Human Factors in Cybersecurity, Vol. 127, 2024, 94–104
https://doi.org/10.54941/ahfe1004769

Mental Firewall Breached: Leveraging Cognitive Biases for Enhanced Cybersecurity

Rebecca L. Pharmer[1], Rosa M. Martey[1], Giovanna L. Henery[1], Ethan Myers[1], Indrakshi Ray[1], and Benjamin A. Clegg[2]

[1]Colorado State University, Fort Collins, CO, USA
[2]Montana State University, Boseman, MT, USA

ABSTRACT

In an effort to examine whether cognitive biases could impact a cyber-attacker's decision making, we created a simplified, game-like task in which participants stole money from vulnerable banking applications. We manipulated both information in pre-task instruction, as well as the values associated with the bank accounts, as two methods to trigger anchoring biases, and included a selection of operating system properties to attack to capture asymmetric dominance effects. We also examined whether choices in this context are related to risk propensity in other tasks. Our findings suggest that cognitive biases can influence decision making in this cyber-attack task, but their effects are mitigated when multiple biases are manipulated concurrently. Potential implications of these findings for enhancing cybersecurity are discussed.

Keywords: Cyber defence, Oppositional human factors, Cognitive bias, Decision making

INTRODUCTION

As the digital landscape evolves, we are increasingly relying on safeguarding our private information. Defensive cybersecurity measures are employed in all sectors, from corporations that have a duty to protect their customers purchasing data, governments that need to ensure sensitive information remains guarded, and social media users whose identities could be stolen. This is done with the aim of countering efforts from individuals and organizations launching cyber-attacks to fulfil a motive of money, status, or blackmail. In a landscape of ever-evolving technologies and tactics, one potential source of consistency could be the general tendencies of the cyber attackers, and identifying methods to leverage these may be one crucial approach to negating the effects of an attack or stopping it altogether. This paper presents results of an experiment conducted using an online, game-like platform in which participants were instructed on how to conduct a series of "attacks" and make decisions about techniques and targets designed to evoke specific cognitive biases and influence decision-making in a manner comparable to cyber-security contexts.

Oppositional Human Factors

Human Factors typically seeks to supply principles to inform system design with the intent of optimizing human performance capabilities. In contrast, the notion of O*ppositional Human Factors* (Gutzwiller et al., 2018) can be thought of as the inverse: applying what we know from Human Factors to intentionally impair performance and usability with the goal of negating or hindering the actions of prospective attackers. As one example, Ferguson-Walter and colleagues (2021) conducted surveys with red team hackers, finding that characteristics of the target system could induce negative affective states in the attackers like frustration, self-doubt, and confusion – potentially reducing the success of cyber-attacks.

A Role for Cognitive Biases

Humans reflexively use judgmental heuristics to make rapid decisions by using intuition and prior information as a rule of thumb, especially in cases where their cognitive resources are being overloaded or they are influenced by a time constraint (Goldstein & Gigerenzer, 2008). However, these "mental shortcuts" can lead to systemic errors in judgment, referred to as cognitive biases, and result in suboptimal outcomes in performance (Tversky & Kahneman, 1974). There is some debate on the number of standalone biases that describe a unique error in decision making, and more research is needed to identify how these biases interact and which overlap substantially (Ellis, 2018; Johnson et al., 2021). Nonetheless, there seems to be a clear potential to deploy manipulations of cognitive biases as part of an oppositional human factors approach. For the scope of this paper, we selected two biases that we considered to be highly relevant in cybersecurity contexts: *anchoring* and *asymmetric dominance*.

Anchoring. The anchoring and adjustment effect describes how people making decisions can show a tendency to be overly influenced by the initial piece of information they are given or think of (i.e., the "anchor"), even when it is irrelevant or misleading, and fail to update their thinking and understanding as new information becomes available (see Furnham & Chu Boo, 2011 for a review). Researchers have also shown an effect of anchoring in general knowledge questions where participants unconsciously use their prior knowledge of things such as average temperature in a specific location or a significant date in history to serve as an anchor or reference point (Tversky & Kahneman, 1974; Epley & Gilovich, 2005). Other literature describes the anchoring bias influencing estimations of probability (Chapman & Johnson, 1999), purchasing decisions (Ariely & Simonson, 2003). For example, a car salesman will often show a very expensive car to a buyer first, to make a less expensive car seem like a better deal, even if the price of the cheaper car is inflated for the market value.

