Research Article

Humans Rapidly Estimate Expected Gain in Movement Planning

Julia Trommershäuser,[1] Michael S. Landy,[2] and Laurence T. Maloney[2]

[1]Department of Psychology, Giessen University, Giessen, Germany, and [2]Department of Psychology and Center for Neural Science, New York University

ABSTRACT—***We studied human movement planning in tasks in which subjects selected one of two goals that differed in expected gain. Each goal configuration consisted of a target circle and a partially overlapping penalty circle. Rapid hits into the target region led to a monetary bonus; accidental hits into the penalty region incurred a penalty. The outcomes assigned to target and penalty regions and the spatial arrangement of those regions were varied. Subjects preferred configurations with higher expected gain whether selection involved a rapid pointing movement or a choice by key press. Movements executed to select one of two goal configurations exhibited the same movement dynamics as pointing movements directed at a single configuration, and were executed with the same high efficiency. Our results suggest that humans choose near-optimal strategies when planning their movement, and can base their selection of strategy on a rapid judgment about the expected gain associated with possible movement goals.***

In the course of a day, people make many decisions. They occasionally make the kinds of explicit economic decisions studied in the decision-making literature, but more frequently they decide how to move in response to the risks and rewards in the environment. Survival can depend on making the latter kind of decision rapidly and well. In this article, we discuss human performance in executing visuo-motor tasks equivalent to decision making under risk, and examine the criteria human decision makers use in rapidly choosing between alternative courses of action.

Previous studies demonstrated that when executing speeded arm movements under risk, humans select movement strategies that are nearly optimal (Trommershäuser, Maloney, & Landy, 2003a, 2003b). In these studies, subjects pointed rapidly at stimulus configurations consisting of a small target and a partially overlapping penalty region. Reaches terminating within the target region yielded monetary reward; those ending in the penalty region could result in a loss. The target's size and the distance between target and penalty regions were small (less than 2 cm), similar in size to the subjects' end-point variability. Yet subjects performed these tasks with high efficiency. This is surprising, as the underlying decision task is complex.

To see why, consider such an experiment when hits on the target and penalty yield gains of +100 and −500 points, respectively (Fig. 1a).[1] Under these conditions, a movement can end in one of four regions: penalty only (Region R_1, gain $G_1 = -500$), target-penalty overlap (Region R_2, gain $G_2 = -400$), target only (Region R_3, gain $G_3 = 100$), or neither target nor penalty (i.e., background; Region R_4, gain $G_4 = 0$). In executing this task, the subject chooses a strategy S. The outcome of executing the strategy is an end point on the display (x, y), and the reward or penalty depends only on the region in which the end point falls. We identify a strategy S with the mean $(\bar{x}, \bar{y})$ of all the end points that would result if it were repeated many times (Trommershäuser et al., 2003b). The choice of strategy fixes the probability $P(R_i|S)$ of hitting each of the four regions R_i ($i = 1, \ldots, 4$). In the decision-making literature, the combination of event probabilities $P(R_i|S)$ and associated gains G_i is called a lottery. We denote lotteries induced by strategies as

Address correspondence to Julia Trommershäuser, Department of Psychology, Otto-Behaghel-Strasse 10F, 35394 Giessen, Germany, e-mail: julia.trommershaeuser@psychol.uni-giessen.de.

[1]Decision making spans many fields, and terminology differs from field to field. We use the terminology employed in Maloney (2002), and we refer to outcomes as gains (denoted G_i), with losses coded as negative gains, as in our previous publications. The term *expected gain* corresponds exactly to *expected value* in the psychological literature.

