# Explore/Exploit: Tradeoffs in Decision Making

Krish Bahel

## Abstract

Explore-exploit tradeoffs–the conflict between *exploring* novel options and *exploiting* familiar ones–is a fundamental decision model adapted from basic and translational science. Striking the right balance between these two strategies is pivotal for achieving efficient outcomes and adapting to varying levels of uncertainty. Individuals must also adapt to varying levels of confidence and external factors that hold implications for their decisions. This review aims to shed light on the influential role of various cognitive (e.g., confidence, bias) and affective processes (e.g., stress) on explore-exploit decision making. We also cover the role modern neuroscience has played in studying this tradeoff and its underlying neural circuitry. This topic holds profound importance in making real-world developments across diverse disciplines. In economics, understanding how confidence impacts decision making can elucidate market behaviors and financial choices. In addition, this research advances models of artificial intelligence and human-computer interaction (HCI), which are highly reliant on understanding principles of decision. Lastly, understanding the underlying brain pathways can provide psychological insights into cognitive flexibility, motivational tendencies, and human learning; indeed, these are critical processes that, if perturbed, underscore the etiology and maintenance of a variety of psychiatric illnesses.

## Keywords

Behavior and Social Sciences, Cognitive Psychology, Decision Making, Explore-Exploit

## Introduction

Dining at your usual spot for burgers or venturing to the new dumpling restaurant down the street is a mundane decision with relatively low stakes. However, at its core, this reflects a conflict between explore and exploit motivations. The conflict splits the highway of our neural decisions into two separately winding roads: Do you want to continue *exploiting* the great taste that you know and love, or do you want to *explore* alternative options and maybe find a new favorite? This dynamic interplay is a theme that courses through our lives, underlying many human behaviors and prompting us to contemplate whether to tread our familiar paths or venture into uncharted territories of inquiry. Interestingly, clinical research suggests studying explore-exploit decision making in humans may help improve our understanding of transdiagnostic features of various psychopathologies[1, 2, 3].

From research in economics to neuroscience, explore-exploit decision making has been extensively explored; indeed, recent work has begun to pursue a mechanistic understanding of these computation[1]. This work not only deepens our comprehension of cognitive processes governing our choices but also broadly contributes to an evolutionary brain-behavior framework for understanding how we traverse the complexity of our environment. In this brief review, we first go over common approaches and paradigms that cognitive neuroscience has used to advance our mechanistic understanding of explore-exploit decision making in humans. Second, we cover the neural underpinnings of the explore-exploit dilemma in addition to cognitive and emotional aspects. Lastly, we will discuss the clinical relevance and real-world implications of this research.

## Results and Discussion

### *Assessing exploration and exploitation in humans*

Neuroscience has advanced our measurement of explore-exploit tradeoffs in exciting ways, adapting human-appropriate paradigms from basic and translational science and leveraging multimodal neuroimaging, eye-tracking, real-time psychophysiology, and other cutting-edge technologies[4]. These methods help achieve a deeper mechanistic understanding of this phenomenon. By utilizing these paradigms and neuroimaging technologies, scientists are able to design experiments and studies that offer valuable insights.

Slot machine tasks and multi-armed bandit paradigms have been go-to experimental paradigms for studying exploration and exploitation in humans. These bandit and slot paradigms primarily rely on the gambling instincts of participants to reflect individual- or group-level patterns of resolving competing motivations. To foster authenticity, tasks often incorporate real monetary-based conflicts and other impacts to probe “real” decision conflicts in digital environments[5]. To this end, many of these slot machine and multi-armed bandit tasks are completed on a screen or even online. Participants must decide whether to persist in *exploiting* rewards offered by their current machine or to cast their gaze outward, *exploring* novel options in the pursuit of more favorable outcomes. In slot machine paradigms, participants gamble as if they were genuinely playing at a set of regular slot machines, aiming for the highest payout[6]. Many papers define the act of switching machines as *exploration* and the act of continuing at the same machine *exploitation*[7,8]. Multi-armed bandit tasks are a bit more complex; here, there are multiple slots that have different reward outcomes. Participants also are typically asked to complete subjective ratings after each trial, measuring contributions of emotion and confidence in decision-making[6].

