Human Visuomotor Decision Making Under Risk is Not Optimally Adapted to Movement Variability

Tyler J. Adkins[1], Richard L. Lewis[1,2,3], Taraz G. Lee[1]*

[1]Department of Psychology, University of Michigan
[2]Department of Linguistics, University of Michigan
[3]Weinberg Institute for Cognitive Science, University of Michigan

*Corresponding author information:
tarazlee@umich.edu
Department of Psychology, University of Michigan
530 Church St., Ann Arbor, MI, 48109, USA

Abstract

From financial choice to movement planning, many human behaviors can be understood as the outcomes of decision processes. Results from studies of decision making under risk in behavioral economics, where risk is manifest in varying probabilities of payoffs in presented options, are taken to indicate that humans fail to maximize expected objective value. In contrast to this research, a recent line of work in psychophysics argues that when such decisions are framed as risky visuomotor decisions, where payoff probabilities in part depend on the variance of reaching end points, people make approximately optimal decisions which maximize expected objective value. We rigorously test this claim by formally comparing three computational models in terms of their ability to predict human behavior in incentivized reaching tasks. Across two experiments, we find that human reaching outcomes are better described by a simple satisficing heuristic than objective or subjective expected-value-maximizing strategies that are optimally adapted each individual's variance in reaching end points, especially when the risk of financial loss is high. The present work provides evidence against the claim that human visuomotor decision making under risk is optimal given the bounds imposed by motor variance, and demonstrates the fruitfulness of including alternative models in analyses of human behavior.

Introduction

Many human behaviors can be understood as the outcomes of risky decision processes, from financial choice to movement planning[1]. A large body of research in behavioral economics indicates that human decision making under risk is shaped by cognitive biases and heuristics, which lead to failures to maximize expected objective value[2]. For example, people tend to assign outsized weight to potential penalties, a tendency known as loss aversion[3–6]. Heuristics, on the other hand, are simple rules or strategies that people leverage to make quick and minimally effortful, yet often satisfactory decisions[2,7–9]. For example, in the face of large and uncertain decision spaces, people often display *satisficing* behavior in which they select an option that is good enough, but potentially suboptimal from an objective standpoint[10,11]. Theoretical explanations of various departures from normative optimality differ in the degree to which they reject or embrace utility maximization in their analysis, but the empirical departures themselves are widely accepted[12].

A recent line of interesting work argues that when risky financial decision-making problems are framed as risky visuomotor decision-making problems, where uncertainty is manifest in the natural variance of reaching actions, people make approximately optimal decisions which maximize expected objective value (Figure 1a)[1,13–16]. This claim is supported by analyses which show a close fit between human performance in an incentivized reaching task and the predictions of an optimal statistical decision model which maximizes expected objective value. This methodology has significant theoretical appeal because, armed with only an estimate of a participants' movement variability, the optimal model generates detailed quantitative predictions about human reaching behavior across a multitude of task conditions. Several follow-up studies have provided evidence for human departures from optimal under specific task conditions, but most assume that the original research provides strong evidence for approximately optimal visuomotor decision making under risk and assume optimality as a baseline condition[17–20]. Our aim is to rigorously test whether human visuomotor decision making behavior is generated as an optimal response to the risk imposed by movement variability, or is better characterized as a process shaped by a cognitive bias or heuristic.

In the present work, we improve on the visuomotor decision-making methodology mentioned above, both empirically and analytically, by devising a version of the canonical reaching task used in prior work with a wider range of incentive values, and by formally comparing three distinct computational models in terms of their ability to predict human behavior. In addition to the optimal model used in prior work, we consider a model incorporating loss aversion and a model incorporating a simple satisficing heuristic. Across two experiments, we find that our participants' reaching outcomes are best described by a simple satisficing heuristic, especially when the risk of financial loss is high. This work contributes to the study of human behavior in three ways:

(1) It provides evidence against the claim that people make visuomotor decisions under risk that are optimally adapted to their reach variability, showing that they more likely rely on a simple satisficing heuristic.

(2) It shows that human performance in a visuomotor decision-making task under risk diminishes (relative to optimal) as the financial risk increases.

(3) It demonstrates the utility of including alternative models in analyses of human behavior.

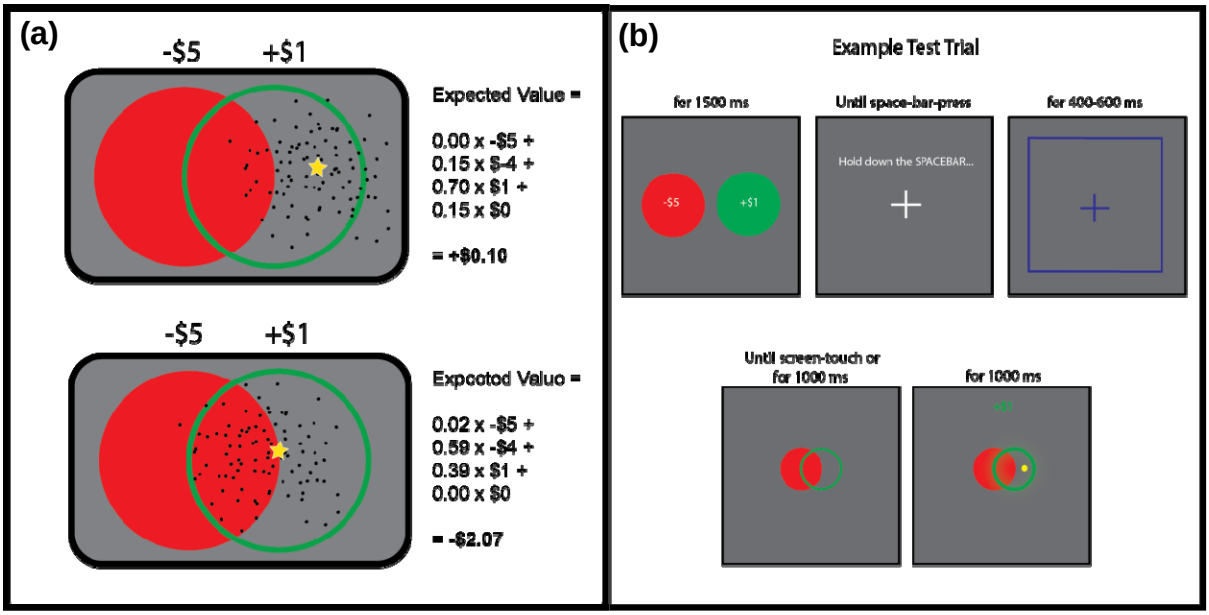

Figure 1. Decision making under risk in an incentivized reaching task. **(a)** Selecting a reaching aim-point is equivalent to accepting a financial gamble. For a given aim-point, the probabilities of gain and loss are imposed by the participant's movement endpoint variability (represented by the distribution of small black points). Thus, each aim-point (represented by the gold stars) has an associated expected value, and the optimal strategy would be to select the aim-point which maximizes expected value. **(b)**

