# Investigation of Multiple Cognitive Biases in Military Contexts

Mark A. Livingston, Prithviraj Dasgupta, Jonathan W. Decker, and John Kliem
Information and Decision Science Branch
Naval Research Laboratory
Washington, DC, USA
e-mail: {mark.a.livingston18, prithviraj.dasgupta, jonathan.w.decker4, john.kliem3}.civ@us.navy.mil

*Abstract*—**We consider the problem of detecting cognitive biases in problems domains that are relevant to military personnel and roles they may have. In particular, we determined that anchoring bias, zero-risk bias, attraction effect, and compromise effect were relevant to military domains. In a user study, we hoped to elicit these biases and determine whether co-occurrence existed. We elicited anchoring bias in a time-extended task, but had limited success eliciting the other biases in small, disconnected scenarios. We did not observe co-occurrence of any of these four biases. We sought, but did not observe, whether visual presentation aids of text scenarios affected the presence of bias. We note some effects of user-identified strategies on biases.**

***Keywords- cognitive biases; anchoring bias; zero-risk bias; attraction effect; compromise effect; user study***

## I. INTRODUCTION

Decision-makers often rely on heuristic strategies, perhaps even without realizing it. These cognitive biases, or unstated bases for decisions, often lead to poor choices, including in military contexts [1], where poor decisions may lead to unnecessary loss of life. Cognitive biases in human decision-making while solving a problem are known to affect the outcome [2][3]. These biases usually degrade the outcome's value to the decision maker and to others that are affected by the problem's outcome. Researchers have proposed several techniques to detect biases.

Our main goal in this work was to develop scenarios, many based on military contexts, that would elicit hypothesized cognitive biases and determine what we could observe that may enable prediction of them. A second goal was to determine if these biases (when they occur) co-occur in individuals. Thus, we also collected information that we hoped would give predictive value for the existence or co-existence of our selected biases. Further, we hypothesize that visual presentation of information may mitigate bias, so we present our scenarios in text only and in text with illustrations of data presented in the text. On these last two questions, we hoped to make new research contributions.

The remainder of this paper is organized as follows. In Section II, we review literature in order to select biases to investigate and scenarios to further our investigation. We also describe differences our work introduces. In Section III, we describe the user study we used to gather data. In Section IV, we will present some data filtering we needed in order to conduct analysis. Our data analysis appears in Sections V and VI, reflecting the variety of forms of analysis we needed. Discussion appears in Section VII. We draw conclusions and recommend ideas for future work in Section VIII.

## II. SELECTION OF BIASES FOR INVESTIGATION

Dimara et al. [2] proposed a task-based taxonomy of 154 cognitive biases. We relied heavily on their taxonomy (mainly as expressed in their Table 2) to decide what biases we would try to elicit. Their list of tasks includes estimation, decision, hypothesis assessment, causal attribution, recall, opinion reporting, and other. We focused on the decision task, because we felt that was most relevant to the military domain that is our focus. They also used an intuitively-developed set of sub-categories as a second level in their taxonomy. This level consisted of *association* (cognition is biased by connections between items), *baseline* (cognition is biased by comparison with a baseline), *inertia* (cognition is biased by the prospect of changing the current state), *outcome* (cognition is biased by how well something fits a desired outcome), and *self perspective* (cognition is biased by a self-oriented view point).

Our initial idea was to examine biases made in the context of a strategy game, thinking this would be a good proxy for military tasks. This caused us to select one task from Dimara et al.'s [2] Estimation category, in the baseline sub-category: the anchoring effect. We felt that an initial solution (shown in a tutorial phase) to a simple game would elicit the effect. However, in order to isolate the biases from each other and be able to better control their elicitation, we opted for writing scenario sets, separately from the game, to elicit the other biases. We wanted something that we believed could be elicited and detected in a straightforward manner, but that would still be representative of decisions made in a variety of military contexts. This led us to select three other biases: the attraction effect (in the Decision/baseline portion of the taxonomy), the compromise effect (also in the Decision/baseline portion), and the zero-risk bias (in the Decision/association portion). The following subsections summarize research on these biases. In Section III, we describe how we created scenarios to (attempt to) elicit each of these biases.

### A. Anchoring Bias

Anchoring bias [3] causes humans to rely heavily on an initial piece of information, called an *anchor*. Because of this, humans tend to overlook information that would lead to better choices in subsequent decisions, and, instead, gravitate towards choices that align with the anchor. Initial research on analyzing anchoring biases focused on single-point decision problems. The main experimental design used for anchoring bias in such single-point decisions is the following: first, a decision maker is exposed to the anchor, about the likely outcome of a decision. Then, the decision maker is asked to make the same or a very similar decision. Anchoring bias is claimed to affect the latter decision if the latter decision's outcome is similar to the initial decision outcome. A canonical example is to anchor the decision maker to a price, e.g., 100 for a certain piece of clothing. Subsequently, the decision maker is shown a similar piece of clothing that is priced well below (or well above) 100, without revealing the price, and asked its worth. If the decision maker says that the clothing is worth around 100, it indicates that they are anchored to the initial price of 100.

Researchers [4][5] have reported the presence of anchoring bias in decision making for time-extended tasks (reviews of books and college applications). However, in these research studies, while making the decision for the current task the decision maker had access to the features of the current task, in addition to their experience from past decisions on similar tasks stored in their memories. In contrast, we ask 'If access to the current task's features while making the decision for the task were to be taken away, and the decision maker had to rely solely on experiences from memory from similar tasks to make decisions, is anchoring bias still present?' This question does not seem to have been investigated well in the literature.

These research settings are complementary to the research in this paper. The two main differences between our work and these are, first, we do not reveal the current problem's features (e.g., current book or college application under review) to the decision maker and the decision maker has to rely only on past task features and decisions from memory to make the current task's decision. Other slight distinctions are that these techniques use offline data that was not generated specifically for the bias studies and there was limited information about the background of the decision maker. On the other hand, the subjects in our study are people that were familiar with computer-game playing and decision-making in scenarios similar to our game. In addition, we report on a study that would have detected co-occurrence of other biases alongside an anchoring bias.

### B. Attraction Effect

Simonson [6] defines the *attraction effect* in a situation in which you have two alternatives and two dimensions that are important to your selection (Fig. 1). These two alternatives (Options 1 and 2 in Fig. 1) create a trade-off between the two dimensions of selection. In addition, a third choice (Option 3 in Fig. 1) has similar values to one of the first two alternatives, but is clearly weaker. This may imply weaker in both dimensions, or clearly weaker in one dimension while equal or even (very) slightly stronger in the other dimension [7][8]. A notional graph of the first form (weaker in both dimensions) appears in Fig. 1; Simonson [6] gives forms in which Option 3 is equal to or even slightly stronger than Option 2 in Selection Dimension 1. We adopt all three forms, and implemented some of each case, depending on whether we thought the form in which the inferior option was weaker would be believable at all. Graphs for these scenarios were much like Fig. 1, with notional labels, not specific values for axis labels.

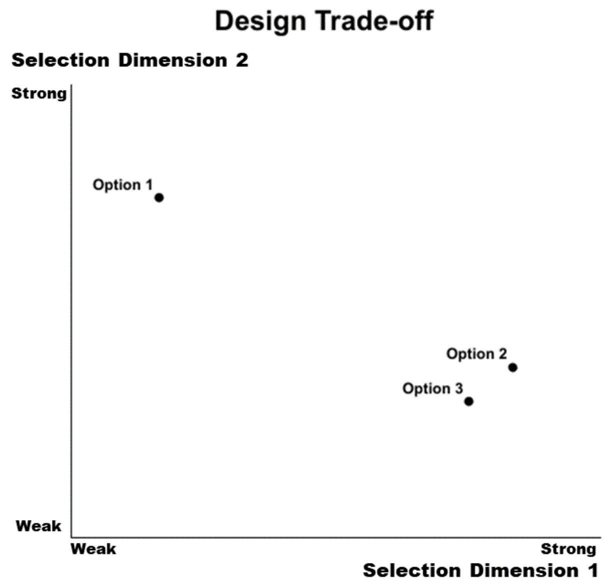


Figure 1. A notional graph showing the *attraction effect* occurring due to the relative placement of options along two dimensions of selection. Option 1 is sometimes known as the *target*; it is generally the best option or a competitor to the selection designer's choice. Option 3 should be seen as clearly inferior to Option 2, but even when Option 1 is objectively better than Option 2, the attraction effect leads a consumer to choose Option 2, because the inferior Option 3 attracts the user's attention.