These paradigms also have their drawbacks. First, there has long been concern about how lab-based paradigms translate in real-world contexts. For instance, people may be more likely to exhibit more explorative tendencies when in the lab which could reflect how participants are appraising stakes. Additionally, we may lose important ecological context by oversimplifying decision conflicts in lab environments that become learned after numerous trials[9]. Other paradigms do exist, however, they may be better suited for other types of decision conflicts (e.g., go/no-go – approach-avoidance conflict)[10]. Indeed, we are due for major advancements in mechanism-focused explore-exploit tradeoff experiments that prioritize ecological validity while preserving rigor and ethics.

As these paradigms and experimental designs facilitate the accumulation of data, various tools of neuroscience are poised to quantify brain activity during explore-exploit decision making. Functional magnetic resonance imaging (fMRI) and electroencephalogram (EEG) have been two widely used neuroscience tools for acquiring brain data while participants perform tasks. While both are useful for detecting brain activity, EEG specifically records electrical responses and is known to have better temporal resolution (i.e., real-time data acquisition). In contrast, fMRI records hemodynamic responses and yields better spatial resolution. Both EEG and fMRI are critical to identifying neural correlates associated with exploration and exploitation strategies[11].

Effectiveness has been found with usage in the animal literature and translation over to humans. For example, in studies involving chimpanzees, EEG has been used to uncover the biology of value setting in primates. Findings from Averbeck and Costa[12] show how setting expected values for specific circumstances involves various regions of the brain, including the orbitofrontal cortex and ventral striatum, which are responsible for emotion and reward processing. These Initial and Final Expected Values (IEV and FEV, respectively) are correlated with how an individual would choose between taking an explorative pathway rather than an exploitative one. A high IEV would suggest significant rewards, potentially influencing exploitation. A low IEV, however, would suggest low or uncertain rewards and may influence exploration. FEVs represent values obtained by participants. High FEVs may incentivize repetitive exploitation. On the other hand, fMRI captures changes in blood flow and oxygenation levels in various brain regions, commonly known as BOLD, or Blood Oxygenation Level Dependent[13]. By studying morphometry and activitation of different brain areas, researchers are able to identify regions associated with exploration, exploitation, and the dynamic interplay between the two. Since decision encompasses complex cognitive processes that likely are supported by distributed neural activity, many researchers use BOLD activation to provide insight into connectivity across different regions. In short, connectivity between regions is inferred when two or more brain regions show similar activity while performing an action[14]. Assessing the physical structure of the human brain is just one of the fMRI's many important applications. By studying structural differences between individuals, researchers can detect variations in different regions and how they could create individual cognitive differences[15]. In the context of Human-Computer Interaction (HCI), researchers' methodology aligns with the objective of understanding how users navigate digital interfaces during exploration and exploitation tasks. By leveraging established HCI paradigms, they aimed to uncover user behaviors and decision-making patterns within real-world digital environments. For example, many of these slot machine and bandit paradigms can be completed on a screen or even online.

### *The neuroscience of explore-exploit decision-making*

Researchers have made significant progress in understanding the biological foundations of explore-exploit tradeoffs. They have identified particular brain areas that are more active throughout each of these decision-making pathways with fMRI. Blanchard and Gershman[7] noted the ventromedial area of the prefrontal cortex (PFC) is significantly more active during exploitative actions. In their study, Lauriero-Martinez and colleagues, the ventromedial PFC consistently emerged as an active region during exploitative tasks, reaffirming its role in encoding Immediate Expected Value[16] .In contrast, when exploring more novel decision-making territory, Daw & colleagues[17] found the frontopolar cortex (FPC) to be relevant. A recurrent theme in these papers and articles is the definition of exploration, often defined as a participant departing from their current course of action, akin to switching from one slot machine to another. Building on these insights, Lauriero-Martinez and her colleagues[16] reported heightened FPC activation during explorative activities. Notably, their findings also revealed activation in other regions associated with attention control, such as the temporoparietal junction and the superior parietal lobule, indicating potential functional connectivity among these brain regions[16]. These collective findings suggest a network of interconnected regions that collaborate in decision-making processes.