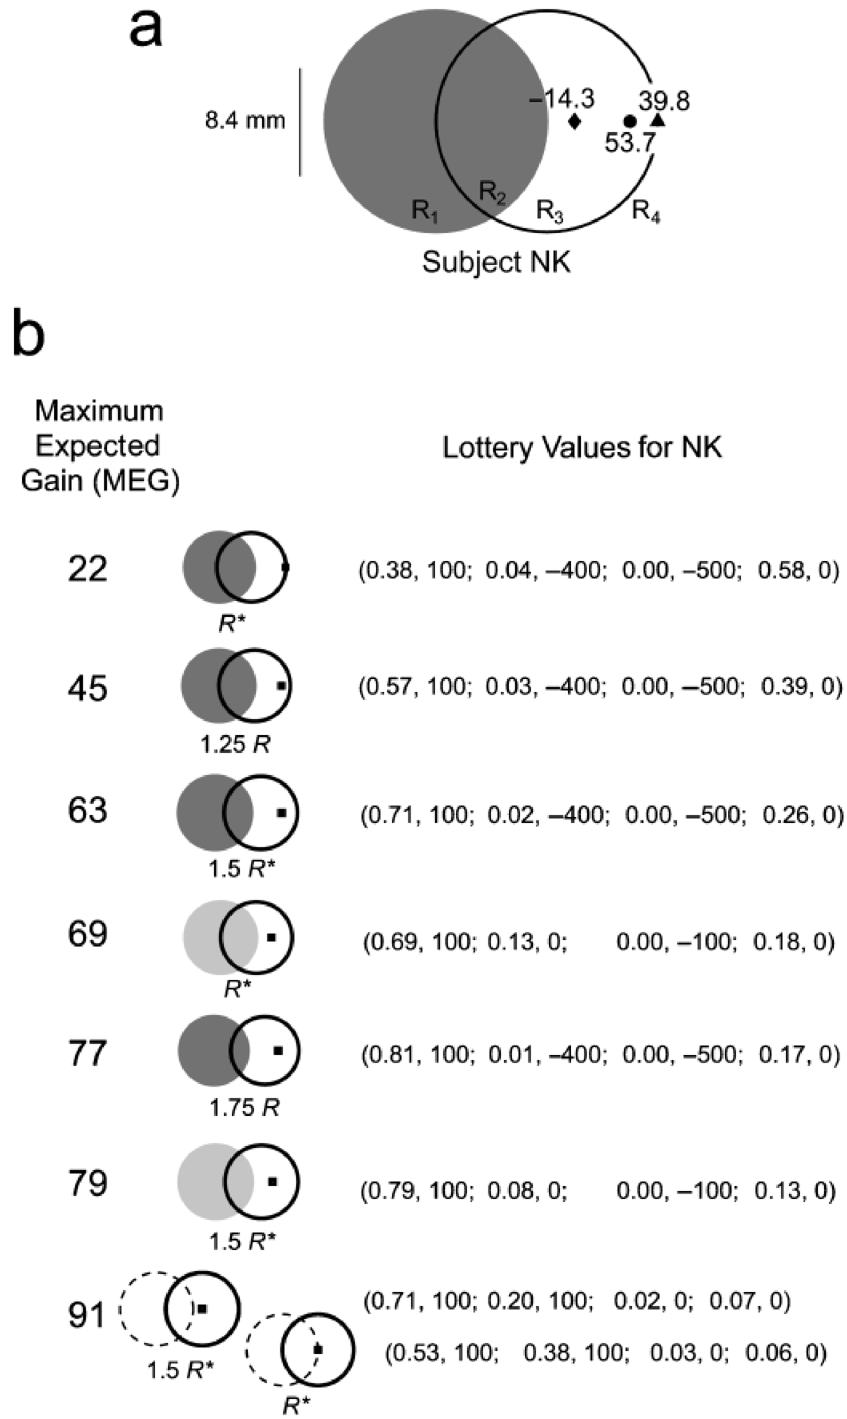

Fig. 1. Equivalence of the rapid pointing task and choice among lotteries. The illustration in (a) shows an example stimulus configuration with the mean end points for three of the many possible visuo-motor strategies. The subject can touch within a shaded penalty region (R_1), an unshaded circular reward region (the target; R_3), both (R_2), or neither (R_4). Each region has an associated gain (penalty). Next to each of the three end points is the expected gain, assuming the penalty for touching the shaded region is −500 points, the reward for touching the target region is 100 points, and the subject's variability matches that of subject N.K. (σ = 3.81 mm). The illustrations in (b) show the stimulus configurations used in Sessions 2 through 5 (selection and key-press movements). The green target (shown here as a solid black outline) and penalty areas were circular, with a radius of 28 pixels (8.4 mm). Penalty values were color coded as follows: gray (shown here as a dashed outline) = penalty of 0 points; blue (shown here as light shading) = penalty of −100 points; and red (shown here as dark shading) = penalty of −500 points. On each trial in these sessions, subjects selected one of two configurations chosen randomly from the eight shown here (or the spatially symmetric configurations with the target on the left side of the penalty region). The displacement of the target relative to the penalty region ranged from R (the radius of the circles) to 1.75 R, as shown below the diagrams. For each configuration, the black square indicates the mean end point that would maximize expected gain for subject N.K. The maximum expected gain (left) and the corresponding lottery (right; see Equation 1 in the text for an explanation) for subject N.K. are also shown. Configurations marked by an asterisk were used in Session 1 (pointing movements directed at a single configuration).