In a typical trial of the incentivized reaching task, a participant is first presented with information about the potential penalties and rewards. Next, they see a fixation cross, which remains until they initiate the trial by pressing and holding down the space bar. After a randomized delay, reward and penalty circles appear and the participant has 1s to reach out and touch the screen. The trial concludes with feedback about the participant's earnings for 1s.

Results

We investigate whether loss aversion or satisficing[3,11] lead human visuomotor decision making to deviate from an optimal adaptation to motor variability in an incentivized reaching task. Furthermore, we expect that possible deviations from optimality will increase as the penalty magnitude increases, in light of research showing that humans often choke under pressure in motor and cognitive tasks with large performance-based incentives[5,21,22]. In what follows, we present human data and modeling results from two experiments using a version of the reaching task with non-cumulative earnings and a wide range of payoff conditions varying by trial (Experiment 1) and a version with cumulative earnings and a narrow range of payoff conditions varying by block (Experiment 2). The results suggest that human visuomotor decision making under risk is best accounted for by a satisficing motor planning heuristic that is not optimally adapted to the risk imposed by the variability in reaching actions. We furthermore find no evidence of loss aversion in the task.

Experiment 1

In Experiment 1, 23 human participants performed an incentivized reaching task in which they quickly reached to targets on a touch-screen monitor. Participants gained and lost money depending on whether they touched a green "reward" circle or a red "penalty" circle, respectively (Figure 1b). The incentivized reaching task had five payoff conditions ([-$0, +$1]; [-$1, +$1]; [-$5, +$1]; [-$3, +$3]; and [-$15, +$5]) and two spatial conditions (32 or 44 pixel spread between the two circle centers, each with a radius of 32 pixels), which varied randomly by trial. We estimated each participant's movement variability and trait-level loss aversion in

separate tasks (See Methods). At the end of experiment, one trial was selected at random from each block and participants gained and lost bonus earnings based on their performance on these randomly selected trials.

Computational models. To better understand the processes that guide human visuomotor decision making under risk, we examine three computational models of human performance in the incentivized reaching task. Each model embodies a unique hypothesis about how people decide where to aim their reaching movements. The models generate precise quantitative predictions about participants' mean reaching outcomes (Figure 2a,c), which we compare to observed trial-level reaching outcomes to compute the models' fits to human data. We define a model's fit as the summed negative log-likelihood of the observed data given the model. We then formally compare models in terms of their fit to human data using a log-likelihood ratio test for non-nested models[23] (see Methods). We hypothesized that our participants' performance would deviate from the optimal expected-value-maximizing model as financial risk increases.

The first model we examined is the optimal expected-value-maximizing model used in prior work[1,13,14,16]. According to this model, a participant will integrate an estimate of their movement variability with the stimuli and the payoff information to estimate the expected values associated with different aim-points. Then, the participant selects the aim-point with the maximum expected value and reaches to that point. Formally, the model predicts that participants select aim-points (x,y) which maximizes the following equation:

$$EV(x, y, \sigma) = V_p * \mathbb{P}(C_P \mid x, y, \sigma) - V_R * \mathbb{P}(C_R \mid x, y, \sigma). \qquad (1)$$

V_P is the value of the penalty, V_R is the value of the reward, $\mathbb{P}(C_P \mid x, y, \sigma)$ is the probability of touching the penalty circle, given the aim-point (x,y) and the participant's movement variability (σ), and $\mathbb{P}(C_R \mid x, y, \sigma)$ is the probability of touching the reward circle, given the aim-point and the participant's movement variability. Using numerical simulations, we derive quantitative predictions about our participants' mean reaching outcomes from this model (see Methods).

The next model we examine is a loss-averse expected utility maximizing model similar to some models explored in prior work[24,25]. This model is identical to the optimal model, except that a participant is supposed to maximize an expected *utility* measure which incorporates their trait-level loss aversion. Formally, the model predicts that participants select aim-points which maximize the following equation:

$$EU(x, y, \sigma, \lambda) = \lambda * V_p * \mathrm{P}(C_P \mid x, y, \sigma) + V_R * \mathrm{P}(C_R \mid x, y, \sigma). \quad (2)$$

All terms in equation 2 are the same as those in equation 1, except that the penalty value is multiplied by the participant's loss-aversion parameter λ, estimated in an independent task[4] (see Methods). As with the optimal model, we use numerical simulations to derive quantitative predictions from the loss averse model about our participants' mean reaching outcomes. For individuals with $\lambda > 1$, the loss-averse model predicts greater shifts away from the penalty region than the optimal model.

The third model that we examine is a satisficing heuristic model (hereafter, the heuristic model). According to this model, a participant's aim-point is determined by the presence/absence of a potential penalty and the width of the pure reward region (the portion of the reward circle that is not overlapping with the penalty circle). Specifically, the model predicts that participants aim at the center of the reward circle when there is no potential penalty and at the center of the 'pure reward region' when there is a potential penalty (regardless of its size). The pure reward region is identified as the portion of the reward circle that is non-overlapping with the penalty circle. It is plausible that participants will judge the center of this pure reward region to be a satisfactory aim-point. Since this model incorporates no unique information about participants, such as their movement variability or loss aversion, the model makes the same predictions for all participants.