### *Cognitive factors*

Extant work has focused on identifying key cognitive processes relevant to resolving explore-exploit conflicts. One fundamental cognitive process at play is our capacity for risk assessment. When faced with the choice of venturing into the unknown or sticking with familiar territory, our brain engages in a sophisticated evaluation of potential risks and rewards. This cognitive weighing of uncertainty versus predictability impacts whether we lean towards exploration or exploitation. Dombrovski and Hallquist[18] highlight how variations in cognitive risk assessment may be associated with addictive tendencies, underscoring the intricate link between cognition and behavior. In addition, delay discounting, favoring (smaller) immediate over (larger) delayed rewards, seems to influence explore-exploit decision making significantly. Those with higher delay discounting tendencies prioritize immediate gains (exploitation), while those with weaker tendencies prefer assured but delayed rewards (exploration), compromising long-term decision-making[18].

Further, memory and immediate processing seem to be influential to our decision-making calculus. When considering exploration, our brain draws upon past experiences stored in long-term memory to assess potential outcomes, while short-term memory aids in processing the immediate information available. In turn, this adds personal biases and learning into the equation. Blanchard and Gershman[7] show how people are more likely to be exploitative after learning patterns or tendencies. For example, if an individual realizes that a game of online poker is patterned to have twice as many rewards in every third round, they are more likely to bet more on the third round to exploit higher rewards. In their study, participants are asked to bet on the color of the light, which has been patterned, that they expect to flash in front of them. After learning the pattern of light colors, they switched to more exploitative decisions and betted more frequently. Lastly, confidence also has been shown to play a role in explore-exploit decision making. According to Boldt & colleagues[6], participants who provide a higher confidence rating report being more sure of their exploitative behaviors. They are more likely to take risks involved with deciding to remain with their current bandit, exhibiting higher exploitative tendencies.

### *The role of emotion*

Since the rise of affectivism, the inclusion of emotion and affect on behavior, there has been an influx of research focusing on the role of emotions in decision-making. Emotions appear to hold a strong basis on whether a person is explorative or exploitative. For example, stress can produce a strong desire for avoidance and a distaste for exploration. Under stress, people may choose to exploit their current source of reward due to the risk of exploration[18]. While weighing their options, people may feel that trying new things leaves them vulnerable to mistakes, harsh judgment, or even danger. However, emotions also motivate exploratory behaviors in some cases. Findings from Kashdan and McKnight[9] suggest that happier moods are connected with explorative tendencies. Those with intrinsic interests and joy from a topic can directly contribute to a pattern of exploration in that field. This relationship between happiness and exploration is profound, as it not only initiates the journey into the unknown but also sustains it over time. When individuals find genuine delight and fulfillment in their exploratory pursuits, positive emotional reinforcement encourages them to delve deeper into their chosen paths of discovery. Once again, the interplay between behavioral and cognitive elements may be crucially moderated by affect, creating a synergy that collectively shapes our decision-making.