### C. Compromise Effect

Simonson [6] also defines the *compromise effect* in a similar situation: you again have two alternatives and two dimensions that are important to your selection. These two alternatives (Options 1 and 2 in Fig. 2) create a trade-off between the two dimensions of selection. Again, there is an additional third choice (Option 3 in Fig. 2), but this time it is very near the mean values of the first two options in each dimension of selection. Hence, it can be seen as a good *compromise* between the two *extreme* alternatives. (These terms will be used to refer to choices in the analysis.)

Kivetz et al. [8] note that the compromise may be slightly better, exactly, or slightly worse than the average of the two extremes. (Respectively, Option 3 would be right of, on, or left of the line connecting Options 1 and 2, in Fig. 2.). We adopted the model of the compromise being slightly worse, because we felt it would more readily elicit the effect. In such a constructed scenario, no one should (rationally) choose the compromise without accepting an overall inferior choice. Again, the graph axis labels were relative terms; precise values were never used (as depicted in Fig. 2).

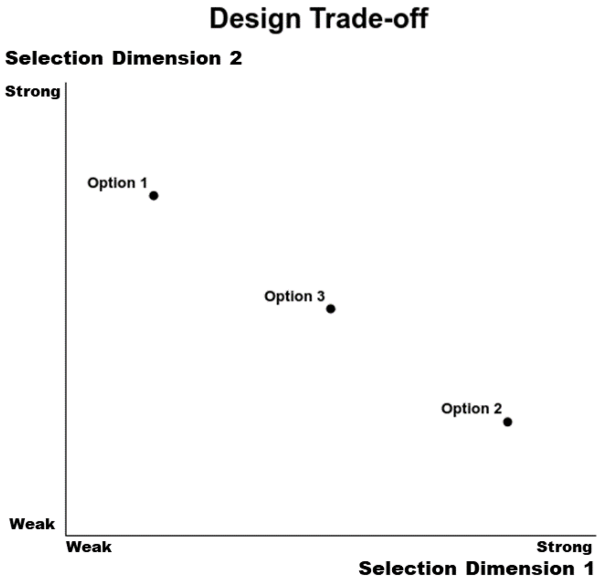


Figure 2. A notional graph showing the *compromise effect* occurring due to the relative placement of options along two dimensions of selection. Options 1 and 2 are seen in some sense as extreme, whereas Option 3 should be seen as good compromise. This leads a consumer to choose Option 3, even if it may not be an exact trade-off (i.e., it would not be on the line connecting Options 1 and 2).

### D. Zero-risk Bias

Baron et al. [9] define *zero-risk bias* as showing a preference for reducing a portion of risk to zero. They do this in the context of how to allocate resources to cleaning up environmentally contaminated sites. They presented three options for cleaning up both sites at varying levels. The authors defined zero-risk bias as expressing that the option that included a reduction to zero risk for one site (out of two) was better. In another version, the zero-risk option actually reduced the cancer cases by fewer incidences; zero-risk bias was defined as choosing this inferior option. This definition of zero-risk bias matches the earlier one of Viscusi et al. [10]. Their choices involved health risk from an existing product versus a new household product (which was presented as real but was only for purposes of the scenario). They found that consumers were willing to accept greater overall "cost" (by whatever metric was defined in the scenario) when the risk in one sub-part of the choice was zero. We followed this model in creating questions we hoped would elicit zero-risk bias.

## III. Data Collection

Under an IRB-approved protocol, we conducted data collection in collaboration with the Naval Aerospace Medical Institute (Pensacola, FL, USA). Various U.S. Navy and Marine personnel volunteered for our study during their free time. Volunteers completed the following steps, per our IRB-approved protocol: (1) informed consent, (2) demographics questionnaire, (3) Cognitive Reflection Test (CRT) [11], (4) Rational-Experiential Inventory (REI) [12], (5) Terrain Orientation Task [13], and then (6--9) four tests designed to elicit the biases described above. The order of these last four sections was determined using a 4x4 Latin square retrieved from an online Latin square generator [14]. This Latin square was counterbalanced for first-order sequence effects [15]. We elected to include the CRT and REI because Sleboda and Sokolowska [16] found each to be predictive of aspects of decision making. We wanted to test whether the Terrain Orientation Task would be predictive of any ability shown on the search-and-destroy tank game, described above.

Our pool of 90 volunteers was skewed to male (69, versus 20 female, with one declining to answer) and college-age or post-college age: 49 were ages 18-22, 34 were ages 23-27, and just seven were age 28 or older. Education level was asked via giving their highest academic degree; 32 said a high school diploma, three said an Associate's degree, 43 said a Bachelor's degree, and 12 said a Master's degree. All were fluent in English (the language of the text portions of the study); most were native speakers, but six identified another language as native, with 2-29 years of speaking English among those, one was native in English and a second language, and five declined to answer. We do not believe language was a barrier, as all participants were members of the U.S. military and thus communicate regularly in English.

### A. Details of the Anchoring Bias Data Collection

Computer-based games have been employed in education and cognitive analysis [17] as an enabler for humans to perform learning or decision-making tasks. Following this, we implemented a game for detecting anchoring bias in a sequential decision-making task. A game player must move a game piece in a grid-based 2D environment. At any point in the game, the player can see only a portion of the game board revealed via a circular viewport centered around the game piece's current location (Fig. 3, top; the red cluster of dots is the game piece). The environment contains objects called tanks that are placed in a cluster around a certain location in the environment. Fig. 3 (bottom) shows the tanks on the game board with the region outside the viewport grayed out for legibility. A tank can be removed or cleared by the player by pressing a key when the game-piece is in the vicinity of the tank. There is also an exit at a fixed location in the environment (elliptical pad on the right edge in Fig. 3, bottom). The exit can be seen only when it is in the player's viewport, but its location is known to the player from the start of the game. The player has two objectives: first, detect and clear all the tanks in the environment, second, after clearing all the tanks, exit the environment.

Due to the limited size of the viewport, a player cannot know beforehand where the tanks are located inside the environment. Consequently, they have to search the environment by moving around the game-piece. Once the tanks are visible inside the viewport, they can move the game-piece to each tank's vicinity, clear the tanks, and finally move to the exit. The game piece could be moved in only the four cardinal directions, Up, Down, Left, or Right. The game board was discretized into a grid-like environment for the purpose of tracking the game-piece's location. Fig. 3 (top) shows a screen capture of the game; Fig. 3 (bottom) shows the full game board (in faded colors) for illustration.

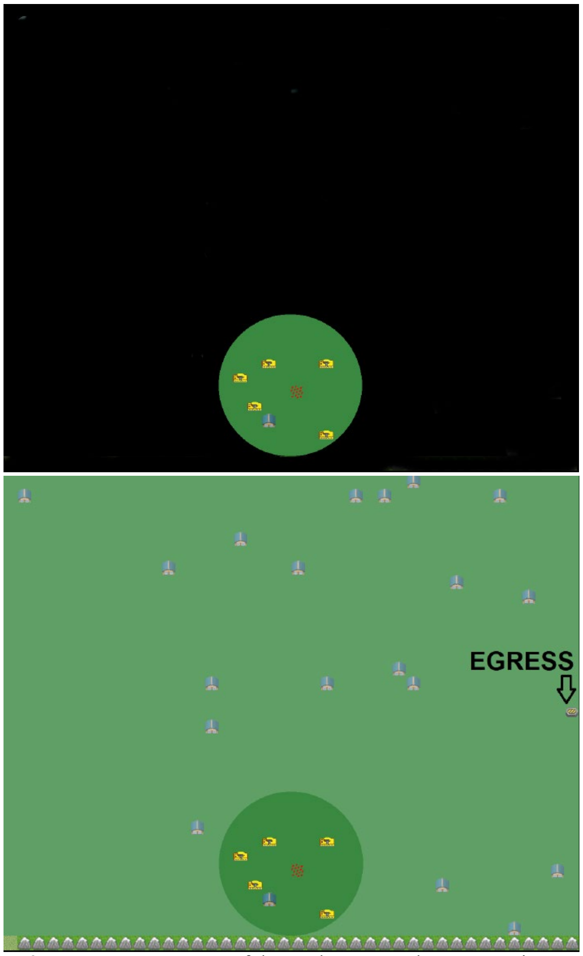


Figure 3. Top: Screen capture of the Tank game as the user saw it, a mostly black field with a viewport centered on the player's current location; the red cluster of dots at the middle of the viewport is the player's game-piece. Bottom: Tank game with grayed map outside viewport (for illustration).