The heuristic model best describes aggregate reaching behavior. To determine the decision process that was most likely driving our participants visuomotor decisions under risk, we compared the three models above in terms of their fit to participants' reaching data. Across data from all payoff conditions, participants' reaching behavior was better described by the heuristic model than the optimal model (circles close: Vuong's Z-statistic = 3.68, $p < 0.001$; circles far: Z = 2.34, $p = 0.02$), and the loss-averse model (circles close: Z = 13.9, $p < 0.001$; circles far: Z = 7.13, $p < 0.001$). Our participants' reaching behavior was also better described by the optimal model than the loss-averse model (circles close: Z = 14.9, $p < 0.001$; circles far: Z = 5.52, $p < 0.001$). These results suggest that our participants' visuomotor decisions were more likely guided by the simple satisficing heuristic, rather than the objective or subjective expected-value maximizing strategies.

The heuristic model best describes reaching behavior when financial risk is high. To determine whether reaching behavior was impacted by the level of financial risk involved, we estimated the fit of each model for each penalty to reward ratio separately (Figure 2b). We ignore data from the zero-penalty condition because the models make identical predictions for this condition. When the penalty to reward ratio was low (1:1) and the circles were close together, participants' reaching behavior was described better by the optimal model than the heuristic model ($Z = 7.89$, $p < 0.001$) and the loss-averse model ($Z = 9.04$, $p < 0.001$). However, when the penalty to reward ratio was high (5:1) and the circles were close together, behavior was better described by the heuristic model than the optimal model ($Z = 10.3$, $p < 0.001$) and the loss-averse model ($Z = 12.1$, $p < 0.001$). Likewise, when the penalty to reward ratio was high and the circles were far apart, participants' reaching behavior was described better by the heuristic model than the optimal model ($Z = 3.84$, $p < 0.001$) and the loss-averse model ($Z = 7.14$, $p < 0.001$). And when the penalty to reward ratio was low and the circles were far apart, participants' reaching behavior was described equally well by the heuristic and optimal models ($Z = 0.458$, $p = 0.647$), the heuristic and loss-averse models ($Z = 1.25$, $p = 0.210$), and the optimal and the loss-averse models ($Z = 0.534$, $p = 0.593$). These model comparison results suggest that our participants more likely used the satisficing heuristic to guide their visuomotor decisions, rather than the objective or subjective expected-value maximizing strategy, particularly when financial risk (penalty to reward ratio) was high.

No participant was best described by the optimal model. Although the results above suggest that the heuristic model outperforms the optimal model at the group level, it remained possible that certain individuals were near-optimal in their behavior. To test for this possibility, we examined the model fits for each participant separately. We declared a model as the *best* description of a participant if and only if the model fit their data significantly better than both of the other two models, where statistical significance was determined by Vuong's test. Overall, when the penalty and reward circles were close together, 6/23 participants were best described by the heuristic model and the remaining 13 participants were described equally well by all three models. When the circles were far apart, 3/23 participants were best described by the heuristic model and the remaining 20 participants were described equally well by all three models. These results suggest that our participants were more likely using the simple satisficing heuristic than a strategy which

maximizes expected objective or subjective value. However, the fact that most participants were described equally well by all three models suggests that the visuomotor task used here may not be well suited for determining the true process that guides human visuomotor decision-making under risk.

Participants' efficiency was lower when financial risk was high. We also examined participants' *efficiency* relative to their optimal counterparts, in line with prior work that has used efficiency as a measure[13,14,16,17,24–26]. We define a participant's efficiency as the ratio of the expected value of their aim-point to the expected value of the aim-point of their optimal counterpart (see Methods). We hypothesized that participants' efficiencies will decrease as the penalty to reward ratio increases. We tested this hypothesis using linear mixed effects models of efficiency, separately for each circle-spread condition (32 px, 44 px). When the circles were close together, the model had a significant negative fixed effect of penalty to reward ratio on efficiency ($b = -0.02$, $p < 0.001$; Figure 3a). This effect was not significant when the circles were far apart ($p = 0.062$), perhaps due to the lower level of financial risk in this condition. This result is consistent with the hypothesis that participants' performance relative to optimal diminished with increasing financial risk.

Experiment 2

To ensure that the results of Experiment 1 were not due to unique aspects of our task design (see Methods)[18,19], we conducted a second experiment to serve as a conceptual replication and as a validation set for hypotheses that we developed while analyzing the results of Experiment 1 (in particular, the heuristic model). To this end, 18 additional participants performed a version of the incentivized reaching task used in the original studies by Trommershauser et al.[13,14]. We took special care to ensure that the task parameters matched those from the original studies to test whether our Experiment 1 results were simply due to unique design choices that we made. In this version, participants were exposed to three payoff conditions ([-0, +100], [-100, +100], [-500, +100]) presented in separate blocks. Points accumulated over time and were converted to cash bonuses at a rate of $0.25/1000 points. The models tested were the same as in experiment 1.

The heuristic model again best describes aggregate reaching behavior. We replicated the initial modeling result from Experiment 1. Across data from payoff conditions, participants' reaching behavior was better described by the heuristic model than the optimal model ($Z = 7.40$, $p < 0.001$) and the loss-averse model ($Z = 19.1$, $p < 0.001$). This provides additional evidence that our reaching behavior is more likely driven by a simple satisficing heuristic rather than an objective or subjective expected-value maximizing decision process.

The heuristic model best describes our participants when financial risk is high. We again assessed which model best described our participants' reaching behavior at the various penalty to reward ratios (Figure 2d). When the penalty to reward ratio was low (1:1), participants' behavior was better described by the optimal model than the loss-averse model ($Z = 13.8$, $p < 0.001$) and the heuristic model ($Z = 4.36$, $p < 0.001$). However, when the penalty to reward ratio was high (5:1), participants' behavior was better described by the heuristic model than the optimal model ($Z = 7.55$, $p < 0.001$) and the loss-averse model ($Z = 15.4$, $p < 0.001$). Although performance appeared to be near-optimal when the penalty and reward to reward ratio was low, large penalty values seemed to elicit behavior that was more akin to a satisficing heuristic.

Only one participant was best described by the optimal model. Overall, 9/18 participants in Experiment 2 were described best by the heuristic model, only one was described best by the optimal model, and the remaining 8 participants were described equally well by all three models. Collectively, these modeling results corroborate the results from Experiment 1 and increase our confidence that participants were most likely using the simple satisficing heuristic, especially when performance contingent-incentives were largest.

