

# The Role of Heuristics and Biases During Complex Choices with an AI Teammate

Nikolos Gurney<sup>1</sup>, John H. Miller<sup>2,3</sup>, David V. Pynadath<sup>1</sup>

<sup>1</sup> Institute for Creative Technologies, University of Southern California

<sup>2</sup> Carnegie Mellon University

<sup>3</sup> Santa Fe Institute

gurney@ict.usc.edu, pynadath@ict.usc.edu, JM7T@andrew.cmu.edu

## Abstract

Behavioral scientists have classically documented aversion to algorithmic decision aids, from simple linear models to AI. Sentiment, however, is changing and possibly accelerating AI helper usage. AI assistance is, arguably, most valuable when humans must make complex choices. We argue that classic experimental methods used to study heuristics and biases are insufficient for studying complex choices made with AI helpers. We adapted an experimental paradigm designed for studying complex choices in such contexts. We show that framing and anchoring effects impact how people work with an AI helper and are predictive of choice outcomes. The evidence suggests that some participants, particularly those in a loss frame, put too much faith in the AI helper and experienced worse choice outcomes by doing so. The paradigm also generates computational modeling-friendly data allowing future studies of human-AI decision making.

## Introduction

The superiority of passive, algorithmic decision making is dogma in the behavioral sciences (Dawes and Corrigan 1974; Dawes, Faust, and Meehl 1989). This reality is unsurprising given that even poorly-tuned models outperform humans (Dawes 1979). Nevertheless, researchers have documented a strong aversion to relying on algorithms (Dietvorst, Simmons, and Massey 2015; Burton, Stein, and Jensen 2020). There is increasing evidence, however, that people are more open to input from algorithmic decision aids (Logg, Minson, and Moore 2019). This shift towards a more trusting stance aligns with intuition: combining the proliferation of algorithmic decision aids and the insight that repeated exposure to stimuli can alter the ways in which people respond to them (Zajonc 1968; Bornstein and D'agostino 1992) creates a scenario in which it would be surprising if people *did not* exhibit more trust.

Increasingly, decision aids rely on adaptive rather than passive algorithms. These artificial intelligence (AI) helpers are optimized to function in settings where passive algorithms, such as linear models, underperform relative to human decision making. Although these technologies function and behave in fundamentally different ways from more passive predecessors, researchers largely still study them by

relying on traditional empirical approaches: either simple choice experiments (e.g., Logg, Minson, and Moore 2019, in which the authors use their evidence to motivate the need for further research, particularly in settings where AI are in use) or *in situ* (in which there are entire sub-fields, such as human-robot interaction). We argue that a middle ground is needed where the experimental setting matches the technological capabilities of AI but avoids the contextual challenges of *in situ* experimentation.

We present results from such an experiment in which participants work with an AI helper to search through complex choice options. We find classic behavioral biases, specifically anchoring and framing effects, are still present in the choice behavior. These biases also predict choice outcomes and highlight opportunities to develop AI helpers that can improve outcomes by accounting for these biases. Critically, we engineered the experimental environment such that it generates AI-friendly data. This feature creates future opportunities for training AI helpers using data from an experimental paradigm in which known artifacts of human cognition are present.

## Related Works

### Complex Choices

A blossoming corner of behavioral science is working to understand the nuance of human judgment and decision making during *complex choice*. A complex choice is one in which the relationships between the variables under consideration are nonlinear. Such a relationship means that even when only a few variables characterize a choice, their interactions can render optimization intractable (Gurney, Miller, and Pynadath 2023). Obvious examples of complex choices are home buying, career moves, and mate selection. In any of these examples, changing a single choice variable can dramatically impact how other variables affect the choice. In the case of buying a home, for example, the opportunity to work remotely changes how commuting distance, home size, and a host of other variables interact. It is not just monumental choices that suffer from complexity. Filling a shopping cart given a budget constraint, something that most adults regularly do, can be complex. A shopper must, for example, consider how the myriad of items they buy combine into meals, when they will perish, whether they will satisfy their taste

preferences in the future, etc.

Behavioral scientists have not ignored the unique nature of complex choices. Herbert Simon, for example, noted that many choices are intractable due to the interaction of variables (1956). Despite this recognition, complex choice experiments have only recently emerged. Behavioral science historically relied on experimental designs that remove complexity from choices because it allows researchers to isolate behavioral nuances. These experiments usually take one of two forms: descriptive, in which all the information needed to make an informed choice is available, or experiential, in which participants must learn about the choice through repeated experience (Jessup et al. 2022).

In descriptive experiments, researchers typically manipulate a single choice feature in a between-subjects design. A classic example from studies of the anchoring heuristic is to ask people about the population density of a city after asking them if the population is more or less than a reference value (Jacowitz and Kahneman 1995). When the reference is low (high), people report that the city is more sparsely (densely) populated. Such experiments are crucial to dissecting human behavior. However, their rigidity limits their usefulness. For example, extensive research documents the dependency of human judgment and decision making on experience, including a person's own thoughts about such experiences (Schwarz 2004). Unfortunately, the very nature of complex choices renders descriptive experimentation near impossible. Moreover, a growing body of work suggests other artifacts of human cognition (such as preference reversals) may undermine descriptive experimentation once people gain more experience (Hertwig et al. 2004; Tsetsos, Chater, and Usher 2012; Erev et al. 2017; Jessup et al. 2022).