follows:

$$L(S) = (P(R_1|S),\ G_1;\ P(R_2|S),\ G_2;\ P(R_3|S),\ G_3;\ P(R_4|S),\ G_4). \tag{1}$$

An alternative movement strategy, S', corresponds to a second lottery:

$$L(S') = (P(R_1|S'), G_1; P(R_2|S'), G_2; P(R_3|S'), G_3; P(R_4|S'), G_4). \tag{2}$$

As illustrated in Figure 1a, every mean end point corresponds to a lottery with an expected gain, that is, the number of points a subject is expected to earn, on average, having adopted a strategy with that mean end point. Figure 1a indicates the expected gain of several mean end points, based on the measured end-point variability of subject N.K. from the experiment we report here. The expected gain of movements aimed at the location marked by the black triangle is less than that corresponding to movements aimed at the location marked by the black circle. However, there are many other possible mean end points and corresponding lotteries, each with its associated expected gain. By choosing among possible strategies, subject N.K., in effect, selected among possible lotteries.

The results of our previous experiments indicate that subjects choose strategies maximizing, or nearly maximizing, expected gain. Efficiency was defined as the number of points won relative to the number expected if an optimal strategy was used. Subjects' efficiencies were typically above 90% (Trommershäuser et al., 2003a, 2003b). This level of efficiency is higher than that for target selection in visual search (Eckstein, Beutter, & Stone, 2001; Najemnik & Geisler, 2005). In traditional decision-making tasks, subjects choosing between lotteries often fail to maximize expected gain (Bell, Raiffa, & Tversky, 1988; Kahneman, Slovic, & Tversky, 1982; Kahneman & Tversky, 2000).

In the study reported here, we tested whether subjects select their motor strategy on the basis of an estimate of expected gain. We forced subjects to rapidly choose one of two configurations differing in expected gain, thereby requiring the subjects to make quick estimates of the expected gains of the two configurations. A configuration's expected gain is a complex function of motor noise, the payoff for each region, and the spatial arrangement of regions (Fig. 1b; see also the section titled "Model of Optimal Movement Planning and Choice Between Lotteries"). We investigated whether subjects would be able to rapidly estimate and compare the expected gains of two configurations.

METHOD

Apparatus

The apparatus was similar to that described in our previous studies (Trommershäuser et al., 2003a, 2003b). Subjects were seated in front of a monitor equipped with a touch screen (AccuTouch, Elo Touchsystems, Menlo Park, CA) and instructed to touch objects on the screen within a predefined time window. In some sessions, subjects pressed a keyboard key instead of touching the screen. The space bar of the keyboard constituted the start position, which was located 27 cm in front of and 34 cm below the center of the screen. The viewing distance was 52 cm. The experiment was run using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) on a Dell Optiplex computer (270GX). A calibration procedure was performed before each session to ensure that the touch-screen measurements were aligned with the visual stimuli.

Stimuli and Experimental Design

The experiment comprised six 50-min sessions, run on 6 separate days. In the first session, subjects learned to respond before the time limit. This training session was followed by a *pointing-movement* session (referred to as Session 1), in which subjects pointed at single configurations. In the remaining four sessions, subjects chose one of two displayed stimulus configurations, either by rapidly pointing at it (Sessions 2 and 3, *selection movements*) or by pressing a button (Sessions 4 and 5, *key-press movements*). The session in which subjects pointed at a single configuration (Session 1) was run before the four sessions in which subjects selected between two configuration (Sessions 2–5) to make sure subjects assigned the correct penalty value to each color (see the next paragraph). We refer to movements directed at one of two configurations as selection movements (in contrast to pointing movements directed at a single configuration).

