Experimental Brain Research
https://doi.org/10.1007/s00221-017-5145-7

RESEARCH ARTICLE

Sensory uncertainty impacts avoidance during spatial decisions

Kevin Jarbo[1,2] · Rory Flemming[3,4] · Timothy D. Verstynen[1,2]

Received: 7 July 2017 / Accepted: 5 December 2017

Abstract

When making risky spatial decisions, humans incorporate estimates of sensorimotor variability and costs on outcomes to bias their spatial selections away from regions that incur feedback penalties. Since selection variability depends on the reliability of sensory signals, increasing the spatial variance of targets during visually guided actions should increase the degree of this avoidance. Healthy adult participants ($N=20$) used a computer mouse to indicate their selection of the mean of a target, represented as a 2D Gaussian distribution of dots presented on a computer display. Reward feedback on each trial corresponded to the estimation error of the selection. Either increasing or decreasing the spatial variance of the dots modulated the spatial uncertainty of the target. A non-target distractor cue was presented as an adjacent distribution of dots. On a subset of trials, feedback scores were penalized with increased proximity to the distractor mean. As expected, increasing the spatial variance of the target distribution increased selection variability. More importantly, on trials where proximity to the distractor cue incurred a penalty, increasing variance of the target increased selection bias away from the distractor cue and prolonged reaction times. These results confirm predictions that increased sensory uncertainty increases avoidance during risky spatial decisions.

Keywords Visually guided action · Target estimation and selection · Spatial risk · Sensory uncertainty · Bias

Introduction

It is fabled that William Tell was forced to use an arrow to precariously shoot an apple placed atop his son's head. Successful completion of his task required Tell to optimally aim his crossbow for the high reward target, i.e., the apple, while avoiding an area with a very high penalty, i.e., his son's head (see also Trommershäuser et al. 2003a). Situations like this can be complicated by environmental conditions. For instance, a thick fog settling into the square would increase the difficulty of Tell's aiming decision. The increased noise in the target estimation process would, in turn, reduce his accuracy and impact the likelihood of striking the apple.

In scenarios like these, where people must execute a visually guided movement with a potentially high cost on feedback outcomes, humans avoid aiming at locations that increase the likelihood of a penalty (Meyer et al. 1988). This spatial avoidance relies on the statistics of both sensory (Whiteley and Sahani 2008) and motor (Trommershauser et al. 2005) signals in the goal of probabilistically estimating the degree of risk associated with actions made to different areas of space (Nagengast et al. 2011). Specifically, humans account for both penalty magnitude and response variability, such that an increase in either will increase their penalty-avoidance bias (Gepshtein et al. 2007; Landy et al. 2007, 2012; Trommershauser et al. 2005; Trommershäuser et al. 2003a, b, 2006; Wu et al. 2006). This process is largely consistent with statistical decision theory, which describes how probabilistic information is incorporated into decision processes (Berger 1985; Maloney and Zhang 2010) to maximize expected gain by minimizing penalty on decisions with uncertain costs. Specifically, the expected gain model of sensorimotor decision-making posits that humans combine probabilistic estimates of spatial targets with estimates of

✉ Timothy D. Verstynen
timothyv@andrew.cmu.edu

1 Department of Psychology, Carnegie Mellon University, Baker Hall 342C, Pittsburgh, PA 15213, USA

2 Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213, USA

3 Department of Neuroscience, University of Pittsburgh, Pittsburgh, PA 15213, USA

4 Department of Psychology, University of Pittsburgh, Pittsburgh, PA 15213, USA

relative reward and penalty associated with those targets to produce a gradient of action value (Trommershäuser et al. 2003a, 2008). Figure 1a depicts the expected gain model in an example trial, where the peak of the expected gain function is shifted away from a penalty-inducing stimulus. One cue represents the target of a spatial action, where selections closer to the mean lead to greater rewards (solid black distribution, Fig. 1a), while the other represents the spatial location of a region, where selections will incur a penalty (gray distribution, Fig. 1a). The expected gain model predicts that the magnitude of selection bias away from the penalty region is estimated as a cost function across space (x) (black dashed line, Fig. 1a), reflecting the difference between these two distributions (Gepshtein et al. 2007; Landy et al. 2007, 2012; Neyedli and Welsh 2013; Trommershäuser et al. 2003a, b, 2006; Wu et al. 2006).

$$\mathrm{MEG}_x = \arg\max\left(\alpha N(x;\mu_T,\sigma_T) - (1-\alpha)(N(x;\mu_{\mathrm{NT}},\sigma_{\mathrm{NT}}))\right). \tag{1}$$

In Eq. 1, the optimal location to reach towards, i.e., the location with the maximum expected gain (MEG_x), is the maximum of a linear function that represents the difference of the target (denoted with subscript T) and non-target (denoted with subscript NT) distributions, each with a mean (i.e., centroid) and standard deviation of μ_{T} and σ_{T}, and μ_{NT} and σ_{NT}, respectively. Thus, selection behavior is reflected as a distribution of endpoints over a series of reaches with a mean centered over the target mean. The value of α ranges from zero to one and determines the weight of the difference between target and non-target distributions in the resulting response (MEG_x). When $\alpha = 1$, the penalty-inducing non-target is ignored and the selections will focus at the mean of the target. As α decreases, the location of the non-target induces a greater avoidance bias, pushing the selection away from the mean of the target. Figure 1b depicts selection bias away from the non-target (downward on the y-axis) in arbitrary units as a function of α (x-axis) at different ratios of target to non-target variance based on Eq. 1. If the target and non-target reflect the spatial location of reward and penalty, respectively, then at smaller values of α, selection bias manifests as a shift in the selection distribution away from the penalizing non-target (and towards a location in the target that is still likely to result in reward).