The intractability of complex choice has contributed to the development of experiential study paradigms that adapt complex systems models. Two popular complex systems models are the "multi-arm bandit" (Berry and Fristedt 1985) and the *NK* (Levinthal 1997) paradigms. One experimental paradigm, for example, asks participants to combine *n* shapes that interact *k* ways to create art for aliens. This paradigm enabled the researchers to investigate how people search, explore, and work together during complex choices (Billinger, Stieglitz, and Schumacher 2014; Wu et al. 2018; Billinger et al. 2021, 2022). Although these experiments did not directly control for the bias, Billinger et al. (2014 did demonstrate that anchoring on past outcomes, using prior experience as an endogenous variable, can inform future complex choices). Study participants exerted more (less) effort when underperforming (overperforming) relative to their previous performance.

### Complex Choice in Teams

Teaming adds another element of complexity to any choice. In the alien art task, human dyads coordinated their search efforts during a complex choice without external guidance or resources through a mix of turn-taking and simultaneous moves (Billinger et al. 2022). Teams can also facilitate better complex choices by expanding the sampled regions of a problem space. Each team member brings unique knowledge, preferences, and beliefs to the task, which their be-

havior may reflect. To illustrate, consider an experiment that tasked teams of participants with a simulated entrepreneurship task: managing a lemonade stand (Sommer, Bendoly, and Kavadias 2020). The teams worked to learn about and combine the interacting variables, such as different ingredients, price levels, stand locations, etc., to make the lemonade stand profitable. When the task complexity was low, i.e., with few interactions and relatively low stochasticity, interactive teams did better than teams of individual performers. This result flipped, however, as complexity increased.

Behavioral scientists are not the only researchers working to improve our understanding of teaming during complex choices. The promise of human-machine teams capable of superior performance is one chased by researchers across computer science disciplines. The earliest examples from the literature are likely better described as cooperative efforts—such as airplane piloting and air-traffic control (Hoc 2000). Early success in these and related domains birthed a sizable, ever-evolving sub-discipline that studies the dynamics of trust in automated systems (Lee and See 2004). This research area now spans settings from games that offer insights into the belief states of human teammates (Siu et al. 2021; Chong et al. 2022) to military scenarios that challenge common models of trust dynamics (Gurney, Pynadath, and Wang 2022). The expansive literature on human-AI teams now encompasses topics as diverse as what expectations people hold of AI teammates (Zhang et al. 2021) to endowing AI with the ability to represent the belief states of their human counterparts (Gurney and Pynadath 2022).

Conspicuously missing from research on human-AI teaming is an investigation of the arguably most salient features of human judgment and decision making: heuristics and biases (Gigerenzer and Todd 1999; Kahneman 2011). The obvious argument for such research is that an AI teammate cannot accurately model its human counterpart unless it considers the cognitive tools used by the human during complex choices. Some prior work on complex choice has endogenously tested the role of heuristics and biases, such as Billinger et al. (2014), who demonstrated that anchoring on previous outcomes informs future search behavior. Our work is the first to exogenously control for anchoring and framing effects during complex choices, two of the most studied phenomena in the heuristics and biases literature.

### Anchoring and Framing Effects

Anchoring in human cognition is the tendency of final choices to reflect starting information, even if that information is uninformative (Tversky and Kahneman 1974; Chapman and Johnson 1999; Epley and Gilovich 2006). To illustrate, imagine walking down a store aisle and seeing a sign above a sale item proclaiming a given limit per customer. The anchoring effect suggests that you will buy more or less of the sale item depending on the stated limit despite your intentions or preferences (Wansink, Kent, and Hoch 1998). Researchers have documented similar effects in a host of domains as diverse as retirement savings (Madrian and Shea 2001), courtroom sentencing (Englich, Mussweiler, and Strack 2006), and even in mate selection (Goller et al. 2018).

Framing effects are similarly robust. The basic idea of framing is that a person has a unique reference value, be it a goal or something more random, for the prospects they encounter. They evaluate outcomes relative to those reference values. Evaluations depend on whether the person views the outcome as a loss versus a gain (Kahneman 2011). The classic empirical demonstration of framing asks participants to consider two interventions to stop the spread of a deadly disease. All participants receive a probability-based option and an alternative. Half receive an alternative framed as a gain (lives saved), while the others receive one framed as a loss (lives lost), but that maintains the same ratio. The observation is that the number of participants selecting the probability-based option in the loss-framing condition increases (Tversky and Kahneman 1981).

## Experimental Paradigm

We adapted the rugged-landscape search metaphor first introduced in biology to describe evolutionary selection (Wright 1932) in an interactive, complex choice task (Gurney, King, and Miller 2022). The fundamental idea of this metaphor is that evolutionary pressure, over many generations, guides organisms in a search for an optimal genotype across a genetic landscape where elevation on the landscape represents the fitness of a genotype. This metaphor was generalized in the *NK* model (Kauffman and Levin 1987; Kauffman and Weinberger 1989) and researchers have applied it to an array of settings, perhaps most famously in the management literature (Levinthal 1997). Our task asks study participants to search such a landscape for its global peak (optimum) by tuning dials that each access a different landscape dimension. The complex choice that participants face is when to stop searching, i.e., what dial setting to submit as the best for a given landscape. Participants earned bonuses for making better choices, i.e., finding higher elevations. Each participant completed four dial tuning tasks, the first two on their own and then two more with an AI helper. Importantly, this task lacks strong contextual cues, unlike prior tasks (making artwork for aliens, managing a lemonade stand, farming, etc.).