Each stimulus configuration consisted of target and penalty regions. The penalty region was circular and colored gray (penalty = 0 points), blue (penalty = −100 points), or red (penalty = −500 points; Fig. 1b). The target region (reward = 100 points) was also circular; it was marked by a green edge and was unshaded so that the overlap with the penalty circle would be readily visible. Target and penalty regions had radii R of 28 pixels (8.4 mm). The target region was displaced horizontally from the penalty region, either to the left or right, with the direction chosen randomly on each trial. The set of configurations used for selection tasks (Sessions 2–5) included the six configurations used in Session 1 and two additional configurations with a penalty value of −500 points and displacements different from those used in Session 1 (Fig. 1b).

In the training session, there were two possible displacements of the penalty circle away from the target circle (*far* displacement = $2R$, *middle* displacement = $1.5R$) and two penalty values (0 and −100 points). The session started with 30 warm-up trials with only a single green target and a time limit of 1.5 s, followed by a block of 36 trials with only a single green target and a shorter time limit of 950 ms. The next two blocks of 36 trials contained stimulus configurations with both target and penalty circles. The penalty was 0 points in the first block and −100 points in the second. Each block contained nine repetitions of each of the four spatial configurations (middle vs. far displacement, penalty displaced left vs. right), in randomized order. In the remaining six blocks of the learning session (36 trials each), the penalty values alternated between 0 and −100 points, and the time limit was reduced to 700 ms, the time limit imposed from then on.

The pointing-movement session (Session 1) started with 12 warm-up trials in which only the target circle was displayed, followed by 12 blocks of 32 trials. Each of these blocks consisted of eight repetitions of each of four stimulus configurations (leftward or rightward displacement of the penalty circle by R or $1.5R$, called *near* and *middle* displacement, respectively). Successive blocks had penalty values of 0, −100, −500, 0, −100, −500, and so on for a total of 12 blocks.

The two following selection-movement sessions (Sessions 2 and 3) each consisted of 342 trials (330 trials preceded by 12 warm-up trials). The 660 trials were 20 repetitions (completely randomized design) of each possible pair of stimulus configurations (chosen from the eight configurations in Fig. 1b), excluding pairs in which both configurations had penalty values of 0 points, but including pairs of identical configurations, to control for possible lateral judgment biases. Thus, there were 33 possible configuration pairs. On each trial, one of the configurations, chosen randomly, was displayed to the left of the screen's center, and the other one to the right; the amount of shift was in the range from 15 to 35 pixels (4.5–10.5 mm) and was chosen randomly and independently for the two configurations. Similarly, one configuration was shifted upward and the other downward (chosen randomly), with the amounts of shift chosen randomly (range: 7–27 pixels, or 2.1–8.1 mm). The displacements of the two configurations were in opposite directions to avoid accidental hits into the "wrong" configuration. The time limit for response was the same as in Session 1 (700 ms).

The final two sessions, the key-press-movement sessions (Sessions 4 and 5), consisted of the same number and sequence of trials as Sessions 2 and 3, to control for possible trial-by-trial effects. The time limit for response was still 700 ms, even though subjects needed less time to complete these button-press responses.

Procedure

In Sessions 1, 2, and 3, a white fixation cross indicated the start of each trial. The subject was required to depress the space bar of the keyboard with the same finger that he or she would use later to touch the screen. Next, the color of the fixation cross changed to blue. After a random delay of between 400 and 600 ms, the stimulus configuration or pair of configurations appeared. The subject viewed the stimulus configuration for 400 ms,[2] after which a 1000-Hz tone indicated that the subject should start the movement. A trial was aborted if the subject released the space bar earlier than 100 ms after presentation of the tone. The subject was required to touch the screen within 700 ms after the tone to avoid a time-out penalty of 700 points. If the subject touched the screen within the area of the target or penalty region, the region that was hit "exploded" graphically, and the subject received the points associated with that region. Then, the subject received feedback about the points scored in that trial, followed by feedback on the total accumulated points for that session.