### *Clinical Implications*

The clinical impact of explore-exploit research promises to be extensive. For instance, work from Dombrovski and Hallquist[19] suggests increasing our understanding of the tradeoff between these behaviors could help lower suicide rates. A better understanding of how people experiencing suicidal ideation prepare themselves and their internal thought processes may initiate the development of better treatment plans or prevention programs. Additionally, research in this field has significant implications for the treatment of addictions and impulsivity. Bechara[20] suggests individuals struggling with addictions display higher impulsivity and lack of patience or control. In this work, Bechara[20] delves into the neural mechanisms underlying impulsive behavior, shedding light on cognitive processes (e.g., executive functioning) that contribute to exploitative tendencies in addicted individuals. By emphasizing the role of impaired executive function and willpower, Bechara's work[20] highlights how addiction can disrupt the balance between exploration and exploitation in decision making. If able to disrupt this pathway of excessive exploitation to addiction, researchers will be able to provide a significant impact on the addiction world, whether it is gambling, substance abuse, or other forms of addiction. In conclusion, if researchers are able to continue uncovering the intricacies of explore-exploit decision making, they can continue paving the way for stronger clinical treatments and resources for complex conditions.

### ***Future Directions***

An integrative approach to studying explore-exploit decision making has been integral to advancing our mechanistic understanding of explore-exploit decision-making, although significant bounds are yet to be made. Even with new technology and methods, many findings are correlational, and little work has attempted to delineate causal relationships between brain processes and explore decisions in humans. In one of a few examples, Raja Beharelle and colleagues[5] suggest that administering anodal transcranial DCT to the frontopolar cortex led to slower, exploratory-type decisions. Further, a cathodal transcranial DCT led to quicker, exploitative-type decisions; indeed, these findings are not well replicated. Nonetheless, more studies looking to draw causal conclusions will undoubtedly fill a major gap in the current literature. Trials like these, ones that manipulate biological variables, are rarely seen far and wide for a reason. To manipulate any biological aspect, the possible benefit of this change must significantly outweigh the risk that the participants are undertaking.

As we continue to advance how we assess explore-exploit pathways, more real-life impact opportunities will appear. For example, the technological advancement that has accompanied this research is especially exciting. With advancements in the field, countless different technological products could be considerably improved. For example, artificial intelligence and machine learning models will be able to quantify human data at a significantly higher rate, understanding their consumers better and improving efficiency. Further, in the realm of content recommendation algorithms, a better understanding of explore-exploit dynamics may yield more personalized and engaging recommendations for users, boosting user satisfaction and platform engagement, while also potentially revolutionizing the way we consume and interact with digital content[21]. In economic decisions, these insights could lead to more efficient resource allocation and investment strategies, ultimately enhancing economic growth and bolstering financial stability.

## Conclusion

The current review offers a brief examination of the complex dynamics underlying exploration and exploitation decisions. Paradigms informed by translational science have been useful for testing explore-exploit tendencies in humans, however, we should consider improving these paradigms to increase ecological validity. Advanced methods from psychology and neuroscience have spearheaded the movement towards a mechanistic understanding of exploration and exploitation, with particular focus on the ventromedial prefrontal cortex (vmPFC) during exploitative behaviors and the frontopolar cortex during more exploratory decisions. Exploring the cognitive and emotional aspects of explore-exploit decision-making has also shown the influence of factors like risk assessment, memory, and emotions on our choices. We now have a mechanistic understanding of how these processes affect people's decision-making as a result of this lengthy examination. Our field has surpassed a mere theoretical perspective of explore-exploit decision-making; advances will improve operations across several disciplines, including economics, artificial intelligence, and clinical psychology.

## Acknowledgments

I would like to thank Polygence for providing me with a platform to connect with mentors and offering resources. I would also like to thank my mentor, Phil Newsome, for his unwavering support in the completion of this project.