## Landscapes as Choice Sets

Landscapes are procedurally generated and unique to each participant. They vary in ruggedness, just like mountain ranges, so even a three-dimension task can be challenging for a person to solve. The only way to know the value of a given location in the landscape is by visiting it. Depending on the smoothness of a landscape, a person may, however, make inferences about different regions based on information they previously uncovered. Smoothness is determined by the range of possible slopes between neighboring points in the landscape. For the present experiment, participants completed tasks with one (simple) or four (rugged) peaks and a slope setting such that two neighboring points never differed in value by more than 10% of the absolute range of values, which was 33 units. The order of the tasks was randomized, but participants always did a simple and rugged task alone and with the AI helper.

Landscape topologies are created by randomly placing the global peak in a plane, which for this experiment encompassed  $24 \times 24$  units. For multi-peaked landscapes, additional peaks are randomly placed by generating a list of candidate locations where they maintain a given level of prominence such that valleys separate peaks. The size and dimensionality of the landscape constrain the number of peaks and their elevations. Once all the peaks are placed, a smoothing algorithm ensures that the slopes between the peaks and valleys meet the needed constraints. Our experimental design implemented an elevation boundary of  $[0, 32]$  and peaks ranging in the  $[26, 32]$  interval with the tallest fixed at 32.

The landscape rendering process ensures they are continuous over the edges, as in the right panel of figure 1. This detail means rolling the east and west landscape edges together will form a smooth transition. The same is true of the north and south edges. Thus, each landscape is a toroidal world characterized by peaks and valleys. Shifting that landscape off of the  $[0, 32]$  interval facilitates studying the impact of earlier trials on future behavior. The incentivized task is to find, via exploration, the tallest peak in a given world. In the right panel of figure 1, the global peak is the light yellow point in the upper center of the image. A dark blue valley is below it and surrounded by other peaks.

## User Interface

Participants cannot see the landscapes. Instead, they uncover information about a landscape by “tuning” on-screen dials, as shown in the left panel of figure 1. Each dial is associated with a different dimension of the task. Participants submit tuned settings by clicking the “Evaluate Dial Settings” button. A two-dial experiment, for example, allows a person to explore a three-dimensional world, which is what we implemented for the present study. In this case, conceptually, one dial moves along the east-west and the other moves along the north-south dimension of the landscape. The feedback given when a setting is submitted is the elevation, or choice quality, of that dial combination. Importantly, the landscape dimensions, thus the dials, are interdependent and their relationships are nonlinear.

Tuning happens by clicking on a dial handle and dragging it around the underlying ring. As a participant moves the dial, the letter in the center of the dial changes. Each position on a dial is associated with a letter, which we implemented to simplify feedback. Feedback about the historical dial settings is available in a numbered list below the dials. Participants are able to scroll through their history, see what settings they previously checked, and decide when they are ready to make their choice by clicking the “Finalize Choice” button.

We implemented dials with 24 discrete settings, which matches the coordinate plane used to generate the landscapes. We based our decision to use this size with fixed steps on pretesting that suggested such a landscape would yield a search space (576 unique combinations) that participants are unwilling to exhaustively search given the incentive structure. In other words, for most participants, uncertainty would remain around their choice.