By design, the expected gain model predicts that changes in stimulus variance should influence estimates of the MEG_x location. This is also illustrated in Fig. 1b, where the topmost curve shows the predictions of the expected gain model (Eq. 1) across all values of α when the variances of the target and non-target are equal. Here, bias is shown as more negative values that reflect stronger avoidance away from non-target mean. The other curves show cases, where the variance of the target is 175, 250, 325, and 400% larger than the variance of the non-target. Note that bias increases as the ratio of target to non-target variance increases, as well as with lower values of α.

The previous studies on risky spatial decisions used filled-in circles with clear boundaries to represent the target and non-target regions of the space (Meyer et al. 1988;

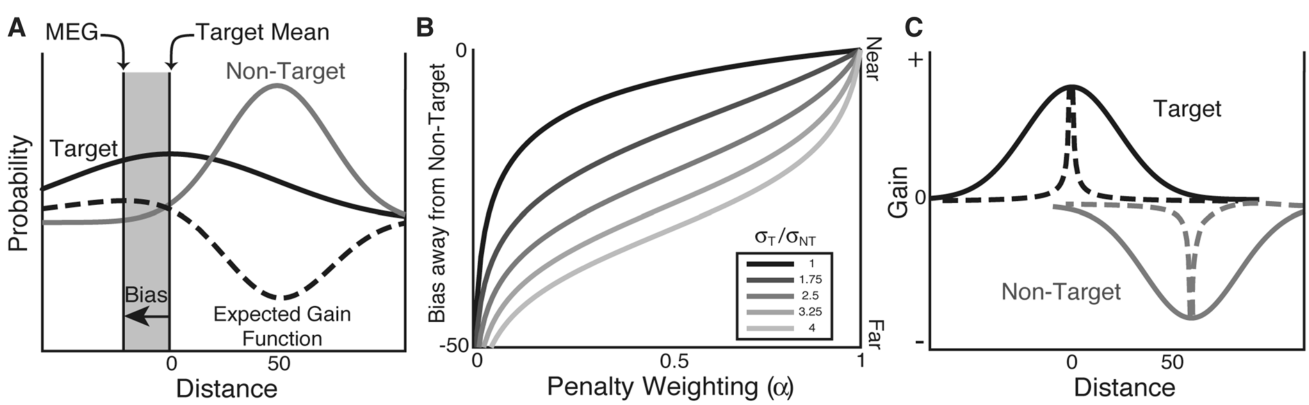

Fig. 1 Illustration of the expected gain function, predicted selection bias, and scoring functions. **a** Target (solid black line) and penalty-inducing non-target (solid gray line) stimuli are represented as Gaussian distributions with means separated by a fixed distance of 50 arbitrary units. The expected gain function (dashed black line) is approximated as a linear combination of the stimulus distributions, weighted by their relative outcomes (i.e., α and $1-\alpha$). On estimation trials (see "Experimental task"), the mean of the target must be selected to receive the highest points possible. Selecting the peak of the non-target minimizes points during penalty blocks. The peak of the expected gain function (MEG) represents the optimal perceptual position to select based on the statistics of the stimulus distributions. The gray area between the MEG and target mean demarcates where selections are biased away from the non-target across trials. **b** Selection bias away from the non-target is plotted as a function of penalty weighting (α) across different ratios of target to non-target variance from 1:1 (black curve) to 4:1 (lightest gray curve). More negative y-axis values reflect larger selection bias. **c** Dashed black and gray lines represent hyperbolic scoring functions (see Eq. 4) for the stimulus distributions. Point gain on a selection increases along the y-axis for the target, but is negative for the non-target.

Trommershäuser et al. 2008). This design introduces two limitations that we address in the current study. First, these previous efforts did not systematically manipulate the effects of stimulus variance and, therefore, sensory uncertainty on spatial decisions, since the target and non-target areas were readily visible to the participants. Put another way, these previous studies were not designed to study the spatial estimation and selection process. Here, we adapted a probabilistic stimulus design, wherein the target and non-target positions must be estimated as the respective means of two sparse Gaussian distributions of dots (Acuna et al. 2015; Bejjanki et al. 2016; Juni et al. 2015; Tassinari et al. 2006). The second limitation involves the feedback payoff structures, in which the reward and penalty values were uniform throughout target and non-target regions of the stimulus. This structure necessarily results in an optimal selection location that is always biased away from the target center and the non-target in penalty conditions (Meyer et al. 1988; Neyedli and Welsh 2013; Trommershäuser et al. 2003a, b; Wu et al. 2006). Thus, there was not a true optimal location based on both the sensory and feedback signals as expected by the expected gain model (Eq. 1), but instead a range of regions, which could be estimated purely from spatial signals, that produced the same reward. Here, we disambiguated the spatial distribution of the feedback signal from the spatial distribution of the visual stimulus to increase attention at estimating the true mean of the target. Finally, the previous studies have relied on ballistic reaches to represent their spatial selections, which can increase urgency in the action, thereby increasing selection noise. To mitigate the influence of motor noise, we allowed participants an unlimited amount of time to respond. Together, these experimental modifications allowed us to extend the past research on risky spatial decisions by better controlling variability in motor behavior.