Participants' efficiency was lower when financial risk was high. We replicated the finding from Experiment 1 using mixed effects models which showed that participants' efficiencies decrease as financial risk increases ($b = -0.01$, $p = 0.04$) (Figure 3b). This result is consistent with the hypothesis that participants' performance relative to optimal diminished as the financial risk increased.

Participants are sensitive to payoffs but not their movement variability. We also checked other aspects of the data to assess whether participants' reaching behavior was driven by expected value maximization, as suggested by the optimal model. If the optimal model is an accurate approximation of human behavior, then participants should reach further from the penalty circle with increases in the penalty to reward ratio (Figure 2a,c). We test this prediction using linear mixed effects models and model comparison (See Methods section for more details), which showed a small fixed effect of penalty to reward ratio for experiment 1 ($b = 0.08$, $p = 0.003$) and experiment 2 ($b = 0.02$, $p < 0.001$). These results suggest that participants adapted slightly to the changing payoffs by aiming further from the penalty circle as the penalty to reward ratio increased, consistent with an expected-value-maximizing model. These shifts in endpoint were smaller than predicted by the optimal model.

According to the optimal model, participants should also reach further from the penalty circle the greater their movement variability. We tested this prediction using a rank-order correlation between participants' mean reaching outcomes and our estimates of their individual movement variability. We do not find evidence for such a correlation in Experiment 1 (32px separation: rho = -0.22, $p = 0.37$; 44px separation: rho = -0.03, $p = 0.90$) or Experiment 2 (rho = 0.15, $p = 0.57$). This suggests that participants may be insensitive to their own movement variability, a result which is inconsistent with the optimal model.

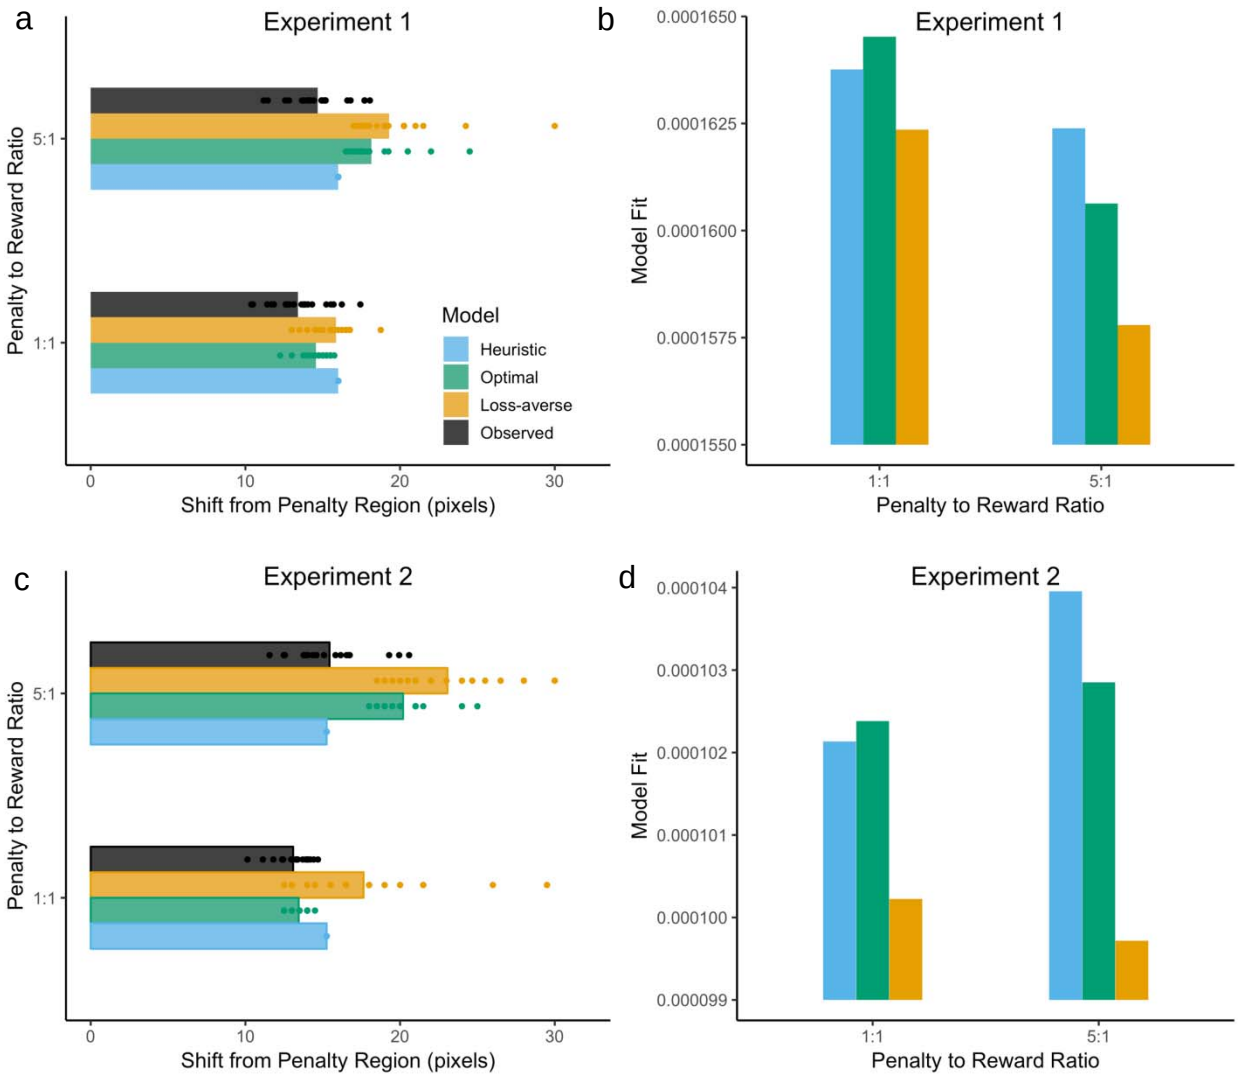

Figure 2. Model predictions and fits by penalty condition. Panels **(a)** and **(c)** show group mean predicted and observed horizontal shifts in aim-point away from the penalty circle for each penalty condition in Experiment 1 and Experiment 2, respectively. The points represent subject means. Across both experiments, we observed that participants aim closest to the heuristic model prediction when the penalty to reward ratio is high (5:1). Panels **(b)** and **(d)** show the fits of each model to reaching data from each penalty condition for Experiment 1 and Experiment 2, respectively. Model fit here is defined for visualization purposes as -1 over the sum of the log-likelihoods of the data, given the model (see Methods). Across both experiments, we find that the heuristic model describes our participants' behavior best when the penalty to reward ratio is high. Data from experiment 1 when the circles were far apart is not shown here.