We define two broad types of dial usage: exploring and

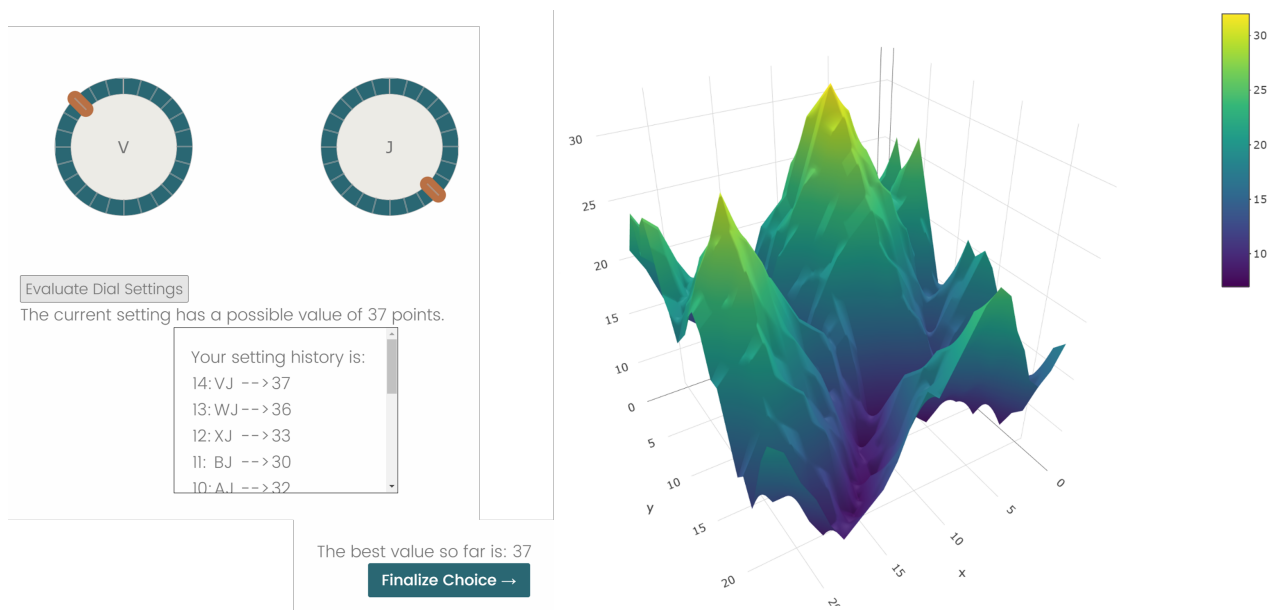


Figure 1: A screenshot from the experiment is on the left, and a landscape rendering is on the right. All participants saw the same dial and feedback interface, but participants in the anchoring treatment saw the best value possible directly above the dials. A two-by-two image of the treatment conditions is available in the online supplement. Participants did not see a rendering of the landscape.

exploiting. Exploration entails looking at regions of the landscape more than two units from any previously observed location. For example, if a participant observed the dial setting [A,A] on their first move, submitting [A,D] or [X,C] on the next move would constitute an explore. Exploitation is when a participant observes a new location less than three units from any previously observed setting.

### AI Helper

After completing two tasks, an AI helper took control of the right dial. It, like study participants, did not have direct access to the landscape, meaning it learned about the topology through sampling. Importantly, participants knew that the AI would help but not how it functions.

The collaborative search happened by the human first adjusting their dial (including leaving it alone) and then clicking the evaluation button. This action initiated the AI helper that took the new participant dial setting and compared two combinations, one using the previous setting of its dial and the other using a new one. The AI was more likely to accept worse values and search distant locations early. As the interaction progressed, however, it decreased the likelihood of accepting worse outcomes and its search distance. After evaluating the combinations and selecting one, the new setting was pushed to the feedback window. Control of the interface was then returned to the study participant who had the option of continuing the search or finalizing the choice.

Even though the task is challenging and relatively time-consuming for humans, it is simple for most AI. Our research questions are about human behavior when working with an AI helper, not developing better helpers, so we chose

to implement a simple agent. The helper uses a stochastic model to determine the likelihood of accepting a worse setting combination and whether to make a big or small leap across the landscape. This model is an adaptation of a simulated annealing algorithm that ensured richer interactions with the participants. It also eliminated unwanted variance in the interactions, thus reducing the needed sample size. Lastly, because it could do the task almost instantaneously, we implemented a brief feedback delay to give the sense that the helper was “working.”

### Experimental Manipulations

The two-by-two design crosses gain-loss framing with anchoring. Participants in the gain frame attempt to gather as many points as possible, which go towards their bonus payment, so the feedback is positive. In the loss frame, participants attempt to reduce the number of points lost, so the feedback is negative. The left panel of figure 1 depicts a gain frame. The task is to maximize the value in both gain and loss framing. Recall that the landscapes are randomly shifted away from the  $[0, 32]$  interval on which they were drawn. Values cover a 33-unit interval in the gain frame on  $[0, 100]$ , while loss frame values are on  $[-100, 0]$ . Participants in the anchoring treatment received a message immediately above the dials that informed them of the best possible dial setting’s value, i.e., as good as they could do on a given task. Participants without an anchor had to discover what constituted a “good” outcome for each task.

Summarizing the present study: each participant completed four incentivized complex choice tasks, two alone and two with an AI helper, in which they tuned dials to find opti-

mal settings. We did not inform participants of the underlying choice landscapes, which were procedurally drawn and unique to each task. Instead, they were left to uncover insights about the space through experience. Participants first did the two solo tasks, but the order of the one and four peaked landscapes was random in both the solo and team tasks. Participants were also randomized into a loss or gain framing treatment conditions. The tuning values of the loss treatment were negative, and participants worked to diminish the number of lost points. Conversely, tuning values of the gain treatment were positive, and participants worked to gain points. The gain-loss framing treatment was crossed with an anchoring treatment in which half of the participants were informed of the best they could find on each task. All participants worked with the same AI helper, but it was stochastic, meaning that even if two participants made the same choice under similar contexts, the helper may have made different choices. Together, these experimental manipulations facilitate studying human-AI teams when a person is operating under a cognitive heuristic or bias.

## Data and Analyses

We recruited 400 participants via Prolific Academic to complete a study on decision making using dials. We only recruited participants with an approval rate of 100 and who could complete the study on a desktop computer. In the months leading up to data collection, there was an influx of new workers on the platform that significantly skewed the composition of the worker population. Thus, we also restricted participation to workers who joined the platform prior to the influx.<sup>1</sup> Funding and IRB constraints restricted us to recruiting English-speaking American citizens. According to Prolific Academic, this left just over 20,000 eligible workers, roughly  $\frac{1}{6}$  of the research pool at recruitment time. The study took about 15 minutes to complete and paid \$2.00 plus a bonus of up to \$2.00 based on effort. The resulting payments yielded an average pay of \$18.40 per hour.