Using this paradigm, we sought to address previously untested predictions about the effects of stimulus variance on selection behavior during risky spatial decisions. First, based on the expected gain model (Eq. 1 and Fig. 1b), we hypothesized that increasing the ratio of target (σ_{T}) to non-target (σ_{NT}) spatial variance should increase selection bias away from the non-target. In other words, when the target mean is harder to estimate, relative to the location of the non-target, participants should be more cautious in their spatial estimations and be more biased away from the non-target stimulus. Second, this effect on spatial variance on avoidance bias should interact with explicit costs (i.e., target and non-target weights, α and $1-\alpha$) to more strongly bias selections than penalty conditions alone. Finally, as demands of integrating spatial signals increases (i.e., target variance increases) and estimating relative value increases (i.e., α gets smaller), then this should increase computational demands on the decision and slow reaction times to initiate the selection. With our paradigm, we were also able to observe the influence of sensory variance on reaction times that were largely unexplored in the previous work.

Methods

Participants

Thirty undergraduate students enrolled in an introductory psychology course at Carnegie Mellon University were recruited through the university's Psychology Research Experiment System. Ten participants failed to perform at or above a criterion of 50% overall accuracy on the catch trials (see "Experimental task"), leaving a final sample of 20 participants (9 females, 11 males). Participant ages ranged from 18 to 23 years of age (mean age = 20.4) and were screened for normal or corrected-to-normal vision and right-handedness. Each eligible participant reviewed and signed a consent form approved by the Carnegie Mellon University Institutional Review Board. All participants who completed the study received credit towards fulfillment of their semester course requirements.

Experimental setup

The experiment was conducted using Psychophysics Toolbox 3.0.10 (Brainard 1997; Kleiner et al. 2007) through MATLAB (Release 2012a, The MathWorks, Inc., Natick, MA, USA) on a desktop computer running Ubuntu 14.04. Participants completed the task seated in a dimly lit room in front of a 23″ computer monitor with a total screen resolution of 1920 × 1080 pixels.

Experimental task

Using a 2 × 2 (low vs. high target variance × no penalty vs. penalty) within-subject design, each participant completed eight blocks of trials (two blocks per condition). The order of block conditions was randomized for each participant. A total of 102 trials (82 estimation trials and 20 catch trials) were presented in each block with a total of 816 trials (656 estimation trials and 160 catch trials) across the entire experiment. The experiment took approximately 45 min to complete.

All trials were self-paced and participants initiated each trial by clicking and holding the left mouse button, while the screen was blank. A fixation stimulus (+) appeared at the center of the screen following the click and hold. On estimation trials (Fig. 2a), after a uniformly sampled period of time between 250 and 2500 ms, the target and non-target stimuli were presented simultaneously on the screen for 300 ms before disappearing. Both the target and non-target