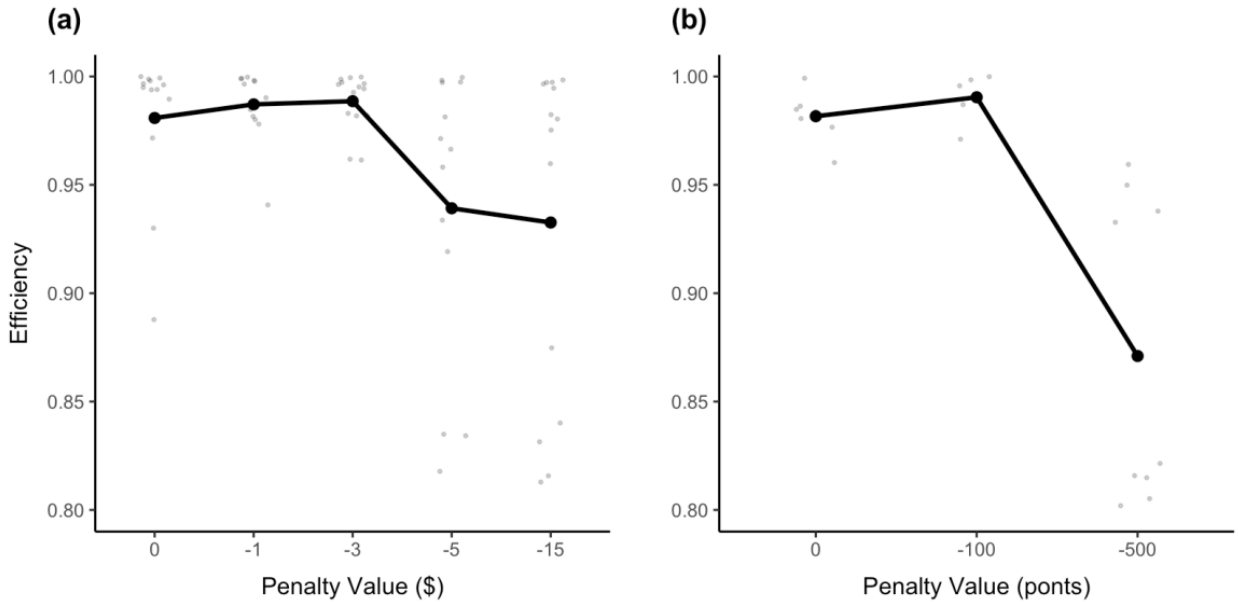

Figure 3. Efficiency by penalty value. Efficiency is defined as the ratio of the expected values of a participant's aim-point and the aim-point of their optimal counterpart. Panel **(a)** shows the group mean efficiencies by penalty value in Experiment 1. Panel **(b)** shows the group mean efficiencies by penalty value in Experiment 2. Each small grey point is an efficiency score for one participant. Across both experiments, we observe that participants' efficiency decreases as the ratio of the penalty to reward magnitude increases.

Discussion

In contrast to prior work, we provide evidence against the claim that people use a strategy optimally adapted to payoffs and movement variability to guide their visuomotor decision-making. While previous studies found that human performance in an incentivized reaching task was close to such an optimal strategy[1,13,14,16], we sought stronger and more sensitive tests of such optimality. We therefore compared human performance to the predictions of three different computational models, all of which embody specific hypotheses about how people perform the task. Of particular interest was the fit of an optimal model compared to two different models that incorporate a known bias (loss aversion) and a known heuristic (satisficing). Across two experiments, we found that the satisficing heuristic model described our participants' behavior better than the optimal and loss-averse models, especially when financial risk was high (Figure 2b,d). This finding is consistent with literature showing that people do not make objectively optimal decisions when faced with large and uncertain decision-spaces[10,11].

We also find that participants' performance relative to optimal diminished as the financial risk increased (Figure 3a,c). A possible explanation for this finding is that our participants "choke" under heightened performance pressure and adopt a suboptimal strategy as a result. This

explanation is plausible in light of literature showing that human performance in motor and cognitive tasks sometimes diminishes paradoxically when performance-based incentives are especially large [5,21,22]. Future research should investigate whether human visuomotor decisions under risk systematically deviate from optimal in high-pressure situations.

Prior work has similarly shown that human visuomotor decision making under risk is suboptimal in some variants of the incentivized reaching task. For example, human performance has been shown to be suboptimal with rapidly varying payoff conditions (like our Experiment 1)[18], complex expected value landscapes[27], and delayed onset of payoff-information[28]. However, much of this prior work assumes that the seminal work using this task provides strong evidence for near-optimal visuomotor decision making under risk. We show that this assumption may be unfounded and that even the original paradigm (our experiment 2) elicits behavior that is better described by a heuristic strategy that is not optimally adapted to payoff and movement variability. Nevertheless, a satisficing heuristic like the one tested here might reflect utility maximization under natural information-processing constraints imposed by the human brain[12]; further analyses would be required to explore this possibility.

While most of our results support the heuristic model over the two other competing models, we are not claiming that the heuristic model is *true* in the sense that it perfectly describes how people make visuomotor decisions under risk. We observed that our participants shifted their aim-points incrementally further from the penalty region as the size of the penalty to reward ratio grew from one to five (Figure 2a,c), consistent with the optimal model and loss-averse models, but inconsistent with the heuristic model. Nevertheless, our participants consistently shifted *less* than predicted by the ideal or loss-averse models, and crucially, our participants' aim-points did not depend on their movement variability, as suggested by these two models. This latter result is consistent with prior work showing that humans do not possess accurate explicit estimates of their own movement variability[29]. Future research should explore additional models that might explain these shifts in reaching aim-point without incorporating individual estimates of movement variability. We also found that the majority of participants examined individually were described equally well by all three models. This is likely a consequence of the fact that these models make very similar (within a few pixels) predictions about a participant's behavior (Figure 2a,c).