Of the 400 participants who completed the experiment, 172 participants identified as male, 218 as female, and eight as other. The average age of a participant was 32 years. 203 participants indicated they were college graduates with a four-year degree or higher. We dropped two observations from the data set because the participants failed to complete the study. One participant was a considerable outlier: evaluating 617 settings for one of their challenges, more than the complete set of combinations. Moreover, this is three times the effort of the next most ambitious participant. Since this behavior would likely have an outsized impact on any analysis, we removed this participant from the data set, leaving 397 observations.

We report participants' solo efforts in a companion paper (Gurney, Miller, and Pynadath 2023). We found that when working alone, the four-peak landscapes were more challenging than the one-peak, that both anchoring and framing yielded significant main effects on participants' choice

<sup>1</sup><https://www.prolific.co/blog/we-recently-went-viral-on-tiktok-heres-what-we-learned> (We also added this blog to the Wayback Machine at [web.archive.org](http://web.archive.org)).

quality, and that doing the one-peak task first was correlated with better outcomes on the four-peak task. Participants in the loss frame submitted more settings, all else equal, and spent more time exploiting versus exploring. Having an explicit anchor, on the other hand, was correlated with submitting fewer dial settings and did not significantly alter participants' search strategies.

## Complex Choice with the AI Helper

Each landscape is unique and, because the space of the landscape is fixed, adding more peaks increases the average elevation. Thus, comparing raw performance on the landscapes does not fully depict the impact of working with the AI helper. We account for this by dividing each landscape score by the landscape's average elevation to create an adjusted score. On average, participants' total adjusted score for the two choice tasks was worse when they worked with the AI helper ( $M = 3.613$ ,  $SD = 0.717$ ) than on their own ( $M = 3.718$ ,  $SD = 0.820$ ). This difference in total adjusted score (0.105, 95% *CI* [0.036, 0.175]) was statistically significant ( $t(396) = 2.985$ ,  $p = 0.003$ ) per a paired sample t-test. As illustrated by figure 2, the difference was primarily driven by participants in the loss framing. Whether there was an anchor present ( $t(90) = 3.464$ ,  $p < 0.001$ ) or not ( $t(101) = 2.270$ ,  $p = 0.025$ ), the difference in adjusted scores for participants in the loss framing was significant.

A factorial ANOVA predicting the difference between solo and team efforts in the total adjusted score using the experimental treatments as interacting independent variables revealed a significant main effect for framing but not anchoring or the interaction. Thus, a two-sample t-test is sufficient for comparing the effect of framing on performance. Participants in the loss framing condition achieved significantly better ( $t(394.420) = 2.230$ ,  $p = 0.026$ ) scores than their counterparts ( $M_{\text{difference}} = 0.156$ , 95% *CI* [0.185, 0.294]).

Participants did worse on the four- ( $M = 1.522$ ,  $SD = 0.259$ ) versus one-peak ( $M = 2.091$ ,  $SD = 0.554$ ) landscape. This difference (0.569, 95% *CI* [0.521, 0.616]) was significant ( $t(396) = 23.431$ ,  $p < 0.001$ ). A factorial ANOVA predicting the score difference between the two landscape types using the experimental treatments as interacting independent variables revealed no main effects for anchoring ( $F(1, 393) = 2.792$ ,  $p = 0.100$ ) and framing ( $F(1, 393) = 0.128$ ,  $p = 0.721$ ), nor an interaction effect ( $F(1, 393) = 0.342$ ,  $p = 0.559$ ).

Relative to their solo efforts ( $M = 2.175$ ,  $SD = 0.631$ ), participants did worse on the one-peak landscape when working with the AI ( $M = 2.091$ ,  $SD = 0.554$ ), a difference (0.084, 95% *CI* [0.025, 0.144]) which was statistically significant ( $t(396) = 2.773$ ,  $p = 0.006$ ). This did not hold for the four-peak landscapes ( $t(396) = 1.342$ ,  $p = 0.181$ ). A factorial ANOVA predicting the difference between solo and team efforts in the adjusted score for the one-peak landscape using the experimental treatments as interacting, independent variables revealed a significant main effect for the framing treatment condition but not for the anchoring or the interaction. Thus, a two-sample t-test is sufficient for comparing the effect of framing on performance on the




Figure 2: Adjusted score accounts for variance in the landscapes and the increase in average elevation created by adding more peaks. The difference in score is simply the total score achieved during the solo effort minus the total score achieved while working with the AI helper.

one-peak landscapes. Participants in the loss framing condition achieved significantly better ( $t(393) = 2.010, p = 0.045$ ) scores than their counterparts ( $M_{difference} = 0.122, 95\% CI [0.003, 0.242]$ ).