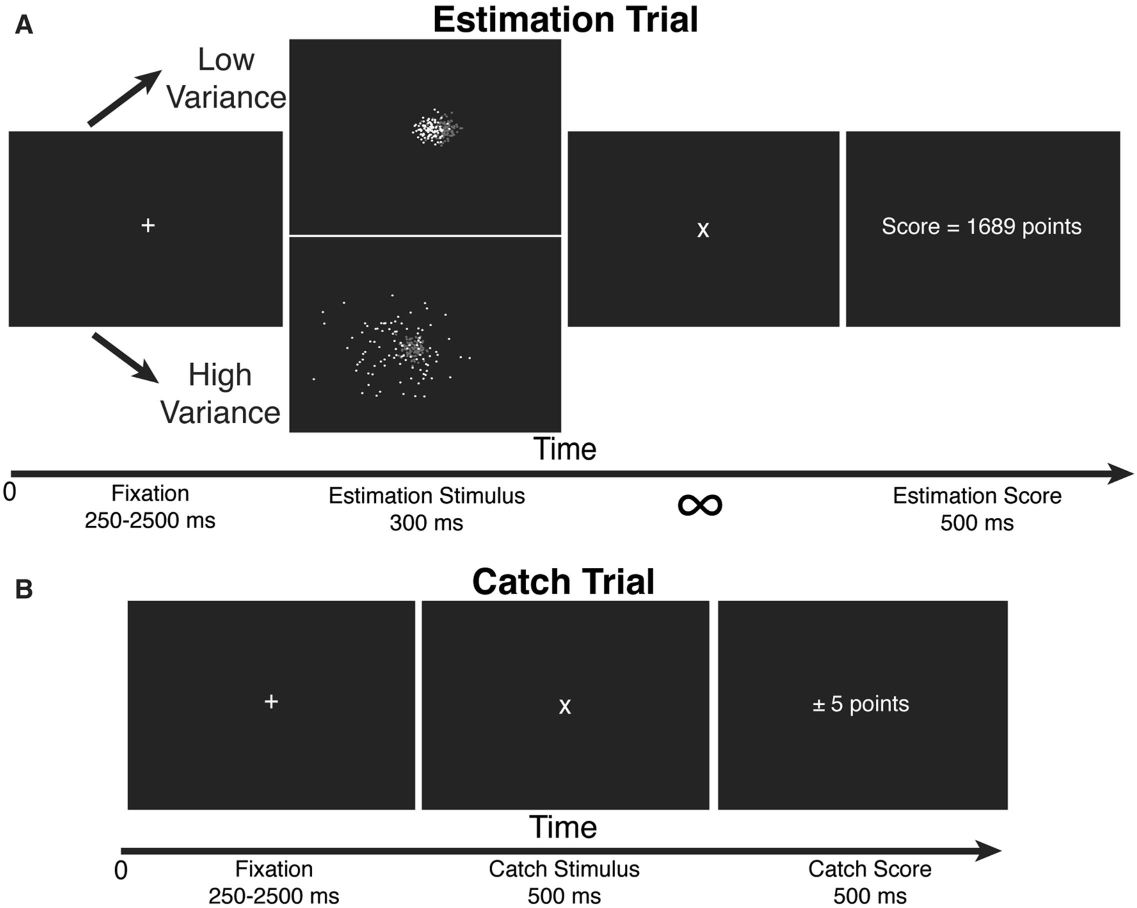

Fig. 2 Experimental trial timeline. Participants clicked and held the left mouse button to initiate all trials. **a** On estimation trials, a fixation (+) was presented (250–2500 ms jittered). The target (referred to as the Target in the task; white) and non-target (referred to as the Danger Zone in the task; gray here, presented in red in the task) stimuli flashed on screen for 300 ms then disappeared. Participants then had unlimited time to indicate their target selection by dragging the cursor (×) and releasing the mouse button. Score on a trial, based on selection distance from the target, was presented for 500 ms. **b** On catch trials, instead of visual cues, an "×" was presented after a random interval. Participants obtained a flat point total for releasing the mouse within 500 ms, or lost points for missed catch trials or responses slower than 500 ms

distributions appeared completely in randomly sampled locations within the rectangular space of 1024 × 768 pixels centered on the screen, with the constraint that both stimuli were completely visible in the workspace. The target stimulus (Target) was presented as a Gaussian distribution of 100 white dots, each 3 pixels in diameter. The non-target stimulus, referred to as the "Danger Zone" on penalty trials (described below), was simultaneously presented as a Gaussian distribution of 100 red dots (gray in Fig. 2a), also with 3-pixel diameters. This distribution could appear either to the right or to the left of the target with equal probability on each estimation trial. The horizontal distance between the mean of the target and mean of the non-target was fixed at 50 pixels. The standard deviation of the target was 25 pixels in the low variance blocks and 100 pixels in the high variance blocks. The non-target always had a standard deviation of 25 pixels.

Once the two stimuli were removed from the screen the mouse cursor was presented as an "×" at the center of the screen. Participants had an unlimited amount of time to drag the cursor to a location and then release the left mouse button to indicate their selection of the mean of the target stimulus. Immediately following the selection, a point total for that trial was presented at the center of the screen for 500 ms. The screen then went blank until the participant initiated the next trial.

On half of the blocks, the reward feedback would be penalized based on the proximity of the participant's selection to the non-target ("Danger Zone"), while this penalty was not applied on the remaining blocks. Participants were cued to the cost condition (i.e., no penalty or penalty) of the upcoming block of trials by onscreen instructions. The block commenced after the participant indicated they were ready to begin by pressing the spacebar on the keyboard.

Regardless of cost condition, selecting the mean of the target stimulus guaranteed the maximum number of points that could be scored on an estimation trial. Points scored on each estimation trial were computed based on the distance of a selection from the target and non-target means. On each trial, the Euclidean distance to the target stimulus (Eq. 2) and non-target stimulus (Eq. 3) were computed based off of the selection location (x_s, y_s) and the means of both the target stimulus (x_T, y_T) and non-target $\left(x_{NT}, y_{NT}\right)$ distributions, respectively:

$$d_T = \sqrt{\sum ((x_s,\ y_s) - (x_T,\ y_T))^2}, \tag{2}$$

$$d_{NT} = \sqrt{\sum ((x_s,\ y_s) - (x_{NT},\ y_{NT}))^2}. \tag{3}$$

The reward feedback score on each trial was computed as the weighted difference between the target (d_T) and non-target (d_{NT}) selection errors, such that

$$\text{Score} = \omega(100/d_T) - (1 - \omega)(100/d_{NT}). \tag{4}$$

In Eq. 4, the feedback score based on selection position was computed to have a hyperbolic $1/d$ falloff, where d equaled the distance between a selection location and mean of the target. Here, the scoring functions are weighted by ω, corresponding to the weight of the spatial distributions specified by the value of α in Eq. 1. In no-penalty blocks, the value of ω was set to 1, so that only selection distance from the target contributed to the score on those trials. In penalty blocks, ω was set to 0.33, so that participants incurred a heftier loss for selections that were closer to the center of the non-target.

The dashed lines in Fig. 1c provide a visual representation of the hyperbolic scoring functions overlaid with Gaussian target and non-target distributions. The highest possible score on any estimation trial was constrained to 200 if a perfect distance was estimated (i.e., $d_T = 0$). To more strongly engage participants in the task, scores were multiplied by 1000 when presented at the end of each trial. The use of the hyperbolic function meant that any spatial error between the selection and target resulted in a steep reduction in points, thereby forcing participants to aim as closely to the mean of the target stimulus as possible. The fixed distance between the target and non-target locations ensured that target selections yielded the greatest number of points on an estimation trial across all blocks, regardless of cost condition.

Twenty catch trials (Fig. 2b) were randomly presented throughout each block as an experimental control to verify that participants were fixating on the center of the screen at the start of each trial. Just like the estimation trials, participants initiated catch trials from a blank screen by clicking and holding the left mouse button. A fixation appeared at the center of the screen for a jittered period of 250–2500 ms after trial initiation. Then, in lieu of the appearance of the estimation trial stimulus, the fixation changed from a "+" to an "×" at the center of the screen. Participants then had to release the mouse button within 500 ms to gain five points, or otherwise lose five points for either failing to respond or responding too slowly.