We partitioned the environment into six equal-size cells (three division horizontally and two vertically); the game piece always began in the lower-left cell. The game had two phases. During the anchoring phase, all tanks were placed in a randomly-selected cell other than the lower-left cell. The player then played the game twice. The game would then silently (i.e., unbeknownst to the player) switch to the evaluation phase, where the player's movements would (or would not) indicate anchoring bias.

## B. *Details of the Attraction Effect Data Collection*

We wrote a mostly custom set of scenarios for our users, beginning with Simonson's consumer scenarios [6] as a basis. We drew inspiration for some scenarios from other studies of the attraction effect [7][18]. We introduced graphs that summarized certain aspects of each scenario; these were shown in the second half of the study for each type of bias. The graphs for the attraction effect clearly showed the closeness of the attractor and the decoy, whereas the target was near the opposite corner of the graph. Since the order of scenarios was determined by a Latin square, all scenarios had a graph available for us to show, and the question counter determined when to show graphs. So, all scenarios were presented with and without graphs across our participants.

## C. *Details of the Compromise Effect Data Collection*

We again wrote a custom set of scenarios for our users, based largely on Simonson's scenarios [6], but also drawing inspiration from other studies of consumer choice [8][19] and other studies of decision-making [20][21]. As with the attraction effect, we introduced graphs that summarized the compromise being made, showing the not-quite-linear relationship of the three choices, with the one in the middle deviating slightly from this relationship in the direction of the slightly worse according to Kivetz et al.'s [8] description. Again, a Latin square determined the order of scenarios, and the question counter showed graphs in the second half of this portion of the study.

## D. *Details of the Zero-risk Bias Data Collection*

As with the attraction effect and compromise effect, we wrote custom scenarios for our users, drawing from previous studies of zero-risk bias [9][22][23] or decision-making [20][24]. We again created graphs which used clustered (pairs of) bars to illustrate the choice to be made. Participants needed to sum the length of bars in a cluster in order to compare to the single bar for the zero-risk option. We chose not to use a stacked bar, in part because they are generally more confusing to readers [25], and in part because we felt that the stacking would make the inferiority of the zero-risk option too obvious. (The stacked bars would take the place of having to sum the component results.) As with the attraction and compromise effects, the order of scenarios was determined by Latin square and the graphs shown once the question counter reached the second half of this portion of the study.

## E. *Further Details of the Text-based Data Collections*

As noted above, each of the tests for the attraction effect, compromise effect, and zero-risk bias were divided in halves. The first half in each showed a text version of the scenario and response options. The second half had this text (scenario and response options) as well as a graph that illustrated what we viewed as the critical data on which participants would want to make their decision for that scenario. At the end of each half, participants were asked if they had any strategy on that section; portions that displayed graphs prompted specifically to indicate whether the participant felt the graph was helpful in that section. On this free text response, we assessed the sentiment towards the graph. The response was categorized as *positive* if the participant indicated using the graph or finding it helpful. The response was categorized as *negative* if the participant indicated ignoring the graph or finding it confusing or otherwise unhelpful. Responses that did not mention the graph at all were categorized as *neutral*. Sentiment will be used as an independent variable in our analyses of these three biases.

We further evaluated these free text responses to determine if the participants indicated a particular strategy to select their response to the scenarios in that portion of the study (nine-question blocks). We identified keywords in the response, and then we coalesced these keywords into five categories; we then added another category for responses that explicitly mentioned keywords from multiple of the other categories. This process yielded the following assessment of strategies (with descriptions of responses that fit into them):

- *None*: Participants expressed clearly that they had "no strategy" or that they "didn't use a strategy"
- *Intuition*: Participants said they "went with their gut reaction" or chose "what felt right"
- *Feature*: Participants said they chose based on a particular feature (that changed with each scenario); examples include least loss of life, lowest value of equipment lost, and lowest cost of devices.
- *Efficient*: Participants said they were aiming for efficiency or the best use of resources
- *Balance*: Participants said they tried to "balance" the competing interests or chose a "middle" option
- *Mixed*: Participants explicitly used multiple words or phrases of the types cited in the preceding descriptions

This yielded an independent variable, Strategy, on which we will report (in Section VI) analyses conducted.

## IV. Data Filtering

Before conducting data analysis, we needed to develop methods to determine which data trials indicated the presence of each bias. We also noted some issues with data that caused us to discard some data trials as unreliable for bias detection. The tank game and the text-based scenarios required separate procedures. These are detailed in this section.

### A. *Method for Detection of Anchoring Bias*

We partition the environments into six equal-size cells (three horizontal divisions and two vertical divisions). The initial position of the player's game piece was always the lower-left cell. The tanks were placed in a randomly-chosen cell other than the lower-left cell. The game proceeded in two phases. In the anchoring phase, the cell containing the tanks did not change, and the player played five iterations of the game. In the evaluation phase, a new cell was again chosen randomly for the tanks from the four remaining cells (not the initial cell for the player and not the previous cell for the tanks). The player did not know the game switched to the evaluation phase, but simply played two more iterations of the game. For detecting anchoring bias, we check whether, during an evaluation run, the player visited the location where the tanks were during the anchoring runs before exploring other regions of the map. Recall that the map of the game board outside the viewport is not visible to the player while playing the game. So, the only reason for a player to go towards the anchoring location would be due to anchoring bias induced by the location retained in their memory during anchoring runs. If the trajectory during an evaluation phase includes visits into cells that contain the previous position of the tanks, then we considered the player to exhibit anchoring bias. Further details of this are available in a previous paper [1].

From the 74 players that played our game for two game sets each, we collected 148 data instances. Each instance was comprised of five anchoring runs followed by two evaluation runs. These data instances were analyzed for detecting anchoring bias. While analyzing, we found that some of the data instances had to be discarded owing to an oversight in the placement of the anchor. If the location of the tanks during the evaluation run was in-between or en-route from the start location to the location of tanks during the anchoring runs, then it was not possible to determine if the player was anchored or not. We discarded 69 of the 148 data points, leaving 79 valid data points.

### B. *Assessing the Responses to Text-based Stimuli*

Since three of our question sets were designed to elicit cognitive biases via scenarios presented (primarily or completely) through prose, we applied a filter based on the reading speed implied by the question word count and the response time. Reading speed has been studied for well over 100 years; however, there still seems to be some concerns raised in the literature about the accuracy of the estimates of reading speed. Brysbaert [26] reviewed 190 studies, dating back to 1898. He concluded that for adults reading silently in English (the language of our study), an average reading speed is 238 words per minute (wpm) for non-fiction and 260 wpm for fiction. Noting the existence of "reliable individual differences," he gives ranges of 175-300 wpm (non-fiction) and 200-320 wpm (fiction). He further noted general agreement that college-age young adults have the highest reading speed. (As noted in Section III, this age group was over half of our participant pool.) In addition, we note that our subjects were reading texts that were markedly below their grade level; our texts were rated with the Flesch-Kincaid Grade Level [27] as being eleventh grade level (approximately age 17 in the U.S.) or lower. Intuitively, this could increase the reading speed, although we lack a good estimate for this increase. Furthermore, it is possible that participants did not read all answers choices (perhaps choosing the first or second that they read), making it hard to estimate reading speed for a full question-and-answer set. Finally, our scenarios and answer choices were generally short; one scenario was 126 words, whereas the remaining 53 (across all bias types) were 35-96 words. Answer choices were 8-66 words. This makes reading speed estimates somewhat sensitive to the short length of the passages. In order to be extremely conservative in disqualifying our participants, we assert a maximum reasonable speed of 1000 wpm. Trials above this reading speed were removed from the analysis; we note that a stricter limit of 640 wpm did not substantially change the results. We also removed trials that were "orphaned" by this filtering, in that very few trials from that participant for a certain condition (e.g., graphs present, or scenario type) remained. Keeping such trials would have made the analysis too sensitive to a small sample.