Search duration (number of submitted settings) and the explore/exploit trade-off (fraction of submitted settings that explore new regions, i.e. search strategy) serve as the main behavioral metrics. On average, participants searched less, i.e. submitted fewer dial settings across the two tasks, when they worked with the AI helper ( $M = 30.368, SD = 27.786$ ) than on their own ( $M = 43.202, SD = 41.521$ ). This difference in search duration ( $12.834, 95\% CI [9.512, 16.155]$ ), per a paired-sample t-test, was statistically significant ( $t(396) = 7.596, p < 0.001$ ). The fraction of those submissions that were explores while working with the AI helper ( $M = 0.612, SD = 0.243$ ) was greater than when they were working on their own ( $M = 0.447, SD = 0.232$ ). This difference in the explore/exploit trade-off ( $-0.165, 95\% CI [-0.186, -0.144]$ ) was also statistically significant ( $t(396) = -15.3456, p < 0.001$ ). The absolute number of submissions that were explores while working with the AI helper ( $M = 13.584, SD = 6.244$ ), however, was not statistically different from the solo effort ( $M = 13.411, SD = 8.617; t(396) = -0.442, p = 0.659$ ). This suggests that participants invested less effort in fine-tuning their submissions, perhaps because they anticipated that the AI would do so for them.

### Search Duration with Anchoring and Framing Effects

Comparing the difference in search duration for

participants' solo and team efforts under the different treatment conditions provides insight into how anchoring and framing effects impact human-AI teams. A factorial ANOVA predicting the difference in search duration (solo effort minus team effort duration) using the experimental treatments as interacting, independent variables revealed significant main effects for both treatment conditions but not the interaction, thus we removed the interaction from the model. The resulting model suggests strong main effects for both anchoring ( $F(1, 394) = 14.971, p < 0.001$ ) and framing ( $F(1, 394) = 16.751, p < 0.001$ ). A post hoc Tukey test showed that the anchoring ( $-12.611, 95\% CI [-21.047, -4.174], p_{adj} < 0.001$ ) and framing ( $13.331, 95\% CI [4.892, 21.769], p_{adj} < 0.001$ ) effects were significant at the  $p < 0.01$  level, although search duration moved the opposite way.

Building on these insights, using an independent sample t-test, loss frame participants adjusted their effort downwards ( $M = -19.964, SD = 39.838$ ) significantly more than gain frame participants ( $M = -6.088, SD = 24.808; t(318.21) = -4.319, p < 0.001$ ) when they started working with the AI. This is the inverse of the anchoring effect: participants with explicit anchors adjusted their effort downwards ( $M = 6.449, SD = 29.426$ ) significantly less than those without ( $M = 19.060, SD = 36.341; t(382.16) = 3.804, p < 0.001$ ). As illustrated in figure 3, the anchoring effect seems to have tempered the framing effect.

On average, participants searched less in each successive task. However, the only significant drop was between their second solo effort and the first team effort ( $M_{difference} = -5.108; t(396) = 5.003, p < 0.001$ ). The difference between the first and second solo ( $M_{difference} = -1.690; t(396) = 1.654, p = 0.099$ ) as well as the first and second team ( $M_{difference} = -0.927; t(396) = 1.300, p = 0.194$ ) efforts were not significant. This points to participants expecting that the AI would reduce the effort they needed to invest in the task to perform well. For gain-frame participants, this appears to be true: as they exerted less effort while working with the AI ( $M_{difference} = -6.088; t(203) = -3.505, p < 0.001$ ), but averaged the same adjusted scores ( $M_{difference} = -0.029; t(203) = -0.574, p = 0.567$ ). Loss frame participants, however, appear to have over-adjusted their effort ( $M_{difference} = -19.964; t(192) = -6,962, p < 0.001$ ) such that their adjusted scores suffered significantly ( $M_{difference} = -0.186; t(192) = -3.877, p < 0.001$ ). For the anchoring treatment, whether a participant saw an anchor ( $M_{difference} = -6.449; t(195) = -3.068, p = 0.002$ ) or not ( $M_{difference} = -19.060; t(200) = -7.436, p < 0.001$ ), the difference in their effort was significantly less when they worked with the AI helper. This was not correlated with a significantly lower adjusted score during the AI tasks for the no-anchor participants ( $M_{difference} = -0.090; t(200) = -1.730, p = 0.090$ ). It was, however, correlated with a significantly lower adjusted score for the anchor participants ( $M_{difference} = -0.121; t(195) = -2.539, p = 0.012$ ).

### Average Search Duration Difference Between Solo and Team Efforts by Condition

95% CI

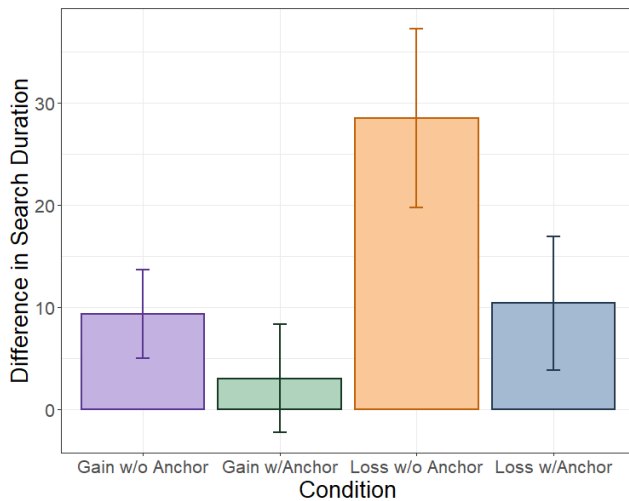


Figure 3: Search duration difference is the total number of submitted settings during a participant’s solo effort minus the total number of submitted settings while working with the AI helper.