Data from ten participants were excluded from further analyses for failure to reach 50% accuracy on the catch trials across the entire experiment. One possible explanation is that the magnitude of points on estimation trials dwarfed that on catch trials, reducing the incentive to respond. Estimation trials were also far more frequent in the task (80% of all trials) and had no limit for responses, whereas catch trial responses had to be made within 500 ms. However, we note that the general pattern of results, including the statistical significance and effect sizes of our reported results, does not change with inclusion of those data. As such, we do not include any further discussion of catch trial performance.

Data analysis

Selection variability, bias, and reaction time were the primary dependent measures. The spatial location of a selection as well as the time between offset of the stimuli and movement onset on estimation trials (i.e., reaction time) was recorded for every trial across all participants. For all analyses, only the position along the x-dimension, i.e., the selection, was used, since this was the dimension along which the adjacent non-target location was manipulated. Selection variability, bias, and reaction time were computed for all 164 estimation trials within the same experimental condition. Selection variability was computed as the standard deviation of the x-coordinate of selections across trials within a condition.

Selection bias was computed as the difference between the selection and target on a trial relative to the position of the non-target stimulus, which could be presented either to the left or to the right of the target stimulus. As illustrated in Fig. 1b, the selection bias score has more negative values with a greater selection distance away from the non-target stimulus, relative to the mean of the target stimulus. Positive values would indicate selections closer to the non-target stimulus.

A two-way repeated measures ANOVA was conducted to observe the main effects of variance and penalty (i.e., cost) conditions, as well as the variance × penalty interaction separately on selection variability, bias, and reaction time (Fig. 3). Paired sample t tests were used as post-hoc measures to determine the directionality of main effects and interactions in significant omnibus tests. Effect sizes were estimated as partial eta squared, η_p^2.

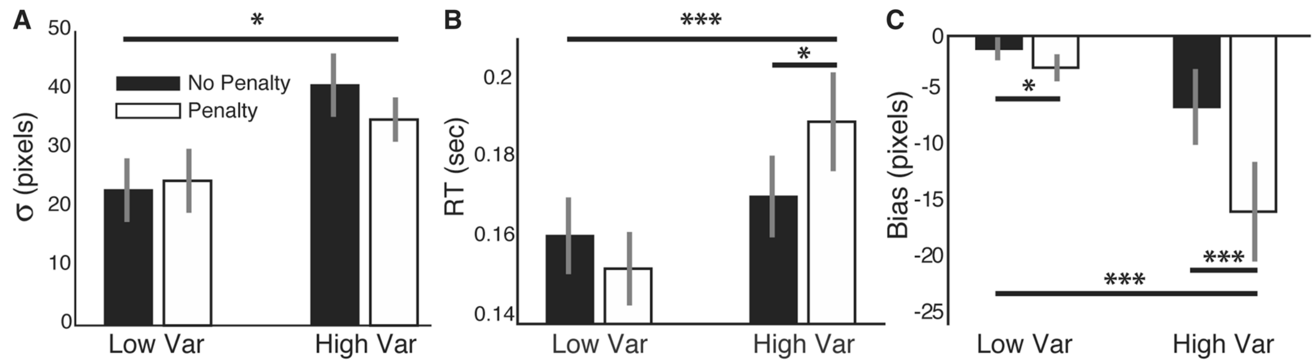

Fig. 3 Selection variability, reaction time, and bias across conditions. Bar color is the same in all panels and error bars represent the standard error of the mean. **a** Average selection variability in pixels, measured as the standard deviation (σ) of selections in low and high variance blocks, was compared across no-penalty (black) and penalty (white) blocks. There was a significant main effect of target variance that resulted in increased selection variability in high variance blocks. **b** Average reaction time (RT) in seconds, measured as the amount of time from stimulus offset and the initiation of movement on an estimation trial. There was a significant interaction between variance and penalty driven by a significant main effect of variance. RTs were slower in high variance conditions with the longest RTs in the high variance × penalty blocks. **c** Average selection bias in pixels, measured as the distance of selections from the target mean. A significant interaction between variance and penalty resulted in the greatest bias away from the non-target (Danger Zone) in the high variance × penalty blocks. Both main effects were significant and showed an increased bias in penalty blocks. Asterisks and hash lines denote significant main effects and interactions (*$p<0.05$, ***$p<0.05$)