Because prior research has clearly established that information presented initially should have an observable impact on subsequent decision making across several domains, in the current study we explore two critical questions: The first is whether it is possible to have two separate anchors operating at different levels - one on a global level that is generated by a manipulation in

task instructions and encountered only once; and a second on a local level presented as part of the first information encountered whenever a round of attacks was conducted. Our second set of questions concerns the time course of these anchors, and how the repeated exposure to information that can drive adjustment from an anchor impacts the level of cognitive bias displayed. The conditions therefore enabled us to examine the extent and durability of the anchors, and to explore whether two different types of anchors can have influence simultaneously and perhaps even additively.

Asymmetric dominance. The asymmetric dominance effect, sometimes referred to as the "decoy effect", is a cognitive bias that occurs when a person is faced with a choice and another option is presented to change their assessment of the original options (Huber et al., 1982). This is often used as a sales tactic to entice customers to spend more money. Take the example of popcorn at the cinema: To increase the likelihood people will choose the large rather than small popcorn servings, they offer a third "medium" size that is close in price to the large but has much less popcorn, making the larger size appear relatively more attractive and a better value for money. Much of the research surrounding the asymmetric dominance effect stems from the field of behavioral economics. Huber et al. (1982) examined this bias by asking participants to make a selection between restaurants with different ratings and locations at varying distances. They found participants were more less likely to settle for a close by, poorly rated restaurant over its high-end counterpart when a mid-tier restaurant was introduced. The asymmetric dominance effect can be applied to cyber-attack contexts when attackers make judgements about which servers to attack, perhaps being influenced by the security measures of each individual server or machine. Indeed, decoy systems are often used as a security measure to entice attackers with false hopes of deceptively easy targets containing misleading information or traps for detection.

Risk-taking. A key element in real cyber-attack scenarios is the risk an attacker incurs when performing illicit hacking behavior. The concept of risky behavior is broadly defined in psychology as behaviors that involve the potential for harm or danger while also providing an opportunity for reward (Leigh, 1999). In the case of cyber-attackers, it is crucial to remain undetected to avoid consequences ranging from legal punishment to loss of access to a system. However, the extent to which general tendencies around risk taking influence the choices made during cyber scenarios remains an open question that we intend to explore in the current study.

The Balloon Analogue Risk Task (BART; Lejuez et al., 2002) captures risk-taking behavior without relying on self-report measures. In the task, participants inflate an on-screen balloon by pumping it with air as many times as possible before it unexpectedly pops. The core mechanism of the task is that the user is rewarded for risky behavior until the point where things become "too risky," and instead of maximizing their earnings, they lose them completely. Decisions about how much to inflate the balloon in each round reflect participants' propensity to engage in risky behavior. The current experiment features a somewhat similar series of decisions where participants can opt for higher risk choices, and we sought to examine whether these decisions were related to the BART scores of participants.

The Current Study. To examine the role of biases as a tool for oppositional cyber security, we began by designing an experiment focused on the two aforementioned biases, *anchoring* and the *asymmetric dominance effect* because of their demonstrated robust effect on decision making in the literature and potential relevance within the domain. We created an experimental platform to imitate a vulnerable banking website where we manipulated the account values presented to the participant, their perceptions of the probability of detection, and their choices for "attacking" the accounts to retrieve the money.

METHODS

Participants

196 undergraduate participants were recruited from an introductory psychology course. They received course credit for their participation. The study took place online, accessed through the participant's own computer (tablet and phone access were not permitted) and lasted approximately 45 minutes. Incomplete data was recorded on approximately 113 of 21,480 trials (less than 1%).

Procedure

CyberReaDec. The study took place online using the "CyberReaDec" (Cyber Reasoning and Decision Making) app developed by our team. This app consists of a set of tasks displayed with graphics and text instructions that lead participants to make a series of decisions to attempt to steal money from a bank. Participants were told that money had originally been stolen from a children's hospital by hackers and their task was to retrieve it. They conducted 10 attacks rounds (blocks) each comprised of viewing a maximum of 10 different bank accounts (trials). During each block, after accessing individual accounts, participants had to determine whether to take the funds ("steal") or move on the next account without taking the funds ("skip"). They were informed that stealing from accounts increased the probability of detection, but they were not provided with exact probabilities. The probability of detection increased cumulatively across each trial of the block. Electing to "steal" from a given account increased the chances they would be caught in that block by 5%, while selecting "skip" increased the chances by 1%. On each trial, if a random number generated was lower than their current detection probability then they were caught, and that attack would end, and they would lose all the money they had stolen from that round. If they successfully got past the tenth account without detection, they secured all the funds stolen during that attack. When they were either caught or got through ten accounts without capture, the block ended, and the probability of detection was reset to zero for the start of the next block. Therefore, participants were required to select *skip* or *steal* from each account within a block - they could not leave the block. They would see 10 accounts if they were not caught but fewer if they were caught. Thus, participants were exposed to a maximum