### Search Strategy with Anchoring and Framing Effects

As noted, the change in search duration echoed through to the strategy participants used when they began working with the AI helper. This shorter search duration meant a higher fraction of submissions that were explorations. A factorial ANOVA predicting the difference in search strategy (solo effort minus team effort strategy) using the experimental treatments as interacting, independent variables revealed no main effects for anchoring ( $F(1, 393) = 1.934, p = 0.165$ ) and framing ( $F(1, 393) = 0.684, p = 0.409$ ), nor an interaction effect ( $F(1, 393) = 1.150, p = 0.284$ ). Again, we interpret this as suggesting that the participants anticipated that the AI would reduce the effort they needed to expend in tuning the dials.

### Discussion

Aversion to algorithmic choice aids is a well-documented phenomenon (Dietvorst, Simmons, and Massey 2015; Burton, Stein, and Jensen 2020). The proliferation of both passive and adaptive algorithms in every corner of life, however, is leaving people increasingly accepting of them (Logg, Minson, and Moore 2019). Historically, the experimental methods used to study human-algorithm teams relied on reducing the complexity of choices or studying them *in situ*. Although both methods have merit, they also have major shortcomings, such as not generalizing to complex choices or choices made in different contexts. The method we developed here maintains choice complexity and is sufficiently abstract to facilitate the generalization of findings.

Study participants generally made worse choices when they worked with the AI helper than when they worked

alone. This result was primarily driven by participants in a loss frame who dramatically decreased their effort once they started working with the AI helper. However, participants in the gain frame performed about the same on their own as with the AI. The decrease in participants’ effort was primarily seen in fine-tuning a choice by exploiting local information rather than exploring new options. The only participants that did not significantly reduce their effort when they started working with the AI helper were those in the gain frame *and* with an explicit anchor, i.e. they knew the value of the best possible choice. These results suggest that people were possibly over-reliant on the AI helper, assumed that it was better at the task than it was, or, relatedly, their lack of knowledge about how it functioned hindered their ability to team with it. These possibilities point to interesting and important topics for future research. Additionally, they suggest that the well-documented phenomenon of algorithm aversion may not be a stable aspect of human cognition.

### Conclusion

AI helpers are making their way into every aspect of life—and people appear to be more willing than ever to allow them to help. Historically, the complex choices for which AI can be most helpful have been reduced to simpler analogs for experiments or studied *in situ*. Although simple choice experiments can easily isolate the effects of heuristics and biases, they ignore the fact that many choices are characterized by astounding complexity. *In situ* experimentation overcomes this, but it may produce results that do not generalize. As demonstrated, the dial tuning task maintains the ability to study well-documented cognitive artifacts during complex choices plus, we argue, it provides general results (and data) that AI systems can use. Specifically, the data generated by this experiment are easy to translate into layered image representations that lend themselves to deep learning models. With sufficient data points, such models could learn to discriminate between biased and unbiased behavior.

### Acknowledgments

Part of the effort behind this work was sponsored by the Defense Advanced Research Projects Agency (DARPA) under contract number W911NF2010011. The content of the information does not necessarily reflect the position or the policy of the U.S. Government or the Defense Advanced Research Projects Agency, and no official endorsements should be inferred. The work was or is also sponsored by the U.S. Army Research Laboratory (ARL) under contract number W911NF-14-D-0005. Statements and opinions expressed and content included do not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

### References

Berry, D. A.; and Fristedt, B. 1985. Bandit problems: sequential allocation of experiments (Monographs on statistics and applied probability). London: Chapman and Hall, 5(71-87): 7–7.

