University Press.
- Wu, C. M., Meder, B., Filimon, F., & Nelson, J. D. (2017). Asking better questions: How presentation formats influence information search. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 43(8), 1274–1297.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, 2(12), 915–924.
- Yang, S. C.-H., Lengyel, M., & Wolpert, D. M. (2016). Active sensing in the categorization of visual patterns. *ELife*, 5. <https://doi.org/10.7554/eLife.12215>
- Yang, S. C.-H., Wolpert, D. M., & Lengyel, M. (2016). Theoretical perspectives on active sensing. *Current Opinion in Behavioral Sciences*, 11, 100–108.
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46(4), 681–692.
- Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1), 45–57.
- Zhou, D., Lydon-Staley, D. M., Zurn, P., & Bassett, D. S. (2020). The growth and form of knowledge networks by kinesthetic curiosity. In *arXiv [q-bio.NC]*. arXiv. <http://arxiv.org/abs/2006.02949>
- Zurn, P., & Bassett, D. S. (2018). On Curiosity: A Fundamental Aspect of Personality, a Practice of Network Growth. *Personality Neuroscience*, 1. <https://doi.org/10.1017/pen.2018.3>

Biography

Abigail Hsiung attended University of Maryland, College Park and in 2015 graduated *magna cum laude* with dual Bachelor of Science degrees in Psychology and Kinesiology. After graduation, she worked for Dr. Christian Grillon and Dr. Monique Ernst at the National Institute of Mental Health in the Section for the Neurobiology of Fear and Anxiety. In her time there, she worked on a variety of studies examining the effects of anxiety on memory, executive functioning, and risky decision-making. In 2017, Abby started graduate school at Duke University through the Cognitive Neuroscience Admitting Program. She completed rotations with Drs. Alison Adcock, Scott Huettel, and John Pearson before joining the labs of both Drs. Adcock and Huettel and affiliating with the Department of Psychology and Neuroscience. Abby's research examines how motivation, in particular curiosity, guides information seeking and learning. During her graduate career, she was awarded the National Defense Science and Engineering Graduate Fellowship by the Department of Defense in 2018, Lafitte Foundation Graduate Research Awards in 2018, 2021, and 2022, and the Lafitte Foundation Undergraduate Mentorship Award in 2021. In 2020, she helped found the Cognitive Neuroscience Research Internship, an organization that promotes equitable access to undergraduate research opportunities in psychology and neuroscience.