of 100 accounts across 10 rounds of attacks; and a minimum of 50 accounts across 10 rounds.

To ensure that participants were equally exposed to the same average account values in the *account anchor* condition (see below), the participants could not be caught in the first 4 trials of a block (although they were not informed of this). The steal/skip choices participants made contributed to the weight of the probability of detection for each block from trial 5 onwards. Once they finished Phase 1 of the experiment, they continued to Phase 2 where they completed the BART Balloon task.

The Balloon Analogue Risk Task (BART). The BART (Lejuez et al., 2002) involves a computerized task where participants were presented with a series of balloons on a computer screen. Participants were instructed to click a button to inflate the balloon, which increases its size and the associated monetary reward. However, the balloon carried a risk of exploding at any point, resulting in the loss of the accumulated reward for that balloon. Participants decide when to stop inflating the balloon and collect the reward before it explodes. The key objective of the BART is to examine the risk-taking tendencies of individuals and their willingness to take risks for potential rewards.

Design

The current experiment utilized a 2 (Instructions Anchor: Vague or Specific) x2 (Asymmetric Dominance: Asymmetry or No-asymmetry) x3 (Account Anchor: High, Standard, or Low) between-subjects design.

Instructions Anchor. The instructions presented to participants manipulated an anchor with respect to the number of accounts accessed in relation to the risk of detection. The anchor was presented only in the initial experiment instructions. The text used for each condition either: 1) Specific ["Stealing money from 3 or more accounts in a single attack round has a very high risk of detection"] or 2) Vague ["Stealing money from too many accounts in a single attack round has a very high risk of detection"].

Asymmetric dominance effect. To examine an asymmetric dominance effect, participants had to choose which operating system (OS) to attack (see Figure 2) at the end of the initial experiment instructions. This selection did not actually affect the chances participants were caught in the subsequent attacks, although participants were not informed about the relationship of any of their choices to the chances of being caught.

The OS options varied system age and recency of updates, and participants were informed older and less updated systems were more vulnerable to exploits, but this was balanced against the possible presence of 'honeypots' on the network ("fake server intentionally set up to appear weak" and resulting in a high probability of detection). Within this choice we manipulated the presence and absence of an asymmetric option to explore whether the presence of a third option impacted the decision between the other two. In this case, the addition of better candidate honeypot (OS7) should direct more people towards away from the most secure server (OS11) because the now intermediate option (OS10) is seen as lower relative risk. In practical terms this might equate to putting a very obvious honeypot alongside a less obvious

honeypot to push people away from a system intended to be protected. The two conditions were:

No asymmetry (two possible server options):

- Windows 10: "moderate security vulnerabilities."
- Windows 11: "minimal security vulnerabilities."

Asymmetry present (three possible server options):

- Windows 7: "some security vulnerabilities"
- Windows 10: "moderate security vulnerabilities."
- Windows 11: "minimal security vulnerabilities."

Account Anchor. A second possible anchor was manipulated in the initial account values participants saw on the first trial of every block. Participants were assigned to conditions such that the same anchor type always occurred for them on the first trial. To ensure that decisions across the block were based on similar opportunities to acquire equivalent funds, each first account anchor was matched with its opposite. Therefore, participants who repeatedly saw low value accounts on trial 1 then always saw high value accounts on trial 4, and vice versa. All other accounts, including all those used for participants in unanchored conditions, were drawn from a standard range. An average across the high and low account accounts was set to be equivalent to the mean of the standard accounts [across first 4 trials all conditions have encountered a mean account value of $110,000].

Table 1. Account anchor sample values.