- Billinger, S.; Benincasa, S.; Baumann, O.; Kretschmer, T.; and Schumacher, T. R. 2022. Learning to search collaboratively: how dyads overcome complexity and misaligned incentives in imperfect modular decompositions. *Industrial and Corporate Change*.
- Billinger, S.; Srikanth, K.; Stieglitz, N.; and Schumacher, T. R. 2021. Exploration and exploitation in complex search tasks: How feedback influences whether and where human agents search. *Strategic Management Journal*, 42(2): 361–385.
- Billinger, S.; Stieglitz, N.; and Schumacher, T. R. 2014. Search on rugged landscapes: An experimental study. *Organization Science*, 25(1): 93–108.
- Bornstein, R. F.; and D’agostino, P. R. 1992. Stimulus recognition and the mere exposure effect. *Journal of personality and social psychology*, 63(4): 545.
- Burton, J. W.; Stein, M.-K.; and Jensen, T. B. 2020. A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2): 220–239.
- Chapman, G. B.; and Johnson, E. J. 1999. Anchoring, activation, and the construction of values. *Organizational behavior and human decision processes*, 79(2): 115–153.
- Chong, L.; Zhang, G.; Goucher-Lambert, K.; Kotovsky, K.; and Cagan, J. 2022. Human confidence in artificial intelligence and in themselves: The evolution and impact of confidence on adoption of AI advice. *Computers in Human Behavior*, 127: 107018.
- Dawes, R. M. 1979. The robust beauty of improper linear models in decision making. *American psychologist*, 34(7): 571.
- Dawes, R. M.; and Corrigan, B. 1974. Linear models in decision making. *Psychological bulletin*, 81(2): 95–106.
- Dawes, R. M.; Faust, D.; and Meehl, P. E. 1989. Clinical versus actuarial judgment. *Science*, 243(4899): 1668–1674.
- Dietvorst, B. J.; Simmons, J. P.; and Massey, C. 2015. Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1): 114.
- Englich, B.; Mussweiler, T.; and Strack, F. 2006. Playing dice with criminal sentences: The influence of irrelevant anchors on experts’ judicial decision making. *Personality and Social Psychology Bulletin*, 32(2): 188–200.
- Epley, N.; and Gilovich, T. 2006. The anchoring-and-adjustment heuristic: Why the adjustments are insufficient. *Psychological science*, 17(4): 311–318.
- Erev, I.; Ert, E.; Plonsky, O.; Cohen, D.; and Cohen, O. 2017. From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological review*, 124(4): 369.
- Gigerenzer, G.; and Todd, P. M. 1999. *Simple heuristics that make us smart*. Oxford University Press, USA.
- Goller, J.; Leder, H.; Cursiter, H.; and Jenkins, R. 2018. Anchoring effects in facial attractiveness. *Perception*, 47(10-11): 1043–1053.
- Gurney, N.; King, T.; and Miller, J. 2022. An Experimental Method for Studying Complex Choices. In Stephanidis, C.; Antona, M.; Ntoa, S.; and Salvendy, G., eds., *HCI International 2022 – Late Breaking Posters*, 39–45. Springer Nature Switzerland. ISBN 978-3-031-19679-9.
- Gurney, N.; Miller, J.; and Pynadath, D. V. 2023. The Role of Heuristics and Biases in Complex Choice. <https://doi.org/10.21203/rs.3.rs-2472194/v1>.
- Gurney, N.; and Pynadath, D. V. 2022. Robots with Theory of Mind for Humans: A Survey. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 993–1000. IEEE.
- Gurney, N.; Pynadath, D. V.; and Wang, N. 2022. Measuring and Predicting Human Trust in Recommendations from an AI Teammate. In *International Conference on Human-Computer Interaction*, 22–34. Springer.
- Hertwig, R.; Barron, G.; Weber, E. U.; and Erev, I. 2004. Decisions from experience and the effect of rare events in risky choice. *Psychological science*, 15(8): 534–539.
- Hoc, J.-M. 2000. From human–machine interaction to human–machine cooperation. *Ergonomics*, 43(7): 833–843.
- Jacowitz, K. E.; and Kahneman, D. 1995. Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21(11): 1161–1166.
- Jessup, R. K.; Busemeyer, J. R.; Dimperio, E.; Homer, J.; and Phillips, A. 2022. Choice is a tricky thing: Integrating sophisticated choice models with learning processes to better account for complex choice behavior. *Decision*.
- Kahneman, D. 2011. *Thinking, fast and slow*. Macmillan.
- Kauffman, S.; and Levin, S. 1987. Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology*, 128(1): 11–45.
- Kauffman, S. A.; and Weinberger, E. D. 1989. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of theoretical biology*, 141(2): 211–245.
- Lee, J. D.; and See, K. A. 2004. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1): 50–80.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management science*, 43(7): 934–950.
- Logg, J. M.; Minson, J. A.; and Moore, D. A. 2019. Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151: 90–103.
- Madrian, B. C.; and Shea, D. F. 2001. The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly journal of economics*, 116(4): 1149–1187.
- Schwarz, N. 2004. Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology*, 14(4): 332–348.
- Simon, H. A. 1956. Rational choice and the structure of the environment. *Psychological review*, 63(2): 129.
- Siu, H. C.; Peña, J.; Chen, E.; Zhou, Y.; Lopez, V.; Palko, K.; Chang, K.; and Allen, R. 2021. Evaluation of human-AI teams for learned and rule-based agents in Hanabi.

*Advances in Neural Information Processing Systems*, 34: 16183–16195.

Sommer, S. C.; Bendoly, E.; and Kavadias, S. 2020. How do you search for the best alternative? Experimental evidence on search strategies to solve complex problems. *Management Science*, 66(3): 1395–1420.

Tsetos, K.; Chater, N.; and Usher, M. 2012. Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109(24): 9659–9664.

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

Tversky, A.; and Kahneman, D. 1981. The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481): 453–458.

Wansink, B.; Kent, R. J.; and Hoch, S. J. 1998. An anchoring and adjustment model of purchase quantity decisions. *Journal of Marketing Research*, 35(1): 71–81.

Wright, S. 1932. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proceedings of the Sixth International congress of Genetics*. na.

Wu, C. M.; Schulz, E.; Speekenbrink, M.; Nelson, J. D.; and Meder, B. 2018. Generalization guides human exploration in vast decision spaces. *Nature human behaviour*, 2(12): 915–924.

Zajonc, R. B. 1968. Attitudinal effects of mere exposure. *Journal of personality and social psychology*, 9(2p2): 1.

Zhang, R.; McNeese, N. J.; Freeman, G.; and Musick, G. 2021. "An Ideal Human" Expectations of AI Teammates in Human-AI Teaming. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3): 1–25.