In summary for these three tests, from 3960 data trials; the wpm filter left 3294 trials for analysis. There were 64 participants who completed the attraction effect scenarios, 61 who completed the compromise effect scenarios, and 58 who completed the zero-risk bias scenarios.

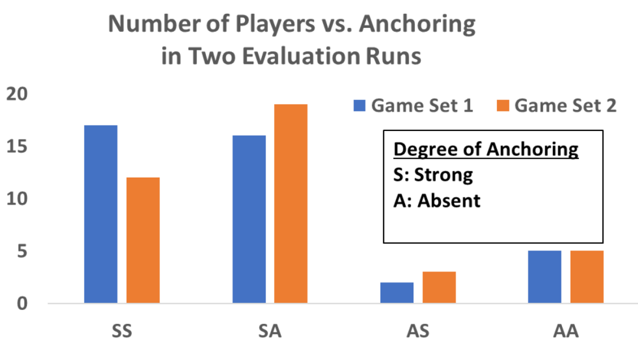


Figure 4. Bar chart showing the number of players (y-axis) that have Strong or no (Absent) anchoring (x-axis) in our two game sets, each consisting of five anchoring runs and two evaluation runs.

## V. Analysis of Anchoring Bias

Because the anchoring bias portion of the study requires a very different method of analysis, we focus this section on analysis of anchoring bias. We had a few research hypotheses and goals, which we present in subsections.

### A. *Exhibiting Anchoring Bias*

*Hypothesis 1*: Participants would exhibit the anchoring bias after five iterations of the tank game.

We detect anchoring bias when the trajectory data from either the first or both evaluation runs meet the criteria above (Section IV.A.). The results show (Fig. 4) evidence of anchoring bias. Out of the 79 data instances, 64 data instances (81%) showed that the player had been anchored (SS and SA in Fig. 4) either in both or only in the first evaluation runs. Across the two game sets, there was very little variation (6%) in the number of subjects displaying anchoring bias. This indicates a strong propensity for anchoring bias among the subjects.

### B. *Duration of Anchoring Bias*

*Hypothesis 2*: When anchoring bias was present, it would last through both evaluation runs.

We determined the number of data instances that showed strong anchoring in the first evaluation run versus those that showed strong anchoring in both evaluation runs (SA versus SS in Fig. 4). We found that in 35 instances players showed that the effect of anchoring waned between the first and second evaluation runs, while the anchoring remained strong between the two evaluation runs for 29 instances. These values indicate that there is small but non-negligible support that the effect of anchoring bias diminishes if the player gets information that contradicts the anchor.

We found that in the first game set, 16 players showed anchoring only in the first evaluation run and 17 showed anchoring in both evaluation runs. In the second game set, these numbers became 19 and 12, respectively. The decrease in strong anchoring in both evaluation runs between the first and second game sets (from 17 to 12), and simultaneous increase in subjects that showed anchoring only in the first evaluation run (from 16 to 19) points further in the direction that, as the player sees more information contradicting the anchor, the effect of anchoring diminishes. Players may have been more fatigued at the start the second set of evaluation runs, after playing 12 runs (five anchoring runs in each of two game sets plus two evaluation runs in first game set) of the game. Conventionally, fatigue would lead to the human brain making shortcuts via heuristics and strengthening the anchoring bias. However, we saw diminishing anchoring bias across game sets. This seems to indicate that the disappointment of not finding the tanks at the anchoring location weakens the anchoring bias and motivates the player to explore in a more objective, less biased manner.

### C. *Building a Model for Prediction of Anchoring Bias*

*Hypothesis 3:* A bias prediction model would enable us to predict exhibition of the bias during evaluation runs from a propensity toward bias exhibited in anchoring runs.

To investigate this hypothesis, we used a bias prediction model based on the work of Jesteadt et al. [28]. We summarize our previous discussion [25] of this model. The model is a linear combination of the stimulus from the current task perception, the stimulus from the task in the previous time-step, and the outcome of the decision in the previous time-step. We mask the current task perception, eliminating one factor. We consider the trajectory length up to viewing the first tank in the viewport as the stimulus from that anchor. This yields a linear model in which each anchor's influence during evaluation is the $J_{eval} = \alpha + \Sigma\ \beta\ J_{anc,i,}$ where the summation is over i=1..5 for the five anchoring runs. We used linear regression with least squares [29] to solve this equation. If the slope of the regression line was less than zero, then we say that the participant had propensity toward anchoring bias.

We compare this bias prediction model with the detection. For the first evaluation run (Fig. 5, parts (a) and (c)), the model was generally accurate (true positive plus false negative of 80% and 77%, respectively). Unsurprisingly, the prediction accuracy of the model diminishes considerably to 52% and 37%, respectively, in the two game sets (Fig. 5, parts (b) and (d)). It appears that the exposure to a different location of tanks than the anchoring runs in the first evaluation run reduced the player's reliance on the anchor to search for the tanks during the second evaluation run. We note that none of these results reach the threshold of statistical significance through Fisher's Exact Test. In our best result (Fig. 5(a)), there was statistically no association between the model prediction and the exhibition of bias (p=0.204). We attribute the disconnect between the apparently high percentage of accuracy and the failure to achieve statistical significance to the low sample size; we again lament the need to remove data, as discussed in Section IV.A. We note that with 100 participants (whose data were not invalidated as described in Section IV.A.) and the percentages we observed on the first evaluation run (Fig. 5(a)), Fisher's Exact Test would yield statistical significance.

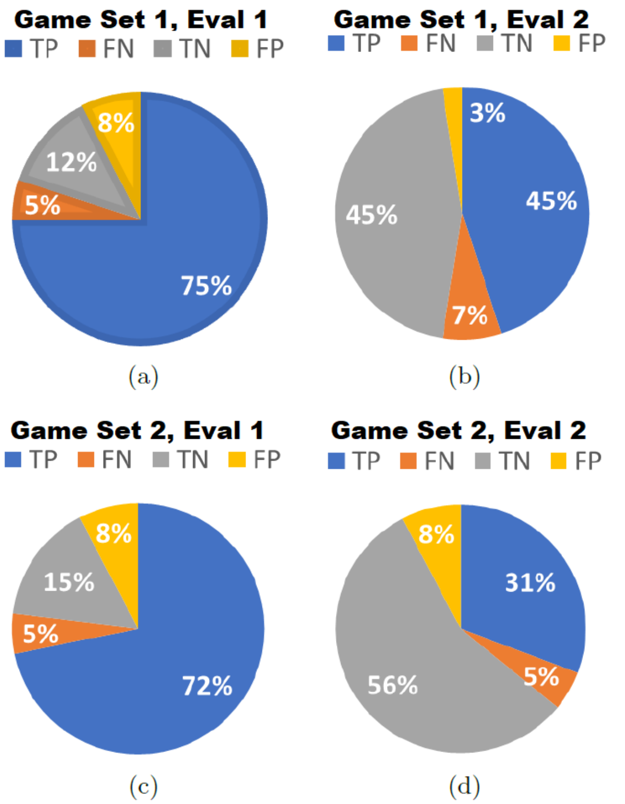


Figure 5. Effect of anchoring bias propensity during anchoring runs on decision in evaluation runs, for game sets 1 (40 trajectories) and 2 (39 trajectories). T/F denote anchoring during anchoring runs as true or false; P/N denote detection of anchoring during evaluation as positive or negative. So, TP and FN are accurate predictions, whereas TP and FN are inaccurate.

Players played the two sets of the game back-to-back without any break. An immediately relevant question is whether the model predicts the anchoring in the second iteration of the game, after having seen the anchor no longer be reliable in the evaluation runs of the first iteration of the game. The answer (Fig. 5(c)) appears to be promising, although this result also does not have statistical support from Fisher's Exact Test; it would appear to require approximately 150 (valid) participants at the percentages indicated in Fig. 5(c) to achieve this threshold. Still, such a result would correspond to findings in other sequential decision-making [5], where the anchoring effect diminished as the decision maker was exposed to more information from successive decision problems that were contrary to the features of the problem in the positive decision instance.