## References

1. Addicott, M., Pearson, J., Sweitzer, M., et al. (2017). A Primer on Foraging and the Explore/Exploit Trade-Off for Psychiatry Research. Neuropsychopharmacology, 42, 1931-1939. DOI: 10.1038/npp.2017.108.
2. Letkiewicz, W.A.M., Kottler, H.C., Shankman, S.A., & Cochran, A.L. (2023). Quantifying aberrant approach-avoidance conflict in psychopathology: A review of computational approaches. Neuroscience & Biobehavioral Reviews, 147, 105103. DOI: 10.1016/j.neubiorev.2023.105103.
3. Linson, A., Parr, T., & Friston, K.J. (2010). Active inference, stressors, and psychological trauma: A neuroethological model of (mal)adaptive explore-exploit dynamics in ecological context. Behavioural Brain Research, 380, 112421. DOI: 10.1016/j.bbr.2019.112421.
4. Roselli, L.R., Almeida, A.T., & Frej, E.A. (2019). Decision neuroscience for improving data visualization of decision support in the FITradeoff method. Operations Research, 1-21.
5. Raja Beharelle, A., Polanía, R., Hare, T.A., & Ruff, C.C. (2015). Transcranial Stimulation over Frontopolar Cortex Elucidates the Choice Attributes and Neural Mechanisms Used to Resolve Exploration–Exploitation Trade-Offs. Journal of Neuroscience, 35, 14544-14556.
6. Boldt, A., Blundell, C., & De Martino, B. (2019). Confidence Modulates Exploration and Exploitation in Value-Based Learning. Neuroscience of Consciousness.
7. Blanchard, T.C., & Gershman, S.J. (2017). Pure Correlates of Exploration and Exploitation in the Human Brain. Cognitive, Affective, and Behavioral Neuroscience, 18, 117-126.

8. Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the exploration–exploitation dilemma: An fMRI study of attention control and decision-making performance. Southern Medical Journal, 36, 319-338.
9. Kashdan, T.B., & McKnight, P.E. (2013). Commitment to a Purpose in Life: An Antidote to the Suffering by Individuals with Social Anxiety Disorder. Emotion, 13(6), 1150-1159.
10. Bari, A., & Robbins, T.W. (2013). Inhibition and Impulsivity: Behavioral and Neural Basis of Response Control. Progress in Neurobiology, 108, 44-79.
11. Nunez, P.L., & Srinivasan, R. (2006). Electric Fields of the Brain: The Neurophysics of EEG. Oxford University Press.
12. Costa, V.D., & Averbeck, B.B. (2020). Primate Orbitofrontal Cortex Codes Information Relevant for Managing Explore–Exploit Tradeoffs. Journal of Neuroscience, 40, 2553-2561.
13. Logothetis, N.K., & Wandell, B.A. (2004). Interpreting the BOLD Signal. Annual Review of Physiology, 66, 735-769.
14. Poldrack, R.A., Mumford, J.A., & Nichols, T.E. (2011). Handbook of Functional MRI Data Analysis. Cambridge University Press.
15. Draganski, B., Gaser, C., Busch, V., Schuierer, G., Bogdahn, U., & May, A. (2004). Neuroplasticity: Changes in Grey Matter Induced by Training. Nature, 427(6972), 311-312.
16. Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the Exploration–Exploitation Dilemma: An fMRI Study of Attention Control and Decision-Making Performance. Southern Medical Journal, 36, 319-338.
17. Daw, N., O'Doherty, J., Dayan, P., Seymour, B., & Dolan, J.R. (2006). Cortical Substrates for Exploratory Decisions in Humans. Nature, 441, 876-879. DOI: 10.1038/nature04766.
18. Aberg, K.C., Toren, I., & Paz, R. (2022). Irrelevant Threats Linger and Affect Behavior in High Anxiety. Journal of Neuroscience, 43, 656-671.
19. Dombrovski, A. Y., & Hallquist, M. N. (2017). The Decision Neuroscience Perspective on Suicidal Behavior: Evidence and Hypotheses. Current Opinion in Psychiatry, 30(1), 7-14. DOI: 10.1097/YCO.0000000000000297.
20. Bechara, A. (2005). Decision Making, Impulse Control, and Loss of Willpower to Resist Drugs: A Neurocognitive Perspective. Nature Neuroscience, 8(11), 1458-1463. DOI: 10.1038/nn1584.
21. Bronzin, T., Prole, B., Stipic, A., & Pap, K. (2021). Artificial Intelligence (AI) Brings Enhanced Personalized User Experience. 44th International Convention on Information and Communication Technology (MIPRO), 1070-1075.