Results

As predicted, variability of participants' estimates of the target stimulus mean was higher in the high variance blocks than the low variance blocks, $F(1,19)=7.29$, $p=0.014$, $\eta_p^2 = 0.28$ (Fig. 3a). We averaged the scaling effect of variance across the penalty and no-penalty blocks and computed the ratio of selection variability in high variance blocks to low variance blocks to be 1.59. There was no main effect of penalty condition on selection variability, $F(1,19)=0.83$, $p=0.37$, nor was there a variance × penalty interaction, $F(1,19)=1.41$, $p=0.25$. This indicates that increasing the spatial variance of the target stimulus reduces the reliability of the spatial estimations. Indeed, we did observe a significant main effect of target variance on reaction time, $F(1,19)=37.80$, $p<0.001$, $\eta_p^2 = 0.67$ (Fig. 3b), wherein reaction times slowed in the high variance conditions. Though penalty did not have a significant main effect on reaction times, $F(1,19)=2.55$, $p=0.13$, there was a significant variance × penalty interaction, $F(1,19)=5.62$, $p=0.029$, $\eta_p^2 = 0.23$. The paired t test confirmed that the high variance condition with penalty resulted in significantly slower reaction times than the high variance condition with no penalty, paired $t(19) = -2.38$, $p=0.014$, Cohen's $d = -0.53$. Reaction times were not significantly different between low variance blocks, paired $t(19)=1.68$, $p=0.1098$, regardless of penalty condition. This effect on reaction times is interesting as the past work motivating the current experiments implemented a time constraint with very short time durations (e.g., < 700 ms), which is often critical for detecting changes in reaction times due to influences on the decision process itself. However, reliable effects of spatial stimulus variance and penalty on reaction times were not observed (Neyedli and Welsh 2013, 2014; Trommershäuser et al. 2003a, b). We elaborate on this difference between studies further in the "Discussion". Taken together, the selection variance and reaction time results confirm that the target spatial variance manipulation impacted the reliability of the spatial estimation process.

The expected gain model predicts that selection bias away from the non-target stimulus should increase in these conditions of low sensory certainty, and this effect should interact with the presence of feedback penalties. Consistent with the previous observations (Neyedli and Welsh 2013, 2014; Trommershauser et al. 2005; Trommershäuser et al. 2003a, b), the introduction of a penalizing cost on selections resulted in a bias away from the non-target (Fig. 3c). The ratio of bias in penalty blocks compared to no-penalty blocks was 2.47 when we averaged the scaling effect of penalty across low and high variance conditions. Both main effects of variance, $F(1,19)=7.72$, $p=0.012$, $\eta_p^2 = 0.29$, and penalty condition, $F(1,19)=18.66$, $p<0.001$, $\eta_p^2 = 0.50$, as well as the variance × penalty interaction, $F(1,19)=10.63$, $p=0.004$, $\eta_p^2 = 0.36$, were significant. In general, selection bias was greater in high variance conditions; one-sample t tests, evaluating the bias effect with respect to a null hypothesis of zero, revealed that this nonzero bias was present in penalty blocks across the low, mean = − 2.84, $t(19) = -2.90$, $p = 0.009$, Cohen's $d = -0.65$, and high variance, mean = − 15.64, $t(19) = -3.72$, $p = 0.002$, Cohen's $d = -0.83$, conditions. A paired t test showed that the magnitude of this penalty-induced bias was significantly greater in high variance blocks, paired $t(19) = -3.23$, $p = 0.005$, Cohen's $d = -0.72$. This confirms our prediction that when the sensory reliability of spatial target estimates is low (i.e., high spatial variance), selection bias away from the non-target increases.

Discussion

Consistent with the predictions of the expected gain model (Trommershäuser et al. 2003a, 2008), we show that sensory reliability of visual targets interacts with spatial cost estimates during goal-directed action. We confirmed that increasing the spatial variance of a visual target reduces the reliability of spatial estimates of the target mean (Bejjanki et al. 2016; Körding and Wolpert 2004; Tassinari et al. 2006). By allowing an unlimited amount of time to make selections, our task was less sensitive to the effects of motor noise on spatial estimates than the previous studies that used a ballistic reaching paradigm (Neyedli and Welsh 2013, 2014; Trommershäuser et al. 2003a, b). Although because our paradigm did not pressure response speed itself, other non-planning processes could contribute to variability in the reaction times (Wong et al. 2017), tempering the interpretation of context influences on response speed. We also replicated the observation that participants biased their selections away from a non-target stimulus that could induce a penalty on feedback scores (Landy et al. 2007, 2012; Trommershauser et al. 2005; Trommershäuser et al. 2003a, b; Wu et al. 2006). Critically, we showed for the first time, under conditions of high target variance and penalty, participants most strongly biased selections away from the non-target stimulus and also took significantly more time to initiate selection movements.

Though our findings are generally consistent with probabilistic models of human spatial estimation (Landy et al. 2007; Neyedli and Welsh 2013; Tassinari et al. 2006; Trommershäuser et al. 2003a, b), we built on past research by providing support for previously unexamined predictions of the expected gain model (Trommershäuser et al. 2003a, 2008) regarding the effect of stimulus variance on spatial decisions. First, we confirmed the prediction that increasing the ratio of target to non-target stimulus variance increases avoidance bias away from the non-target stimulus (Figs. 1a, b, 3c). While this prediction comes out of the normative form of the expected gain model (see Eq. 1), it was not evaluated in the previous studies, because the stimuli used did not allow for the manipulation of spatial certainty. Our novel implementation of 2D Gaussian distributions as target and non-target stimuli, rather than circles, allowed for systematically manipulating the spatial precision of sensory signals and, consequently, the variance of the estimation process itself. Second, we found an interaction between stimulus reliability and penalty-induced avoidance bias, the greatest selection bias was away from the penalizing non-target in high target variance conditions. Again, this follows from the prediction of the normative form of the expected gain model (see Fig. 1b). Finally, by dissociating the feedback function from perceptual distributions of the target (Fig. 1c), we were able to show that the avoidance bias reflects a purely perceptual estimation process, rather than a feedback learning process. Specifically, had participants only been using trial-by-trial reward feedback signals to find the optimal selection location, then the mean of their selections would have centered on the mean of the target stimulus (i.e., zero spatial bias). The fact that participants still showed a bias in non-penalty conditions and that this bias scaled with perceptual reliability of the spatial location of the target, confirms that trial-by-trial reinforcement learning has little impact on the estimation process itself (Trommershäuser et al. 2003a, b). We should point out, however, that Neyedli and Welsh (2013) found evidence that reward feedback signals may moderate the shape of this bias over time. While we failed to see evidence of such learning in our participants (analysis not shown), this is likely due the fact that our experiment was not designed to explicitly test for learning effects.