Overall, our findings of the anchoring bias prediction model indicate that a more robust prediction model, based on a larger data set, would be worth investigating for longer-term prediction of anchoring bias effects.

### *D. Potential for Real-time Detection of Anchoring Bias*

We had one other long-term goal for which we could not form a hypothesis. We hoped to identify a method which would indicate in real-time whether a participant appeared to be getting anchored, rather than relying on a post-hoc assessment of whether the participant was anchored. It

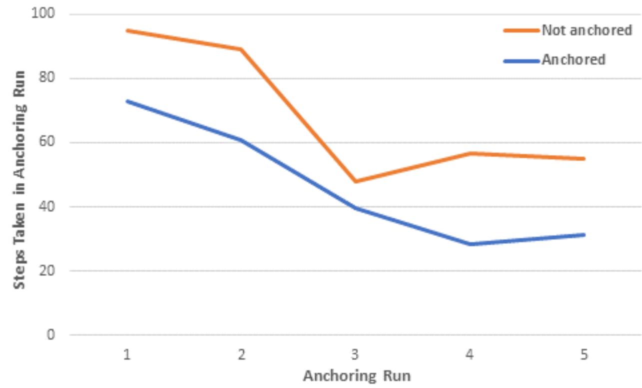


Figure 6. Time steps (moves) used during the five anchoring runs for participants who were later judged to be anchored (blue) and those who were judged to be not anchored (orange). The separation of these two graphs and the significant differences reported lead us to hypothesize that this could be a way to detect anchoring bias in real-time.

appears that the number of time steps (moves) used in our game is a potential real-time indicator of whether a participant is getting anchored. We observed a significant correlation between the count of anchoring runs and the number of time steps (moves) taken in the game until the tanks were within the viewport. This was true for both those judged post-hoc to have been anchored – Pearson R=-0.98, $t(3)=-7.99$, $p < 0.005$ – and those judged to be not anchored – Pearson R=-0.88, $t(3)=-3.29$, $p<0.047$. But these two observations were different. Analysis of variance (ANOVA) showed that, for the first iteration of the game, those who were not anchored were significantly slower than those who were – $F(1,38)=5.793$, $p<0.022$, generalized effect size $\eta=0.132$. This difference disappeared for the second iteration of the game – $F(1,37)=0.393$, $p>0.534$. Looking at the graph of the moves taken (Fig. 6), we can see the separation, which leads us to hypothesize that this may be a way to detect anchoring bias in real time. Further data would have to be gathered to validate this hypothesis.

## VI. Analysis of Biases through Text Scenarios

The attraction effect, compromise effect, and zero-risk bias were hoped to be elicited through text-based scenarios. The analysis of these three biases follows a similar pattern, and we present these analyses in the following subsections.

As a preliminary result, we noted a strong correlation – Pearson R=0.85, $t(16)=6.336$, $p<0.001$ – between the trial number (1-18) and reading speed (Fig. 7). This correlation measured reading speed under the assumption that participants read the scenario and all three answer choices completely. This measure was averaged across all three decision-making tasks, which "folds over" the trial number three times, if participants completed all three question sets. Some of this could be attributed to using the graphs as a shortcut, which some participants indicated that they were doing. This could also indicate that participants were getting fatigued, since they "read" faster and faster as they progressed through the trials. Note that the last three trials (and at least four of the last five) are faster than the speed that the literature on reading comprehension indicates is likely for a reader to be able to read for comprehension (see

TABLE I: DISTRIBUTION OF RESPONSES TO THREE SCENARIO TYPES

| Attraction Effect | | Compromise Effect | | Zero-risk Bias | |
|---|---|---|---|---|---|
| Attractor | 401 | Compromise | 384 | Zero-risk | 365 |
| Decoy | 373 | Extreme | 714 | Balanced risk | 679 |
| Target | 378 | | | | |

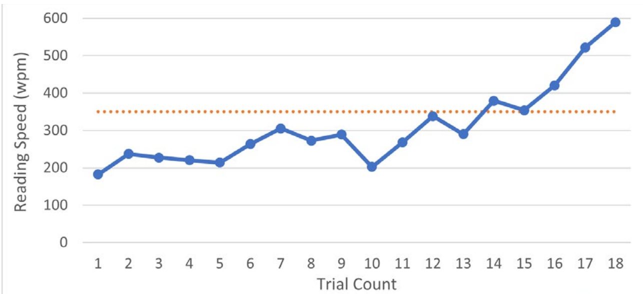


Figure 7. Measured reading speed versus trial count. We saw a strong correlation (Pearson R=0.846, $t(16)=6.336$, $p<0.001$) between trial number and reading speed. Trials 10 through 18 showed graphs, which may have influenced "reading speed" by enabling a shortcut to reading prose for the information needed to make a decision.

Section IV.B.). Perhaps participants were not reading all the answer choices, or they were just scanning. Some of this may be attributed to most of the scenarios being at a reading grade level well below the participants' respective abilities (based on the self-reported education levels, reported in Section III).

Overall, we hypothesized a skewed distribution of the responses to the three types of scenarios (for the three biases discussed below). We believed this would indicate the presence of the bias. As shown in Table I, we got a nearly balanced set of responses to each of the three question types. Thus, we cannot conclude that the bias was elicited by our scenarios. Therefore, our analysis focuses on whether there were certain features of the scenarios or of the participants that may be associated with eliciting the bias in a subset of the data. We note that not all participants completed all sections of the study, so the degrees of freedom in the ANOVA measures below are not the same as in the test above for the anchoring bias, nor do they match each other.

## A. Detection of Attraction Effect

Following [16], we measured the correlation between the REI scores (both the rational and experiential, as well as the ability and engagement sub-scales within each score) and the selection of the unpreferred options. We did not find a significant correlation between any REI score (the two main scales or any of the four sub-scales) and the rate of selecting unpreferred responses. We also measured the CRT score for each participant and measured the correlation of these scores against the rate of selection of the unpreferred options. Again, we found no significant results for the attraction effect. We had hoped that we would not only elicit the biases described above, but that we could help mitigate these by the presence of the graphs. However, there was no main effect of the presence of the graphs in the second half of the set of questions for the attraction effect.

We wrote scenarios of multiple types that reflected many roles members of the military might encounter in their duties. This encompasses mundane issues like purchasing non-military equipment, purchasing military equipment, and even more weighty matters such as aspects of military strategy. We noted a main effect – $F(3,189)=9.465$, $p<0.001$, $\eta=0.030$ – of this cost measure on the response time. However, we note that the generalized effect size $\eta$ was very small. The questions were ordered according to a Latin square using the question ID; this led to some variation in the relative placement of the type of cost measure that created a potential confound of this effect. In addition, some of this effect may reflect the general behavior of participants to get faster as the question count rose. Finally, we categorized scenarios into multiple types, inflating the degrees of freedom and potentially the effect. Therefore, we note this effect, but do not yet consider it to be a reliable result.

We had hoped to elicit the biases, and also mitigate these by the presence of the graphs illustrating the data. Unfortunately, there was no main effect of the presence of the graphs in the second half of the study on the rate of selecting the attractor or decoy – $F(1,63)=0.689$, $p=0.409$. We also did not observe a main effect of the presence of the graphs on response time – $F(1,63)=0.977$, $p=0.326$. Foreshadowing the contrast with the compromise effect and zero-risk bias (which both show this main effect), we note that the attraction effect scenarios tended to have a lower word count, which could have confounded the "shortcut" of using the graph and not reading the answer choices. There was a main effect of graph sentiment on the rate of selecting the various responses to the attraction effect – $F(2,61)=3.423$, $p<0.039$, $\eta=0.101$. Participants whose sentiment about the graphs appeared to us to be negative selected unpreferred options on the attraction effect (attractor or decoy) 61.1% of the time. Participants whose sentiment appeared neutral selected unpreferred options 66.7% of the time. Participants whose sentiment appeared positive selected unpreferred options 73.5% of the time. Follow-up t-tests indicate that all differences between these values are statistically significant. (For the smallest difference, negative to neutral, $t(47)=3.236$, $p<0.003$.) It appears that those participants with positive sentiment toward the graphs were led to exhibit the bias with *greater* frequency than those who did not use (or actively avoided) the graphs. In retrospect, perhaps eye tracking would have been advisable, so that we could know exactly what portions of the graph those who said they used it or liked were reading in order to make their decision. That might have given us greater insight to this result.