Finally, although the models embody very different hypotheses about the processes which generate participants' behavior, their predictions for reaching end points are quite similar—within a few pixels of each other. This suggests that incentivized reaching task used here may not be well suited to test assumptions about optimality of human visuomotor decision making under risk. It would likely be more fruitful for future work to explore alternative visuomotor tasks for which an array of plausible alternative computational models make substantially distinct predictions about human behavior.

Methods

Experiment 1

Participants. 23 students (16 female) from the University of Michigan participated in experiment 1. All participants gave written consent to participate in the study. All participants were right handed, free of neurological disorders, and had normal or corrected-to-normal vision. At the end of the experiment, participants were compensated at a rate of $10/hour plus performance bonuses and their outcome from the loss aversion task described below. A participant's base hourly compensation was not affected by their performance in any task. All procedures were approved by the University of Michigan Institutional Review Board of Health Sciences and Behavioral Sciences.

Apparatus. Participants sat in an immobile chair in a small room in front of a 23" touchscreen computer monitor, and a keyboard positioned such that the spacebar was 21.5 cm away from the screen. We programmed the reaching task using PsychoPy[30] and the loss aversion task using the PyGame (www.pygame.org).

Reaching task. Participants performed an incentivized reaching task that has been argued to elicit movement strategies that are near-optimal adaptations to movement variability and payoffs[1,13–16]. At the start of each trial (Figure 2b), the participant was presented with the reward and penalty values for 1500 ms. Next, the participant looked at a white fixation cross and then pressed and held the space bar with their right index finger when they were ready to begin. Once the space bar was pressed, the fixation cross became blue, and a blue box (88 mm x 88 mm) was presented

at the center of the screen for a randomized duration of 400 to 600 ms to indicate the region where the imperative stimulus could appear. Next, overlapping green and red circles (radii = 8.5 mm) appeared at a random location within the blue bounding box and the participant had 1000 ms to reach out and touch the screen.

If a participant touched within the boundary of the green "reward" circle before the time limit (1000 ms), they earned a reward. If they touched within the boundary of the penalty circle before the time limit, they incurred a penalty. If they touched within the boundaries of both circles before the time limit, they incurred the sum of the penalty and reward. A touch outside both the penalty and reward circles resulted in neither a penalty nor a reward. If a participant completed their reach after the trial time limit had elapsed, they incurred a large penalty (-$5) for that trial and were shown feedback stating, "Too slow". If a participant released the space bar before the reward and penalty circles appeared, the trial was aborted, and they were shown feedback stating, "Too soon".

The experiment consisted of training and test sessions completed on two separate days and lasting approximately one hour each. The training session was divided into three parts. In Part 1, which consisted of 12 blocks of 36 trials, participants reached at a single unchanging yellow dot. This provided a distribution of end-points from which to estimate each participant's reaching end-point variability. Part 2 consisted of 6 blocks of 36 trials. During Part 2, participants were exposed to six conditions: 3 payoff conditions ([-$0, +$1]; [-$1, +$1]; [-$5, +$1]) with one of two separation distances between the penalty and reward regions (1R; 1.4R, where R = 8.55 mm). Participants were informed that performance in training did not affect their total earnings in the task and was just for practice. Conditions varied randomly on a trial-by-trial basis. Part 3 of training consisted of 2 additional blocks of 36 trials each in which two additional payoff conditions were added: [-$3, +$3] and [-$15, +$5]. These 'high-stakes trials' were essential to the comparison between the loss-averse model and the optimal model as the former would predict more drastic shifts away from the penalty region on these trials.

On Day 2, before the test session, participants took a "pre-test" to ensure they understood the task instructions. This pre-test presented a series of hypothetical reaching end-points for every possible payoff condition and asked the participant how much would be earned for each item. In order to advance to the test phase of the experiment, participants needed to complete the pre-test with 100% accuracy. If participants did not understand an item, we explained it to them.

Only one participant failed to complete the test with 100% accuracy on their first attempt. After the pre-test, participants performed 10 blocks of 60 trials each, with all 10 conditions varying randomly on a trial-by-trial basis. At the end of the test session, one trial was selected at random from each block and participants gained and lost bonus earnings based on their performance on these randomly selected trials. After the reaching task was complete, participants performed the loss aversion task described below.

Pre-processing. We pre-processed our raw data for ease of data analyses. This included shifting all end-points to a standard space with the center of the reward circle as the origin, removing overall horizontal biases, and removing remove temporal/spatial outlier trials. We define a spatial outlier trial as a trial with a reaching outcome that fell outside a +/-3 standard deviation range around the subject's mean. Temporal outliers were the "too slow" or "too soon" trials. Of the 10,560 of total end-points observed, we removed a total of 494 temporal outliers (late and early) (per subject removals: min. = 6, max. = 117, mean = 27.4) and 604 spatial outliers (per subject removals: min. = 10, max. = 119, mean = 33.6). We defined a participant's movement variability as the standard deviation of their horizontal reaching end-points during the first part of training, when they repeatedly reached at a single point. We assume that a participant's aim-point in a condition is equal to their mean reaching outcome in that condition.

Loss aversion task. After participants completed the last block of the test session, they completed a gambling task which measured their loss-aversion, understood as the extent to which they weighted losses more than objectively equivalent gains[3,4]. In this task, participants were presented with a randomized series of 50/50 gambles. Potential gains ranged from +$10 to +$30, while potential losses ranged from -$5 to -$15. Participants were instructed to evaluate each gamble independently and to accept or reject each of these gambles. They were further instructed to avoid the use of simple decision-making rules, such as a rule to accept all gambles with potential rewards greater than $10. Participants were told that one of the gambles would be randomly selected after completing the task and that if they had accepted the chosen gamble, the participant would win or lose the amount displayed based on the outcome of a computerized coin flip. Participants were instructed that the outcome of the gamble could add to or take away from the bonus they earned from the reaching task.