It is worth noting that the avoidant selection behavior shown here likely reflects a top–down strategy rather than a simple computation on bottom–up sensory inputs. The past risky spatial decision-making studies manipulated the degree of target and non-target circle overlap (Meyer et al. 1988; Neyedli and Welsh 2013; Trommershäuser et al. 2003a, b), thereby constraining the spatial region that would produce a reward. From a statistical perspective, there are two ways to compensate behavior when the rewarded spatial region shrinks: improve spatial precision by reducing selection variability or improve accuracy by shifting the mean of the selection to being near the center of this constrained reward region. As in the current study, these previous experiments showed that rather than constraining motor output variability based on sensory estimates to ensure that selections fell within the available space of the target, participants shifted the mean of their selections away from the penalizing non-target to a degree that corresponded to target and non-target overlap to avoid losses. In some ways, this is consistent with the principle of loss aversion in Prospect Theory, which posits that "losses loom larger than gains" in that individuals are more sensitive to a potential loss than a sure gain of equal or greater expected value (Tversky and Kahneman 1992). However, it is worth noting that our experimental design was able to dissociate the actual optimal feedback position with the expected optimal position given a weighted combination of the two stimuli (see Fig. 1c). If participants were simply using a mixture of the incoming sensory signals and previous reward feedback to learn an optimal location to select, they would always select the mean of the target stimulus. The fact that we also observed a strong avoidance bias suggests that the maximum expected gain estimation is a purely perceptually driven spatial estimate and not an optimal decision given the reward feedback delivered.

Although selection bias was not significantly different from zero on low target variance blocks with no penalty, selections still trended away from the non-target, rather than varied symmetrically about the target (i.e., zero bias). While participants may have been primed to always avoid the non-target stimulus across the experiment, this nonzero bias on no-penalty blocks when α was fixed at a value of 1 suggests that participants did not fully discount the non-target, even though it should have had no influence on their spatial estimates. Our experimental design was limited in determining whether this observation was due to some carryover of α values when no-penalty blocks followed penalty blocks, or whether there is some other source of noise in the spatial estimation process that should be considered in the expected gain model. Indeed, there is evidence that both learning of expected costs on spatial decision outcomes (Neyedli and Welsh 2013, 2014) and noisy spatial estimates (Juni et al. 2015) can lead to biased selection behavior. Future paradigms can use a counterbalanced block structure and manipulate α parametrically to quantitatively assess any potential effects of carryover (e.g., learned value of α) or noisy spatial estimates.

We also found that increased sensory stimulus variance interacts with penalty to further slow the time it took to initiate the selection decision. Under conditions of high target variance, participants took significantly longer to initiate their movements with the slowest reaction times occurring in high variance blocks with penalty. When penalty was added along with high demands on selection precision, i.e., high sensory variance, participants might have taken longer to respond as a precaution to mitigate larger than expected costs. This cautionary behavior may reflect some subjective difficulty or uncertainty either in a strategic plan to reduce known costs or in the implicit spatial estimation process itself. This presents another avenue of research wherein expected costs, reflected by trial-by-trial fluctuations in the value of α, and estimates of stimulus variance can be considered together during spatial decisions. As such, new models of sensorimotor integration can relate explicit (e.g., costs) and implicit (e.g., sensory variance) aspects of estimation processes that underlie spatial decision behavior (McDougle et al. 2016; Summerfield and Tsetsos 2012; Taylor et al. 2014).

Taken together, our findings clearly show that estimates of sensory variance contribute to the degree to which individuals attempt to avoid penalties during risky spatial decisions by biasing their action selections away from regions that induce feedback penalties. Based on our results, it largely appears that estimates of stimulus variance and cost conditions along with expected feedback are considered together while people make spatial judgments in an attempt to maximize gain. However, we note that some limitations in our experimental design prohibited us from quantitatively determining how much increased selection bias and slower reaction times can be attributed to noisy, or less reliable, estimates of costs or stimulus variability. To address remaining questions in future work, modifications to the present paradigm and the development of quantitative models can more deeply explore the interaction of explicit and implicit processes that support spatial decision-making behavior in the context of risk.

Acknowledgements The authors would like to thank Drs. Roberta Klatzky, Marlene Behrmann, Algernop Krieger, and Robert Whitwell for their consultation and helpful comments on the manuscript.

Compliance with ethical standards

Funding Kevin Jarbo was funded by the National Institutes of Health predoctoral training grant 5T32GM081760-08 during all stages of research and manuscript submission. This research was supported in part by the PA Department of Health Formula Award SAP4100062201.

Conflict of interest The authors have no conflicts of interest to report.