The Strategy we inferred based on the free text responses had a significant main effect on their rate of selecting the unpreferred option on the attraction effect portion of the study – $F(5,84)=3.228$, $p<0.011$, $\eta=0.161$ (Fig. 8). Using post-hoc t-tests, we determined that there were essentially two groups of three strategies. The strategies (see Section III.E.) *None*, *Intuition*, and *Mixed* were not significantly different from each other, and the remaining strategies of *Balance*, *Efficient*, and *Feature* were not different from each other (though some differences with the second set showed a

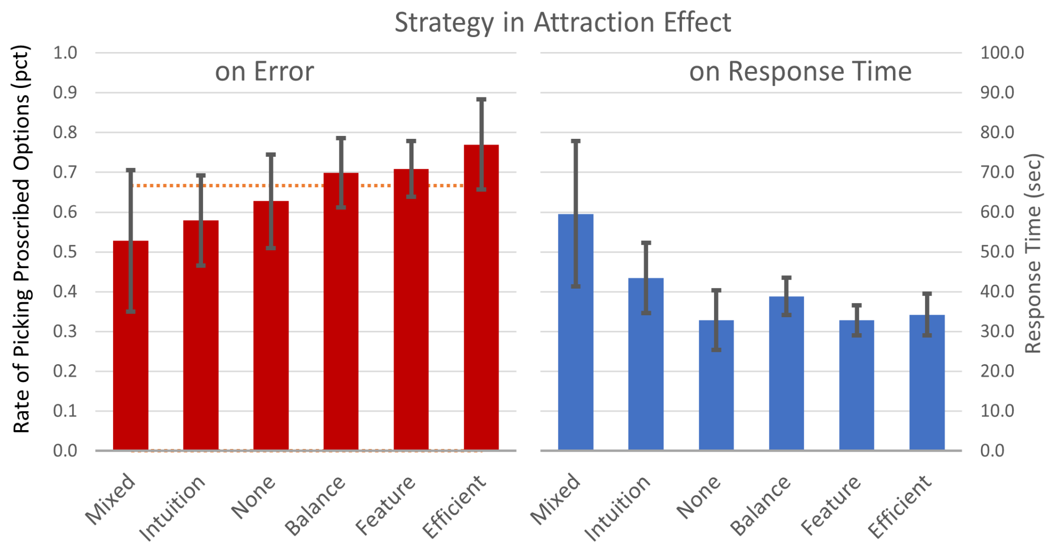


Figure 8. This pair of graphs shows the effects of strategy in the attraction effect portion of the study. The strategy was assessed by the research team based on the free text responses of participants. The assessed values showed a main effect of strategy on the rate of selecting the unpreferred option (left, red bars). The orange line represents random selection; as noted in the text, the three strategies with the least error were statistically different from the three with the most error. We also noted a trend towards difference in response time based on strategy for the attraction effect scenarios (right, blue bars), largely due to the slow responses using the *Mixed* strategy.

trend in the t-test). But all strategies in the first set led to significantly different performance than strategies in the second set. The first column of Table II shows that the usage of these strategies was not equally distributed amongst our participants on the attraction effect scenarios.

There was also a trend for the inferred strategy to lead to differences in response time for the attraction effect scenarios – $F(5,84)=1.990$, $p<0.089$, $\eta=0.106$. However, the grouping is quite different. The *Mixed* strategy was notably slower than all the others, with the *Intuition* strategy being slightly slower than the remaining four. While the graphs do not mirror each other, we do note that the *Mixed* strategy had the lowest rate of selecting the unpreferred option, took the longest time, and was the least-often used. We did not find an interaction between graph sentiment and inferred strategy.

There was statistically no association between anchoring bias and the attraction effect. There was no association between the attraction effect and the compromise effect, nor between the attraction effect and the zero-risk bias ($p>0.409$ for all associations).

### B. *Detection of Compromise Effect*

As with the attraction effect, we measured the correlation between REI scores and CRT scores. Again, we did not find a significant correlation between REI score (or any subscale) and choosing the compromise option. Similarly, we did not find a correlation with CRT scores and selection of the compromise option. We had hoped that we would not only elicit the biases described above, but that we could help mitigate these by the presence of the graphs. However, there was no main effect of the presence of the graphs in the second half of the set of questions for the compromise effect.

Regarding the type of scenario, we once again observed a main effect of the type of cost measure on response time – $F(1,60)=27.970$, $p<0.001$, $\eta=0.032$. As with the attraction effect, there was uncontrolled variation in the ordering (we did not counterbalance the order of questions by type), and participants generally got faster, which may confound this effect. With the compromise effect, we had only two scenario types, but even without this potential inflation, we consider it a weak (small η) and, for now, unreliable result.

TABLE II. FREQUENCY DISTRIBUTION OF STRATEGIES IN SCENARIOS

| Strategy | Attraction Effect | Compromise Effect | Zero-risk Bias |
|---|---|---|---|
| None | 17 | 13 | 12 |
| Intuition | 19 | 14 | 12 |
| Feature | 42 | 33 | 51 |
| Efficient | 14 | 8 | 6 |
| Balance | 28 | 41 | 29 |
| Mixed | 8 | 13 | 6 |

Again, as we did with the attraction effect, we had hoped that the presence of the graphs would reduce the rate of selecting the compromise option. Unfortunately, we did not observe a main effect – $F(1,60)=2.000$, $p=0.162$. However, we did observe a main effect of the presence of the graphs on response time – $F(1,60)=23.943$, $p<0.001$, $\eta=0.029$. We see a statistically significant but small effect, and like the previous main effect of scenario type on response time, because the graphs were always the second half of each portion of the study, this effect could be due to faster ingest of information through the graph, or may be confounded with the general pattern of participants getting faster with the increasing number of data trials to complete.

There was no main effect of graph sentiment on the distribution of choices among the options – $F(2,58)=1.081$, $p>0.453$. Neither was there a main effect on the response time – $F(2,58)=1.500$, $p>0.231$. There was no main effect of strategy on the rate of selecting the compromise option – $F(6,94)=1.501$, $p>0.186$ – or on the response time – $F(6,94)=0.778$, $p>0.589$. However, examining the distribution of the inferred use of each strategy (middle column of Table II) and the definition of the *Balance* strategy (Section III.E.), we see that the most commonly named strategy was choosing a "balance" between the features or the option that was "in the middle." This explicitly expresses a bias for the compromise option, and is the strongest evidence we have that at least some participants made their choice on the basis of this bias. However, this strategy was used on approximately one-third of the data trials (33.6%), so while it indicates that some users chose on the basis of this bias, it does not constitute a majority strategy, and it does not appear to have led to significantly better performance. We did not find an interaction between graph sentiment and inferred strategy.

There was no statistical association between anchoring bias and the compromise effect. Nor was there an association between the compromise effect and zero-risk bias ($p>0.748$ for both associations). We observed a trend toward statistical association between the inferred strategy of *Balance* (defined in Section III.E. as participants saying they tried to balance the competing interests or features, or that they chose the middle option), $p<0.058$. So, it appears quite possible that participants were accurate in assessing their own strategy, at least with the Balance strategy.

### C. Detection of Zero-risk Bias

As with the two previous biases, we measured the correlation between the REI scores and the rate of zero-risk bias, as well as the CRT scores and the rate of the zero-risk bias. In contrast to the previous results, we did see a significant correlation between the CRT score and the rate of selecting the zero-risk option, $t(56)=2.5425$, $p=0.014$. This would appear to extend the prior results, since zero-risk bias was not a part of the previous work [16]. We get a similar result evaluating this association with Fisher's Exact Test, yielding an association with $p<0.022$. A significant portion of the incorrect responses we saw to the CRT were the intuitive-but-incorrect responses [11]. Therefore, it should not be a surprise that this effect also manifests itself in a significant (negative) correlation between the number of the intuitive (but incorrect) responses our participants gave and their rate of selecting the zero-risk option, $t(56)=2.954$, $p<0.005$. The zero-risk option can be seen as an intuitively best choice, even though deeper analysis shows it is inferior.