Trial	High Anchor		Low Anchor	
	Value	*Type*	*Value*	*Type*
1	$214,947.60	High	$4,689.93	Low
2	$102,572.60	Standard	$102,572.60	Standard
3	$152,805.80	Standard	$152,805.80	Standard
4	$4,689.93	Low	$214,947.60	High

Account values were randomly generated by the app within a specified range of possible amounts uniquely for each participant. Values were designed to generate approximately the same average total value across the 10 trials within each round, regardless of condition, approximately an average of $110,000. Table 1 provides an example of values by condition for the first 4 trials of each block.

High Anchor (high account value): $215,000 (+/- a random value between $0 and $500) in 1st account; a random value between $20,000 and $200,000 for accounts 2-3; $5,000 (+/- a random value between $0 and $500) in 4th account. Accounts 5 – 10 a random value between $20,000 and $200,000.

Low Anchor (small account value): $5,000 (+/- a random value between $0 and $500) in 1st account; a random value between $20,000 and $200,000 for accounts 2-3; $215,000 (+/- a random value between $0 and $500) in 4th account. Accounts 5 – 10 a random value between $20,000 and $200,000.

Standard (No Anchor - moderate account value): All accounts a random value between $20,000 and $200,000.

RESULTS

Participants took an average of 2.11 seconds to make their decision on each trial. They elected to steal 38.7% of the accounts encountered. Across the 10 blocks they were caught an average of 6.73 times. To examine the effects of the two types of anchors 2 (Instructional Anchor: specific or vague) by 3 (Account Anchor: high, low, or standard) ANOVAs were conducted across all participants, regardless of asymmetric dominance condition.

Anchoring: Total Times Caught. Looking at the metric of total number of times participants were caught across the ten blocks in the experiment, there was a marginally non-significant effect of account anchor ($F(2, 190) = 2.29$, $p = 0.10$, $\eta_p^2 = 0.02$; $M_{low} = 7.1$, $M_{standard} = 6.5$, $M_{high} = 6.6$); no effect of instruction anchor ($F(1, 190) < 1$; $M_{vague} = 6.6$, $M_{specific} = 6.9$); and no account anchor by instruction anchor interaction ($F(2, 190) = 1.42$, $p > 0.10$, $\eta_p^2 = 0.01$).

Despite the absence of an interaction, to assess whether account anchors alone were effective in influencing the propensity to take actions leading participants to get caught, we next examined their effects in the absence of a second (prior) anchor in the instructions. Limiting the analyses to just the participants who were given the vague (no anchor) instruction now produced a main effect of account anchors ($F(2, 104) = 4.04$, $p = 0.02$, $\eta_p^2 = .07$; $M_{low} = 7.3$, $M_{standard} = 6.1$, $M_{high} = 6.4$). In contrast, when looking only at participants whose instructions did provide a specific anchor, there was no hint of an effect of subsequent account anchors ($F(2, 86) < 1$; $M_{low} = 6.9$, $M_{standard} = 6.8$, $M_{high} = 6.8$).

In examining the effects of the instruction anchor in the absence of the account anchors (i.e., limited to just the standard account condition participants), there was no effect of the initial instruction manipulations across the experiment ($F(1, 60) = 2.07$, $p = 0.16$, $\eta_p^2 = .03$; $M_{vague} = 6.8$, $M_{specific} = 6.2$).

Combined these findings suggest that it is possible to exert control over the initial information present to set an anchor for subsequent decisions, but the results also hint that such effects may be diminished where other anchors are present.

Anchoring: Account Values. Turning from the global properties to the more local, individual decisions, we next examined the value of accounts selected. Across the experiment for the average value of accounts that were chosen to steal from, there was no effect of account anchor ($F(2, 190) < 1$; $M_{low} = \$139{,}926$, $M_{standard} = \$134{,}911$, $M_{high} = \$136{,}758$); no effect of instruction anchor ($F(1, 190) < 1$; $M_{vague} = \$137{,}668$, $M_{specific} = \$136{,}729$); and no account anchor by instruction anchor interaction ($F(2, 190) = 1.56$, $p > 0.10$, $\eta_p^2 = 0.02$).