In order to calculate a measure of loss aversion from this task, logistic regression was performed with the size of the potential gain and the size of the loss as independent variables and acceptance/rejection as the dependent variable. This analysis was performed separately for each participant. Behavioral loss aversion (λ) was computed as $\lambda = \beta_{loss}/\beta_{gain}$ where both β_{loss} and β_{gain} are the unstandardized regression coefficients for the gain and loss variables separately. This analysis makes the simplifying assumption that subjective value is linearly related to objective value for both losses and gains. Eight participants (four in each experiment) had performance that resulted in an implausible λ value that was less than one (loss-seeking) or greater than 15 (extremely loss-averse). These participants appeared to be following a simple decision-making rule rather than evaluating each gamble independently.

Endpoint analysis. To test for effects of payoff condition on trial-level reaching outcomes, we used linear mixed effects models and model comparison to infer significance. These models were implemented and compared using the R package lme4[31]. The null model was specified as $x \sim spread + (1 + prr + spread\,|subject)$, where *spread* is the distance between the two circles' centers and *prr* is the ratio of penalty to reward magnitude. The alternative model was specified as $x \sim spread + prr + \ (1 + prr + spread\ |subject)$. Since these models are nested, their fits can be compared using a simple log-likelihood ratio test. If the alternative model fits significantly better than null, then we infer that the fixed effect of *prr* significantly increases the predictive power of the model.

Computational modeling analysis. All computational modeling analyses were done with R. To decide which of the three computational models (optimal, loss-averse, or heuristic) best describe our participants, we derive model predictions using simulations, assess the fit of these predictions to the observed data, and compare the models in terms of their fit. To derive predictions from the optimal and loss-averse models, we estimate solutions to equations 1 and 2 over a grid of possible aim-points for each combination of participant and task condition. To solve equation 1 or 2 for a given aim-point, participant, and condition, we start by generating a bivariate gaussian distribution of 100000 points around the aim-point using an independent estimate of the participant's movement endpoint variance-covariance matrix:

$$\mathrm{v} = \begin{pmatrix} \mathrm{var}[x] & \mathrm{cov}[x,y] \\ \mathrm{cov}[x,y] & \mathrm{var}[y] \end{pmatrix}$$

Next, we determine the earning of each point, based on where whether it fell in the virtual reward or penalty circles, and compute the expected value of the present aim-point as the mean earnings across all 100000 samples. The predicted aim-point for a given condition is the one with the maximum expected value. For the loss-averse model, sample losses are multiplied by the participant's independently estimated loss-aversion parameter, λ.

After the predictions of the models have been derived, we compute a vector of log-likelihoods of participants' reaching outcomes (x_p, y_p), conditional on the model's predictions (x_m, y_m) and the participants' variance-covariance matrices (v_p):

$$\vec{\ell}_m = \log \mathbb{P}\big((\vec{x}_p, \vec{y}_p) \mid (\vec{x}_m, \vec{y}_m), \vec{\mathrm{v}}_p\big)$$

We then compare our models using these log-likelihood vectors and Vuong's log-likelihood ratio test for non-nested models[23]. Vuong's stat is based on a vector of likelihood ratios for two models, defined as:

$$\vec{r} = \vec{\ell}_{m_1} - \vec{\ell}_{m_2}$$

Vuong's stat is then computed as follows, where *n* is the length of *r* (sample size):

$$Z = \sqrt{n} * \mathrm{mean}(\vec{r})/\mathrm{sd}(\vec{\mathrm{r}})$$

We infer that one model fits better than another if the corresponding Vuong stat has a probability under the null of $p < 0.05$ (i.e., if pnorm(Z) < 0.05). When comparing the optimal and heuristic models to the loss-averse model, participants with implausible loss-aversion parameters ($\lambda > 15$ or $\lambda < 1$) were excluded, since these extreme values would negatively skew the loss-averse model's performance. Lastly, we declared a model as the *best* description of a participant if and only if the model fits their data significantly better than the other two models.

Efficiency analyses. In line with prior work, we examine how similar participants' earnings in the test session were to the expected earnings predicted by the optimal model[13,14,16,17,24–26]. We were particularly interested in whether this 'efficiency' measure decreased as the penalty to reward ratio increased, perhaps reflecting choking under pressure. We compute a participant's efficiency by estimating the expected value of their mean end-point in each condition and dividing this expected value by the expected value of the mean end-point predicted by the optimal model. We then used mixed effects models and model comparison to test whether this efficiency score

decreased as the penalty to reward ratio increased. The null model was specified as $x \sim (1 \mid subject)$, where *prr* is the ratio of penalty to reward magnitude. The alternative model was specified as $x \sim prr + (1 \mid subject)$. Random slopes were not included in the models because there was only one observed efficiency score for each pair of subject and *prr* level. The models were compared using a log-likelihood ratio test, allowing us to assess whether adding the *prr* predictor significantly increased the model's fit.

Experiment 2

Participants. 18 students (12 female) from the University of Michigan. All participants gave written consent to participate in the study. All participants were right-handed and had normal or corrected-to-normal vision. At the end of the experiment, participants were compensated at a rate of $10/hour plus performance bonuses and their outcome from the loss aversion task described below, if applicable. A participant's base hourly compensation was not affected by their performance in any task. The University of Michigan Institutional Review Board of Health Sciences and Behavioral Sciences approved all procedures used in this experiment.

Apparatus. Participants sat in an immobile chair in a small room in front of a 23" touchscreen computer monitor, and a keyboard positioned such that the spacebar was 21.5 cm away from the screen. We programmed the reaching task using OpenSesame[32] and the loss aversion task using the PyGame (www.pygame.org)

Reaching task. Participants performed a task similar to the incentivized reaching task used in Experiment 1. All task parameters were identical except those described below. Experiment 2 consisted of training and test sessions completed on two separate days and lasting approximately one hour each. The training session was composed of 10 blocks of 30 trials each. In training, participants always saw the overlapping reward and penalty regions, but the penalty value was zero. The time-limit for each reaching movement decreased from unlimited in block 1 to 850 ms in blocks 2 – 4 to 700 ms in blocks 5 – 10. The only penalties in the training session were due to reaches completed after the time-limit. The separation distance between the reward and penalty

circles was fixed at 1R (32 pixels). We estimated participants' movement variability from their data in blocks 5 through 10 of training.

The test session was composed of 12 blocks of 30 trials each. In the test session, participants accrued points throughout the experiment that were translated into cash earnings at a rate of 25 cents per 1000 points at the conclusion of the experiment. The reward value was fixed at +100 points and the penalty value was fixed within each block at either zero, -100, or -500 points (4 blocks each, in random order). The time-limit was 700 ms and the penalty for completing a movement after the time-limit was -700 points.