Again, in parallel to the previous two bias types, we noted a main effect of scenario type on response time – $F(4,228)=7.665$, $p<0.001$, $\eta=0.041$. All the caveats described for the attraction and compromise effects apply to the zero-risk bias portion of the study (uncontrolled variation in ordering of scenario type, general behavior of participants to get faster, large degrees of freedom due to multiple type categories). Again, the effect is not yet considered reliable.

Turning to the presence of the graphs, we did not observe a main effect on the rate of selecting the zero-risk option – $F(1,57)=0.040$, $p=0.841$. But we observed a main effect on the response time – $F(1,57)=9.325$, $p<0.021$, $\eta=0.021$. As with the compromise effect, this main effect is potentially confounded with the general tendency for participants to get faster as the number of data trials increased (since the graphs were always the second half of this section of the study).

We found that graph sentiment had a main effect on the response time in the zero-risk portion of the study – $F(2,55)=5.557$, $p<0.007$, $\eta=0.168$. Participants with positive sentiment were fastest (58.3 sec), followed by participants who expressed no sentiment about the graphs (66.1 sec). Participants who said they found the graph unhelpful or said they did not use it were slowest, at 94.2 sec. All these pairwise differences are significant with $p<0.001$. As with the result for the attraction effect, we lament the missed opportunity to have recorded eye movements and gain greater insight into this result. There was no main effect on the selection of responses to the scenarios intended to elicit the zero-risk bias – $F(2,55)=1.476$, $p>0.710$.

There was a trend for these strategies to lead to different rates of selecting the unpreferred option on the zero-risk bias portion of the study – $F(5,76)=2.188$, $p<0.065$, $\eta=0.126$. There was a trend for the various strategies to lead to different response times. Since we chose (as described above) to make the zero-risk option objectively worse (i.e., a greater total loss), the *Feature* strategy should have helped identify this aspect of the scenario's data and thus have been successful at avoiding the unpreferred option. But participants had to look for the minimal loss summed over the two portions of the outcome. Using post-hoc t-tests, we found that this strategy had a statistically lower rate than only the *Balance* strategy. There was a trend for the *Feature* strategy to be better than the *Efficient* strategy. There was main effect of strategy on the response time – $F(5,76)=1.77$, $p>0.128$. We did not find an interaction between graph sentiment and inferred strategy.

There was statistically no association between anchoring bias and zero-risk bias ($p>0.745$). Nor was there association between inferred strategy and the zero-risk bias ($p>0.406$).

## VII. Discussion

We investigated using game-playing and decision-making exercises to induce cognitive biases, with limited success. We were able to induce anchoring bias. However, for a small fraction of the players (1 out of 74 instances in set 1 and 3 out of 74 instances in set 2), we found that they initially showed influence of the anchor during the first few anchoring runs, but in subsequent anchoring runs and in the evaluation run, the anchoring effect went away. They started exploring the map instead of heading to the anchor location. Fig. 9 shows an example where the first two anchoring runs (top) show anchoring but the other anchoring runs (bottom) do not. This de-anchoring was more pronounced in set 2. Perhaps the evaluation runs in game set 1 reduced the reliance of the player on the anchor during game set 2 even after they found it, prompting general exploration again.

The movement of game-piece in our computer-based game for anchoring bias was controlled by keyboard arrow keys; thus, it was limited to the four cardinal directions. This resulted in players using long horizontal or vertical tracks to

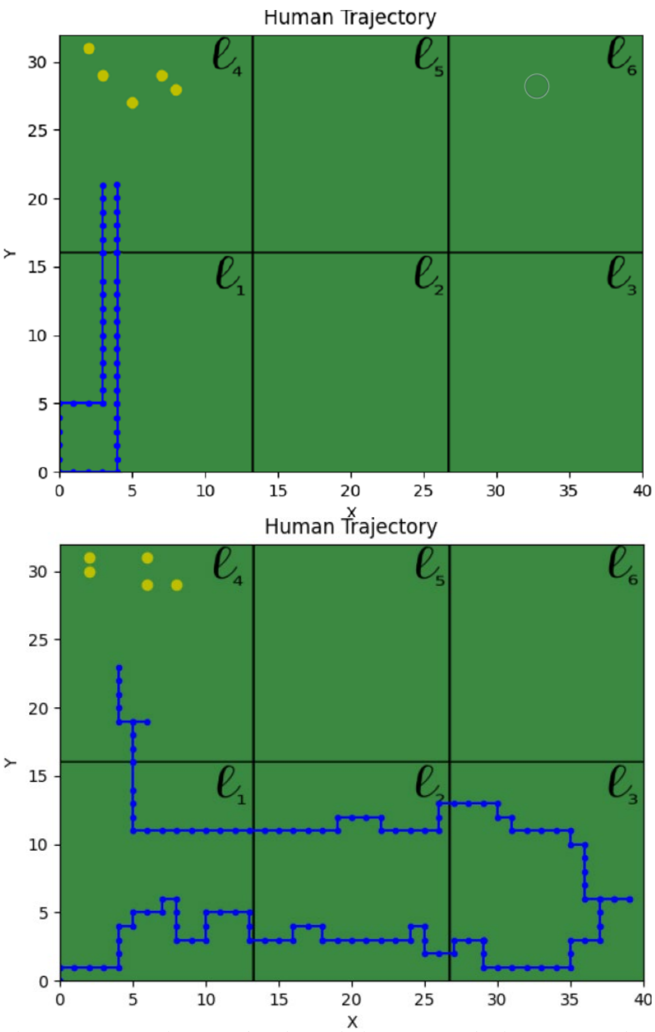


Figure 9. Top: Trajectory of a player during set 2 anchoring runs 1 and 2. Bottom: Trajectory during anchoring runs 3-5 (bottom). These images also illustrate the six cells of the game board to detect anchoring bias.

explore the environment. The number of key presses made by players in the game was not recorded and there is a possibility that some players purposefully reduced the number of keystrokes by holding keys to continue in the same direction for long periods. This could also have stemmed from psychological factors like motivation, interest, and engagement with the game and overall experiment.

Anchoring bias, as we have used the term, intersects with other biases. Sequential bias deals with repetitive decision outcomes in sequential (not necessarily time-extended) tasks. Experiential bias considers the reliance of humans on experience from past decision outcomes on the current decision. It would be interesting to analyze our results with appropriate theoretical models for these biases, to understand overlap, similarity, and divergence between these biases.

What causes humans to depend on anchors for making decisions? The conventionally accepted theory is the human brain is inclined to make shortcuts via heuristics [2] due to boredom, motivation, repetitiveness and other factors. In contrast, the selective accessibility model [30] proposed an alternative theory that the brain made information related to the anchor more readily accessible to its decision process. The difference is subtle but consequential, as the former attributes the cause of anchoring bias to the internal working of the brain's decision-making process while the latter attributes it to the information presented to the brain's decision-making process. A deeper understanding, fortified with appropriate mathematical models for these two theories, would help with a clearer understanding of anchoring bias.

## VIII. CONCLUSION AND FUTURE WORK

We elicited anchoring bias with a time-extended task. We were statistically unable to elicit the attraction effect, compromise effect, or zero-risk bias. We have some evidence that the zero-risk bias and compromise and attraction effects exist due to selection of options that should not have been selected at all and due to statements from the participants in which they indicate the compromise effect. However, we had limited success in identifying factors that contribute to the bias. The stated strategy and the use of graphs seems to have influenced the attraction effect, and the stated strategy also influenced the zero-risk bias. Further research would be needed to determine how one might systematically elicit (and therefore, how to systematically mitigate) these biases. As noted above, we recommend that future research on the potential for graphs to mitigate bias should incorporate eye tracking to determine the extent of participants' use of graphs and which options they considered the longest. We recommend counterbalancing the usage of graphs to investigate the reliability of our results on the response time in the presence of graphs for both the compromise effect and zero-risk bias.