One possibility is that the anchoring effects presented only in the initial instructions were not durable across the repeated exposures to accounts and attack outcomes. To explore this possibility, we repeated the previous analysis looking only at the 1st block of trials. There was a main

effect of account anchor ($F(2, 190) = 4.66$, $p = 0.01$, $\eta_p^2 = 0.05$; $M_{low} = \$107{,}864$, $M_{standard} = \$106{,}198$, $M_{high} = \$128{,}354$); an effect of instruction anchor ($F(1, 190) = 7.82$, $p < 0.01$, $\eta_p^2 = .04$; $M_{vague} = \$123{,}464$, $M_{specific} = \$104{,}812$); and a marginally non-significant account anchor by instruction anchor interaction ($F(2, 190) = 2.34$, $p = 0.10$, $\eta_p^2 = 0.02$)

Once again limiting the analyses to just the participants with a single anchor (i.e., the vague condition for instructions), the 1st block showed a large effect of account anchors ($F(2, 104) = 8.23$, $p < 0.001$, $\eta_p^2 = 0.14$; $M_{low} = \$109{,}858$, $M_{standard} = \$113{,}270$, $M_{high} = \$147{,}265$). Again, when another anchor was also present (i.e., the specific instructions), there was no hint of an effect of subsequent account anchors ($F(2, 86) < 1$; $M_{low} = \$105{,}869$, $M_{standard} = \$99{,}126$, $M_{high} = \$109{,}443$).

The lack of durability of these anchoring effects is reflected when the analyses explored performance on the 2nd block of trials. For average account values stolen, there was now no main effect of account anchor ($F(2,190) < 1$); no effect of instruction anchor ($F(1, 190) = 1.84$, $p > 0.10$, $\eta_p^2 = 0.01$); and no account anchor by instruction anchor interaction ($F(2, 190) < 1$).

Overall, these findings are generally consistent with the literature, suggesting the potential to use anchors within these contexts. However, there is no evidence that there was a layering of multiple, concurrent anchors. Moreover, these data suggest that adjustment occurs somewhat rapidly (here, effects are confined to only the initial block), and that repetition of the anchors in later blocks did not continue to re-establish that anchor as a foundation for comparisons in subsequent decisions.

Asymmetric dominance effect. In examining the percentage of participants who selected each operating system, the overall preference for which systems to attack showed in the no asymmetric dominance condition a preference for the more vulnerable system (OS11: 45%, OS10: 55%), which would be consistent with a high proportion of participants ignoring the honeypot risk noted in the instructions. The introduction of the asymmetric in the three OS options condition (OS7) had little impact on moving people away from the most secure system (OS11: 41%, OS10: 42%).

However, there was evidence that the instructions condition of the anchoring manipulation (despite being an anchoring-based manipulation) influenced the impact of the asymmetric dominance effect. For those in the specific instructions condition, in the asymmetry absent case there was a tendency to avoid the potential honeypot (OS11: 57%, OS10: 43%), but, as predicted by the asymmetric dominance effect, this changed once a better candidate honeypot was also present (OS11: 35%, OS10: 40%). In other words, when the instructions anchor was specific (rather than vague), the asymmetric dominance effect did show the predicted effect on participant selection of OS. We will return in the discussion section below to the question of why the presence of specific information in the initial instructions might have played an important role in whether the asymmetric dominance effect influenced choices or not.

Risk Taking. Post-session surveys examined participant propensity in risk-taking. Overall, the relationship between the total number of times caught

and BART performance showed only a weak, non-significant correlation ($r = 0.12$; $p = 0.50$). However, when broken down by the instruction conditions, for the specific anchor condition there was some sign of a stronger correlation between times caught and BART performance ($r = 0.44$; $p = 0.06$), whereas for the vague conditions the relationship was negligible ($r = -0.19$; $p = 0.42$). This suggests there may be some potential role that risk taking propensity plays in the likelihood individuals would move beyond the conservative approach that the specific anchor would supply.

DISCUSSION

The current study examined whether cognitive biases can be observed operating in a simulated cyber-attack scenario. To understand how multiple biases can be manipulated concurrently, we created a cyber-attack task in which participants were exposed to an anchor in the pre-task instructions, another anchor in the form of a first account value in each block, and a selection among either two or three operating systems to serve as our asymmetric dominance effect manipulation. We sought to understand whether manipulation of the initial information encountered can impact the choices made, and more specifically lead to decisions that increase the possibility of detection.