Analyses. All analyses were identical to those used in experiment 1, except for obvious changes to accommodate the distinct task conditions (e.g., only one spatial condition, only three payoff conditions, etc.).

REFERENCES

1. Trommershäuser, J., Maloney, L. T. & Landy, M. S. Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences* **12,** 291–297 (2008).

2. Tverksy, A. & Kahneman, D. Judgement under uncertainty: Heuristics and Biases. *Science* **185,** 1124–1131 (1975).

3. Tversky, A. & Kahneman, D. Prospect theory: An analysis of decision under risk. *Econometrica* **47,** 263–291 (1979).

4. Tom, Fox, C., Trepel, C. & Science, P. R. The neural basis of loss aversion in decision-making under risk. (2007).

5. Chib, V. S., Martino, B., Shimojo, S. & O'Doherty, J. P. Neural mechanisms underlying paradoxical performance for monetary incentives are driven by loss aversion. *Neuron* **74,** 582–94 (2012).

6. Chib, V. S., Shimojo, S. & O'Doherty, J. P. The Effects of Incentive Framing on Performance Decrements for Large Monetary Outcomes: Behavioral and Neural Mechanisms. *The Journal of Neuroscience* **34,** 14833–14844 (2014).

7. Hutchinson, J. M. & Gigerenzer, G. Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural processes* **69,** 97–124 (2005).

8. Shah, A. K. & Oppenheimer, D. M. Heuristics made easy: an effort-reduction framework. *Psychological bulletin* **134,** 207–22 (2008).

9. Gigerenzer, G. & Gaissmaier, W. Heuristic Decision Making. *Annual Review of Psychology*

62, 451–482 (2011).

10. Simon, H. A. A behavioral model of rational choice. *The quarterly journal of economics* (1955). doi:10.2307/1884852

11. Schwartz, B., Haim, Y. & Dacso, C. What makes a good decision? Robust satisficing as a normative standard of rational decision making. *Journal for the Theory of Social Behaviour* **41,** 209–227 (2011).

12. Howes, A., Warren, P. A., Farmer, G., El-Deredy, W. & Lewis, R. L. Why contextual preference reversals maximize expected value. *Psychological Review* **123,** 368 (2016).

13. Trommershäuser, J., Maloney, L. T. & Landy, M. S. Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America. A, Optics, image science, and vision* **20,** 1419–33 (2003).

14. Trommershäuser, J., Maloney, L. T. & Landy, M. S. Statistical decision theory and trade-offs in the control of motor response. *Spatial vision* **16,** 255–75 (2003).

15. Trommershäuser, J., Gepshtein, S., Maloney, L. T., Landy, M. S. & Banks, M. S. Optimal Compensation for Changes in Task-Relevant Movement Variability. *The Journal of Neuroscience* **25,** 7169–7178 (2005).

16. Trommershäuser, J., Landy, M. S. & Maloney, L. T. Humans rapidly estimate expected gain in movement planning. *Psychological science* **17,** 981–8 (2006).

17. Neyedli, H. F. & Welsh, T. N. Optimal weighting of costs and probabilities in a risky motor decision-making task requires experience. *Journal of Experimental Psychology: Human Perception and Performance* **39,** 638 (2013).

18. Neyedli, H. F. & Welsh, T. N. People are better at maximizing expected gain in a manual

aiming task with rapidly changing probabilities than with rapidly changing payoffs. *Journal of Neurophysiology* **111,** 1016–1026 (2014).

19. Neyedli, H. F. & LeBlanc, K. A. The Role of Consistent Context in Rapid Movement Planning: Suboptimal Endpoint Adjustment to Changing Rewards. *Journal of Motor Behavior* 1–11 (2017). doi:10.1080/00222895.2016.1271296

20. Dekker, T. M. & Nardini, M. Risky visuomotor choices during rapid reaching in childhood. *Developmental Science* **19,** 427–439 (2016).

21. Ariely, D., Gneezy, U., Loewenstein, G. & Mazar, N. Large Stakes and Big Mistakes. *Review of Economic Studies* 451–469 (2009). doi:10.1111/j.1467-937X.2009.00534.x

22. Lee, T. G. & Grafton, S. T. Out of control: diminished prefrontal activity coincides with impaired motor performance due to choking under pressure. *NeuroImage* **105,** 145–55 (2015).

23. Vuong, Q. H. Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses. *Econometrica* 307 (1989). doi:10.2307/1912557

24. Wu, S.-W. W., Delgado, M. R. & Maloney, L. T. Economic decision-making compared with an equivalent motor task. *Proceedings of the National Academy of Sciences of the United States of America* **106,** 6088–93 (2009).

25. Jarvstad, A., Hahn, U., Rushton, S. K. & Warren, P. A. Perceptuo-motor, cognitive, and description-based decision-making seem equally good. *Proceedings of the National Academy of Sciences of the United States of America* **110,** 16271–6 (2013).

26. Jarvstad, A., Hahn, U., Warren, P. A. & Rushton, S. K. Are perceptuo-motor decisions really more optimal than cognitive decisions? *Cognition* **130,** 397–416 (2014).

27. Wu, S.-W. W., Trommershäuser, J., Maloney, L. T. & Landy, M. S. Limits to human

movement planning in tasks with asymmetric gain landscapes. *Journal of vision* **6,** 53–63 (2006).

28. Trommershäuser, J., Mattis, J., Maloney, L. T. & Landy, M. S. Limits to human movement planning with delayed and unpredictable onset of needed information. *Experimental Brain Research* **175,** 276–284 (2006).

29. Zhang, H., Daw, N. D. & Maloney, L. T. Testing whether humans have an accurate model of their own motor uncertainty in a speeded reaching task. *PLoS computational biology* **9,** e1003080 (2013).

30. Peirce, J. W. PsychoPy--Psychophysics software in Python. *Journal of neuroscience methods* **162,** 8–13 (2007).

31. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J Stat Softw* **67,** (2015).

32. Mathôt, S., Schreij, D. & Theeuwes, J. OpenSesame: an open-source, graphical experiment builder for the social sciences. *Behavior research methods* **44,** 314–24 (2012).