References

Acuna DE, Berniker M, Fernandes HL, Kording KP (2015) Using psychophysics to ask if the brain samples or maximizes. J Vis 15(3):7. https://doi.org/10.1167/15.3.7

Bejjanki VR, Knill DC, Aslin RN (2016). Learning and inference using complex generative models in a spatial localization task. 16:1–13. https://doi.org/10.1167/16.5.9.doi

Berger JO (1985) Statistical Decision Theory and Bayesian Analysis. Springer, New York. https://doi.org/10.1007/978-1-4757-4286-2

Brainard DH (1997) The psychophysics toolbox. Spat Vis 10(4):433–436. https://doi.org/10.1163/156856897X00357

Gepshtein S, Seydell A, Trommershäuser J (2007) Optimality of human movement under natural variations of visual-motor uncertainty. J Vis 7(5):13.1–18. https://doi.org/10.1167/7.5.13

Juni MZ, Gureckis TM, Maloney LT (2015) Information sampling behavior with explicit sampling costs. Decision 3(1):41. https://doi.org/10.1037/dec0000045

Kleiner M, Brainard DH, Pelli DG, Broussard C, Wolf T, Niehorster D (2007). What's new in Psychtoolbox-3? Perception 36:S14. https://doi.org/10.1068/v070821

Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427(6971):244–247. https://doi.org/10.1038/nature02169

Landy MS, Goutcher R, Trommershäuser J, Mamassian P (2007) Visual estimation under risk. J Vis 7(6):4. https://doi.org/10.1167/7.6.4

Landy MS, Trommershäuser J, Daw ND, Trommershauser J, Daw ND (2012) Dynamic estimation of task-relevant variance in movement under risk. J Neurosci 32(37):12702–12711. https://doi.org/10.1523/JNEUROSCI.6160-11.2012

Maloney LT, Zhang H (2010) Decision-theoretic models of visual perception and action. Vis Res 50(23):2362–2374. https://doi.org/10.1016/j.visres.2010.09.031

McDougle SD, Ivry RB, Taylor JA (2016) Taking aim at the cognitive side of learning in sensorimotor adaptation tasks. Trends Cognit Sci 20(7):535–544. https://doi.org/10.1016/j.tics.2016.05.002

Meyer DE, Abrams R, Kornblum S, Wright CE, Smith JE (1988) Optimality in human motor performance: ideal control of rapid aimed movements. Psychol Rev 95(3):340–370. https://doi.org/10.1037/0033-295X.95.3.340

Nagengast AJ, Braun D, Wolpert DM (2011) Risk-sensitivity and the mean-variance trade-off: decision making in sensorimotor control. Proc Biol Sci R Soc 278(1716):2325–2332. https://doi.org/10.1098/rspb.2010.2518

Neyedli HF, Welsh TN (2013). Optimal weighting of costs and probabilities in a risky motor decision-making task requires experience. J Exp Psychol Hum Percept Perform 39(3):638–645. https://doi.org/10.1037/a0030518

Neyedli HHF, Welsh TNT (2014) People are better at maximizing expected gain in a manual aiming task with rapidly changing probabilities than with rapidly changing payoffs. J Neurophysiol 111(5):1016–1026. https://doi.org/10.1152/jn.00163.2013

Summerfield C, Tsetsos K (2012) Building bridges between perceptual and economic decision-making: Neural and computational mechanisms. Front Neurosci 6(MAY):1–20. https://doi.org/10.3389/fnins.2012.00070

Tassinari H, Hudson TE, Landy MS (2006) Combining priors and noisy visual cues in a rapid pointing task. J Neurosci 26(40):10154–10163. https://doi.org/10.1523/JNEUROSCI.2779-06.2006

Taylor J, Krakauer JW, Ivry RB (2014) Explicit and implicit contributions to learning in a sensorimotor adaptation task. J Neurosci 34(8):3023–3032. https://doi.org/10.1523/JNEUROSCI.3619-13.2014

Trommershauser J, Trommershäuser J, Gepshtein S, Maloney LT, Landy MS, Banks MS, Trommershauser J (2005) Optimal compensation for changes in task-relevant movement variability. J Neurosci 25(31):7169–7178. https://doi.org/10.1523/JNEUROSCI.1906-05.2005

Trommershäuser J, Maloney LT, Landy MS (2003a) Statistical decision theory and the selection of rapid, goal-directed movements. J Opt Soc Am A Opt Image Sci Vis 20(7):1419–1433. https://doi.org/10.1364/JOSAA.20.001419

Trommershäuser J, Maloney LT, Landy MS (2003b) Statistical decision theory and trade-offs in the control of motor response. Spat Vis 16(3–4):255–275. https://doi.org/10.1163/156856803322467527

Trommershäuser J, Landy MS, Maloney LT (2006) Humans rapidly estimate expected gain in movement planning. Psychol Sci 17(11):981–988. https://doi.org/10.1111/j.1467-9280.2006.01816.x

Trommershäuser J, Maloney LT, Landy MS (2008) Decision making, movement planning and statistical decision theory. Trends Cognit Sci 12(8):291–297. https://doi.org/10.1016/j.tics.2008.04.010

Tversky A, Kahneman D (1992) Advances in prospect-theory—cumulative representation of uncertainty. J Risk Uncertain 5(4):297–323. https://doi.org/10.1007/Bf00122574

Whiteley L, Sahani M (2008) Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. J Vis 8(3):2.1–15. https://doi.org/10.1167/8.3.2

Wong AL, Goldsmith J, Forrence AD, Haith AM, Krakauer JW (2017). Reaction times can reflect habits rather than computations. eLife 6:e28075. https://doi.org/10.7554/eLife.28075

Wu S-W, Trommershäuser J, Maloney LT, Landy MS (2006) Limits to human movement planning in tasks with asymmetric gain landscapes. J Vis 6(1):53–63. https://doi.org/10.1167/6.1.5