In key-press-movement trials (Sessions 4 and 5), the displays were identical, and subjects were again instructed to select one of two configurations (as in Sessions 2 and 3), but in this case they made their choice by pressing one of two keys. They were told that rewards and penalties would be based on performance in Sessions 2 and 3. Once a subject made a choice, the points for that trial were awarded by simulating that subject's movements (on the basis of the subject's performance in Sessions 2 and 3). The simulated end point was sampled from a bivariate Gaussian distribution whose mean and variance matched those estimated for the corresponding configuration from Sessions 2 and 3 (end-point variability ranged from 3.3 to 3.8 mm across subjects). If the simulated end point fell within the area of the target or penalty region, the corresponding region "exploded" graphically, and the points associated with that region were added to or subtracted from the subject's winnings. Then, the subject received feedback exactly as in Sessions 2 and 3.

Subjects and Instructions

Six subjects participated in the experiment. The subjects were 4 male and 2 female students at the University of Giessen, Germany, and ranged in age from 22 to 29. All but 1 were right-handed, and all used their dominant hand to perform the experiment. All subjects had normal or corrected-to-normal vision. They had given their informed consent before testing and were paid for their participation. All were unaware of the hypothesis being tested. Subjects were informed of the payoffs and penalties before each block of trials. They were told that the total score over the six sessions would be converted into a bonus payment of 10¢ per 1,000 points; the purpose of this payment was to motivate fast, accurate responses.

Model of Optimal Movement Planning and Choice Between Lotteries

In previous work, we developed a model of optimal movement planning based on statistical decision theory (Trommershäuser et al., 2003a, 2003b). Here, we briefly summarize the ideas behind our model and explain how it applies to the choice between lotteries.

Our model is based on the finding that motor responses, particularly if executed under a tight time constraint, are variable (Fitts & Petersen, 1964; Meyer, Abrahams, Kornblum, Wright,

[2]The interval of 400 ms matches the time required for a representation of expected value to build up in the lateral intraparietal area (LIP) in monkeys, as measured through single-unit recording (Platt & Glimcher, 1999; Roitman & Shadlen, 2002; Sugrue, Corrado, & Newsome, 2004).

& Smith, 1988). We identify a visuo-motor strategy S with the mean end point $(\bar{x}, \bar{y})$ on the screen. In our experiments, movement end points are distributed around this mean end point according to a bivariate Gaussian distribution,

$$p(x, y|\bar{x}, \bar{y}, \sigma_x, \sigma_y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp\left[-(x-\bar{x})^2/2\sigma_x^2\right] \cdot \exp\left[-(y-\bar{y})^2/2\sigma_y^2\right]. \quad (3)$$

Furthermore, once subjects are practiced in the task, the variance is isotropic (i.e., $\sigma = \sigma_x = \sigma_y$) and constant throughout the experiment, independent of the mean end point.

Consider a configuration for which the penalty value is -500. A movement can end in one of four regions: penalty only (Region R_1, gain $G_1 = -500$), target-penalty overlap (Region R_2, gain $G_2 = -400$), target only (Region R_3, gain $G_3 = 100$), or neither target nor penalty (i.e., background; Region R_4, gain $G_4 = 0$). The probability of hitting inside region R_i is defined by

$$P(R_i|\bar{x}, \bar{y}, \sigma) = \int_{R_i} p(x, y|\bar{x}, \bar{y}, \sigma)\, dx\, dy \quad (4)$$

Thus, the choice of $(\bar{x}, \bar{y})$ simultaneously fixes the probability $P(R_i|S)$ of hitting each of the four regions R_i ($i = 1, \ldots, 4$). Hence, each choice of mean end point $(\bar{x}, \bar{y})$ on the screen corresponds to a lottery:

$$L(\bar{x}, \bar{y}, \sigma) = (P(R_1|\bar{x}, \bar{y}, \sigma), G_1;\ P(R_2|\bar{x}, \bar{y}, \sigma), G_2;\ P(R_3|\bar{x}, \bar{y}, \sigma), G_3;\ P(R_4|\bar{x}, \bar{y}, \sigma), G_4). \quad (5)$$

In selecting among infinitely many possible mean end points on the screen, the subject in effect selects among an infinite number of lotteries. For the stimulus configurations of our experiment, there is a single mean end point $(\bar{x}_{opt}, \bar{y}_{opt})$ corresponding to the lottery $L(\bar{x}_{opt}, \bar{y}_{opt}, \sigma)$ with maximum expected gain (see Trommershäuser et al., 2003a, for typical examples of distributions of expected gain). Our previous studies indicate that subjects choose mean end points corresponding to lotteries that nearly maximize expected gain.