We found that providing an instructional anchor at the start of the session did not affect the number of times participants were caught, nor did it influence the average amount of money they stole from accounts. There was only some evidence for a very short-term influence on the first encounter with the system. This implies that this type of anchoring may not be optimal for influencing decisions that could increase the possibility of detection, given its modest and transient impact. We did find evidence to support the impact of account value anchors in both the average amount of money stolen and the number of times participants got caught. However, these behaviors appear only significantly impacted by the account value anchor in the absence of an instructional anchor. This finding suggests that concurrently manipulating these biases may serve to wash out the effectiveness of each. Moreover, the effects of account value anchors were found in blocks presented early in the process, but ongoing exposure to more relevant information reduced the impact of the anchoring bias on subsequent decision making. Hence although anchors presented at the start of a process can be influential, individuals begin to use more relevant and specific information somewhat quickly as a task progresses to adjust, and do not repeatedly re-anchor on the initial information encountered.

The asymmetric dominance effect manipulation was at best only partially successful. Asymmetric dominance effects were only found in the conditions that were given a specific instructional anchor as part of the anchoring manipulation. There is no inherent reason to believe in this case that manipulating biases concurrently is impacting the effectiveness of each manipulation. Rather, this might reflect that the order of the instructional content and attack selection played a role in attention. Our post-hoc explanation would be that participants assigned to conditions without a specific anchor in the instructions may have inferred from the lack of concrete information that

the details were largely unimportant, and hence adopted a strategy of paying less attention to the other information. Later, in making a decision about which operating system to attack, those in the vague instructions condition may have sought less understanding of the goals, risks, and/or parameters of their tasks. In contrast, having previously encountered valuable information, participants in the specific anchor condition may have paid more attention to the options they were provided in the system choice, and thus obtained the required understanding of risk and trade-offs of vulnerability and possibilities of being a honeypot in each OS for the asymmetric dominance effect to occur. The potential role of attention to information in asymmetric decisions should be addressed in future research.

Analysis of participants behavior on the BART test offered some support for the notion that the presence of specific information in a task could be related to risk-taking propensity. Because performance on the BART was most related to task decisions among those who received the specific information (in the instruction anchoring condition), it may be the case that additional detail about risk helps individuals calibrate their risk tolerance more effectively. That is, providing more specific information about the parameters in which failure occurs ("3 or more accounts") is related to how people make decisions about related tasks. This could indicate that a reduced tolerance for risk could be induced when a specific anchor is provided.

REFERENCES

Ariely, D., & Simonson, I. (2003). Buying, bidding, playing, or competing? Value assessment and decision dynamics in online auctions. *Journal of Consumer psychology*, *13*(1-2), 113–123.

Chapman, G. B., & Johnson, E. J. (1999). Anchoring, activation, and the construction of values. *Organizational Behavior and Human Decision Processes, 79*(2), 115–153.

Ellis, G. (2018). So, what are cognitive biases? In G. Ellis (Ed), *Cognitive biases in visualizations* (pp. 1–10). Springer.

Epley, N., & Gilovich, T. (2005). When effortful thinking influences judgmental anchoring: differential effects of forewarning and incentives on self-generated and externally provided anchors. *Journal of Behavioral Decision Making, 18*(3), 199–212.

Ferguson-Walter, K. J., Gutzwiller, R. S., Scott, D. D., & Johnson, C. J. (2021). Oppositional human factors in cybersecurity: A preliminary analysis of affective states. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering Workshops* (ASEW), pp. 153–158.

Furnham, A., & Boo, H. C. (2011). A literature review of the anchoring effect. *The Journal of Socio-Economics, 40*(1), 35–42.

Goldstein, D. G., & Gigerenzer, G. (2009). Fast and frugal forecasting. *International Journal of Forecasting, 25*(4), 760–772.

Gutzwiller, R., Ferguson-Walter, K., Fugate, S., & Rogers, A. (2018, September). "Oh, Look, A Butterfly!" A framework for distracting attackers to improve cyber defense. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting Vol. 62, No. 1* (pp. 272–276). Sage CA: Los Angeles, CA: SAGE Publications.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research, 9*(1), 90–98.

Johnson, C. K., Gutzwiller, R. S., Gervais, J., & Ferguson-Walter, K. J. (2021). Decision-Making Biases and Cyber Attackers. *36th IEEE/ACM International Conference on Automated Software Engineering Workshops (ASEW)*, Melbourne, Australia, 2021, pp. 140–144.

Leigh B. C. (1999). Peril, chance, adventure: concepts of risk, alcohol use and risky behavior in young adults. *Addiction, 94*(3), 371–383.

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L.,... & Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: the Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied, 8*(2), 75–84.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science, 185*(4157), 1124–1131.