One could explore links between learning style, problem-solving strategies, and these biases. Perhaps the biggest difference between the text-based scenarios (with or without graphs) is the level of engagement participants had with the exercises. Our initial design concept included interleaving the tank game with responding to decision-making scenarios at stopping points in the game. This design felt contrived; thus, we worried that our data would suffer. Clearly, there is room for improvement on our design. We repeat the lament of our tank game design and the placement of anchors that caused us to discard many data trials because the placement of the tanks rendered it impossible to determine if the path was due to anchoring or not.

We note one potential avenue to identify bias in real-time. The anchoring effect appears potentially predicted by the time spent in the anchoring runs. The stated strategy of the participant matched the exhibition of the compromise effect. In what appears to be a new result in the literature, we found a statistical effect of the responses to the CRT and the exhibition of the zero-risk bias. We believe these observations hold promise for future research. Study designs, with larger participant pools, that explicitly investigate these relationships could potentially validate our results and would represent logical next steps for the detection and prediction of each of these biases.

ACKNOWLEDGMENT

The authors thank Zhuming Ai, Bryan Brandt, Joseph Coyne, Kayla Russell, and Kaylin Strong. The authors also thank the editors and reviewers for their many contributions to the improvement of this manuscript. The Office of Naval Research funded our work via the NRL Base Program.

REFERENCES

[1] P. Dasgupta, J. Kliem, M.A. Livingston, J. Decker, "Inducing and Detecting Anchoring Bias via Game-Play in Time-extended Decision-Making Tasks," IARIA Annual Congress on Frontiers in Science, Technology, Services, and Applications, July 2025

[2] E. Dimara, S. Franconeri, C. Plaisant, A. Bezerianos, P. Dragicevic, "A Task-Based Taxonomy of Cognitive Biases for Information Visualization," IEEE Transactions on Visualization and Computer Graphics, vol. 26, no. 2, pp. 1413-1432, 2020

[3] A. Tversky and D. Kahneman, "Judgment under Uncertainty: Heuristics and Biases," Science, vol. 185, no. 4157, pp. 1124-1131, 1974

[4] P.T. Sukumar, R. Metoyer, S. He, "Making a Pecan Pie: Understanding and Supporting The Holistic Review Process in Admissions," Proceedings of the ACM Conference on Human-Computer Interaction, Volume 2 CSCW, Article no. 169, 2018

[5] J.M. Echterhoff, M. Yarmand, J. McAuley, "AI-Moderated Decision-Making: Capturing and Balancing Anchoring Bias in Sequential Decision Tasks," Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, Article no. 161, 2022

[6] I. Simonson, "Choice Based on Reasons: The Case of Attraction and Compromise Effects," J. of Consumer Research, vol. 16, no. 2, pp. 158-174, 1989

[7] J. Huber, J.W. Payne, C. Puto, "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis," J. of Consumer Research, vol. 9, no. 1, pp. 90-98, 1982

[8] R. Kivetz, O. Netzer, V. Srinivasan, "Extending Compromise Effect Models to Complex Buying Situations and Other Context Effects," J. of Marketing Research, vol. XLI, pp. 262-268, Aug 1989

[9] J. Baron, R. Gowda, H. Kunreuther, "Attitudes Toward Managing Hazardous Waste: What Should Be Cleaned Up and Who Should Pay for It?," Risk Analysis, vol. 13, no. 2, pp. 183-192, 1993

[10] W.K. Viscusi, W.A. Magat, J, Huber, "An Investigation of the Rationality of Consumer Valuations of Multiple Health Risks," RAND Journal of Economics, vol. 18, no. 4, pp. 465-479, Winter 1987

[11] S. Frederick, "Cognitive Reflection and Decision Making," J. of Economic Perspective, vol. 19, no. 4, pp. 25-42, Fall 2005.

[12] R. Pacini and S. Epstein, "The Relation of Rational and Experiential Information Processing Styles to Personality, Basic Beliefs, and the Ratio-Bias Phenomenon," J. of Personality and Social Psychology, vol. 76, no. 6, pp. 972-987, 1999

[13] J.T. Coyne et al., "Evaluation of New Measures of Spatial Ability and Attention Control for Selection of Naval Flight Students," Human Factors in Design, Engineering, and Computing (Proceedings of AHFE Intl.), vol. 159, pp. 2006-2016, 2024

[14] D. Masson, "Balanced Latin Square Generator," Online at https://damienmasson.com/tools/latin_square, Last accessed 11 Oct. 2025

[15] J.V. Bradley, "Complete Counterbalancing of Immediate Sequential Effects in a Latin Square Design," J. of the American Statistical Association, vol. 53, no. 282, pp. 525-528, 1958

[16] P. Sleboda and J. Sokolowska, "Measurements of Rationality: Individual Differences in Information Processing, the Transitivity of Preferences and Decision Strategies," Frontiers in Psychology, vol. 8, pp. 1844, 2017

[17] M.W.B. Zhang, J.B. Ying, G. Song, R.C.M. Ho, "A Review of Gamification Approaches in Commercial Cognitive Bias Modification Gaming Applications," Technology and Health Care, vol. 26, no. 6, pp. 933-944, 2018

[18] S.A. Malkoc, W. Hedgcock, S. Hoeffler, "Between a rock and a hard place: The failure of the attraction effect among unattractive alternatives," J. of Consumer Psychology, vol. 23, no. 3, pp. 317-329, 2013

[19] R. Dhar, S.M. Nowlis, S.J. Sherman, "Trying Hard or Hardly Trying: An Analysis of Context Effects in Choice," J. of Consumer Psychology, vol. 9, no. 4, pp. 189-200, 2000

[20] A. Mintz, N. Geva, S.B. Redd, A. Carnes, "The Effect of Dynamiac and Static Choice Sets on Political Decision Making: An Analysis Using the Decision Board Platform," American Political Science Review, vol. 91, no. 3, pp. 553-566, Sep. 1997

[21] L.K. Marusich, N. Buchler, J.Z. Bakdash, "Human Limits to Cognitive Information Fusion in a Military Decision-Making Task," 19th Intl. Command and Control Research and Technology Symposium, June 2014

[22] E. Schneider, B. Streicher, E. Lermer, R. Sachs, D. Frey, "Measuring the Zero-Risk Bias: Methodological Artefact or Decision-Making Strategy?" Zeitschrift für Psychologie (J. of Psychology, Germany), vol. 225, no. 1, pp. 31-44, 2017

[23] D.J. Mezzio, V.B. Nguyen, A. Kiselica, K. O'Day, "Evaluating the Presence of Cognitive Biases in Health Care Decision Making: A Survey of U.S. Formulary Decision Makers," J, of Managed Car & Specialty Pharmacy, vol. 24, no. 11, pp. 1173-1183, Nov. 2018

[24] L. Dodd and J. Moffat and J. Smith, "Discontinuity in decision-making when objectives conflict: a military command decision case study," J. of the Operational Research Society, vol. 57, no. 6, pp. 643-654, 2006

[25] C. Nobre, K. Zhu, E. Mörth, H. Pfister, J. Beyer, "Reading Between the Pixels: Investigating the Barriers to Visualization Literacy," Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, May 2024

[26] M. Brysbaert, ""How Many Words do we Read per Minute? A Review and Meta-analysis of Reading Rate," J. of Memory and Language, vol. 109, Dec. 2019

[27] J.P. Kincaid, R.P. Fishburne Jr., R.L. Rogers, B.S. Chisson, "Derivation of New Readability Formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy Enlisted Personnel," Research Branch Report 8-75, Naval Air Station Memphis, Feb. 1975

[28] W. Jesteadt, R.D. Luce, D.M. Green, "Sequential Effects in Judgments of Loudness," J. of Experimental Psychology: Human Perception and Performance, vol. 3, no. 1, pp. 92-104, 1977

[29] Y. Nievergelt, "Total Least Squares: State-of-the-Art Regression in Numerical Analysis," SIAM Review, vol. 36, no. 2, pp. 258-264, 1994

[30] T. Mussweiler and F. Strack, "Comparing is Believing: A Selective Accessibility Model of Judgmental Anchoring," European Review of Social Psychology, vol. 10, no. 1, pp. 135-167, 1999