In the present study, we asked whether subjects make choices based on a representation of expected gain. Choosing between configurations involves two steps. First, the subject must judge which configuration is more "promising." Second, the subject has to point at the chosen configuration.

If subjects consistently select the configuration with higher maximum expected gain (MEG), we can conclude that they effectively have an internal ordering of the configurations based on an estimate of each configuration's expected gain. We compared the subjects' preference for each configuration with that predicted by MEG. This prediction was different for each subject and was computed using each subject's end-point variability σ. Figure 1b shows the ordering of the configurations by MEG, the corresponding optimal mean end points $(\bar{x}_{opt}, \bar{y}_{opt})$, and the lotteries corresponding to MEG for 1 of the subjects (N.K., with $\sigma = 3.81$ mm).

Data Analysis

In Sessions 1 through 3, for each trial, we recorded the *reaction time* (time between the tone and release of the space bar), *movement time* (time from movement onset until the screen was hit), *screen position* that was hit, and *score*. Trials in which the subject initiated the movement less than 100 ms after presentation of the start signal or hit the screen later than 700 ms after presentation of the start signal were excluded from analysis. Data points that were more than 2 cm from any target center were classified as errors (e.g., knuckle hits) and were excluded from analysis. Each subject contributed approximately 384 data points in Session 1, approximately 660 data points for Sessions 2 and 3 combined, and 660 data points for Sessions 4 and 5 combined (20 repetitions of each pair of configurations). Movement end points were recorded relative to the center of the target circle. The end points for symmetric configurations (i.e., configurations that differed only in whether the target was displaced to the left or right of the penalty circle) did not differ significantly, so the data were collapsed across symmetric conditions.

Efficiency for Different Types of Judgments

To examine whether efficiency in the motor task differed for pointing movements compared with selection movements, we computed efficiency for performance in Session 1 and in Sessions 2 and 3 combined. We define efficiency in our task as the actual score divided by the optimal score derived from the MEG movement-planning model. The actual score was computed for each subject individually, summed over all conditions (configurations or pairs of configurations). Performance was classified as significantly different from optimal when the actual score fell outside the 95% confidence interval of optimal performance (see Trommershäuser et al., 2003a, for a discussion of how to compute this range of efficiencies).

Preference for Configurations With Higher MEG

We computed the MEG (based on each subject's variability and the optimal strategy) for each configuration used in Sessions 2 through 5. We emphasize that the subjects were choosing between two configurations, each of which had an MEG that depended on the configuration and the subject's motor uncertainty. To test whether subjects preferred the configuration with higher MEG, for each subject, we computed the proportion of trials, pooled over all configuration pairs, in which the configuration with higher MEG was chosen and tested whether this proportion was significantly greater than .5 (binomial test). Next, we computed the choice probability (the proportion of times the configuration with higher MEG was chosen) for each subject and each configuration pair for which the two MEGs differed. We then tested whether the majority of these choice probabilities were above .5 using a one-tailed binomial test.

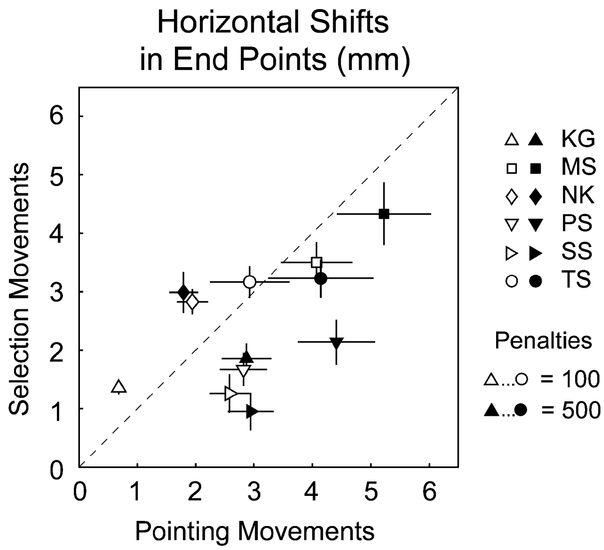

Fig. 2. Horizontal shift in end point away from the center of the target for pointing versus selection movements in the near condition (i.e., displacement of the target relative to the penalty region = 1 radius). Positive values indicate shifts away from the penalty region. Results for penalties of −100 and −500 are graphed separately for the 6 subjects. Error bars indicate ±1 *SEM*.

RESULTS

Pointing and Selection Movements Rely on the Same Movement Plan

We considered whether pointing movements directed at a single configuration differed from movements to select one of two configurations. The latter movements involved a choice about which configuration to point at. We therefore examined whether reaction and movement times for pointing movements differed from those for selection movements. Reaction times were shorter for selection (Sessions 2 and 3) than for pointing (Session 1) movements (Wilcoxon signed rank test, $p < .01$ and $p_{rep} > .95$[3] or better for all subjects). It is surprising that reaction times were slightly faster for selection movements, as average reaction times typically increase with the number of response alternatives (Hick, 1949; Hyman 1953). As Sessions 2 and 3 were always run after Session 1, we attribute this reduction in reaction time to increased practice, not to a difference in movement plan.

We next compared movement times to determine whether movement dynamics differed between pointing and selection movements. Movement times differed significantly across spatial and penalty conditions for pointing and selection movements for all subjects (Wilcoxon signed rank test, $p < .05$ and $p_{rep} > .88$ or better for all subjects), showing no consistent trend.

We also compared end points for pointing and selection movements. Most subjects moved their mean end point further away from the penalty region when the penalty increased (Fig. 2) and when the penalty region was closer to the target region (data not shown), as we have found previously (Trommershäuser et al., 2003a, 2003b). This shift was similar for movements directed at a single configuration compared with movements selecting one of two configurations ($r = .439$, $p = .032$, $p_{rep} = .91$; correlation computed across end points from Session 1 and from Sessions 2 and 3 for all conditions in which the penalty value was nonzero).

Efficiency was equally high for pointing and selection movements (pointing: 94.4–118.4% across subjects; selection: 97.8–107.1% across subjects). No subject's performance in either condition differed significantly from MEG. We conclude that pointing and selection movements rely on the same movement plan.

Subjects Prefer Configurations With Higher Expected Gain

We examined whether subjects based their selection of a configuration on an estimate of expected gain. We first determined the MEG for each configuration (again based on the subject's variability σ; the value of σ ranged from 3.3 to 3.8 mm across subjects). We calculated for each subject the proportion of trials on which the higher-MEG configuration was chosen and found that subjects chose the configuration with higher MEG on the majority of selection and key-press trials (selection: proportion ranged from .73 to .87 across subjects; key-press: .66 to .91 across subjects; $p < .001$ for all subjects for both selection and key-press trials). Next, we computed the proportion of choice probabilities (the proportion of times the configuration with higher MEG was chosen) that were greater than .5, pooled over subjects and configuration pairs that differed in MEG. This proportion was significantly greater than .5 (fraction of configuration pairs; selection movements: .883, $p < .001$; key-press movements: .877, $p < .001$).

Finally, we asked whether subjects based their selection judgments on estimates of MEG or on some monotonic transformation of MEG, $MEG' = \psi(MEG)$. We modeled the choice process as a signal detection observer that attempts to select the higher-MEG configuration given estimates of the two values of MEG' perturbed by equal-variance Gaussian noise ε. We introduced ε to model the subject's uncertainty in assessing MEG' rapidly and assumed that, on each trial, the subject formed the decision variable Δ:

$$\Delta = \psi(MEG_2) - \psi(MEG_1) + \varepsilon, \tag{6}$$

where MEG_i ($i = 1, 2$) denotes the MEG of each configuration, ψ is an increasing transformation, and the subject's uncertainty ε: *Normal* $(0, \sigma_G^2)$ is modeled by additive Gaussian noise. We assumed the subject chose the second configuration if and only if the value of Δ was greater than 0. We fit this model to the choice data for each subject and used a nested hypothesis test (Mood, Graybill, & Boes, 1974, pp. 441 ff.) to test the hypothesis that the transformation ψ is the identity.

[3]The p_{rep} statistic is the probability of finding the observed effect in the same direction (e.g., experimental mean higher than control mean) if the study were repeated without changes in methods (Killeen, 2005).

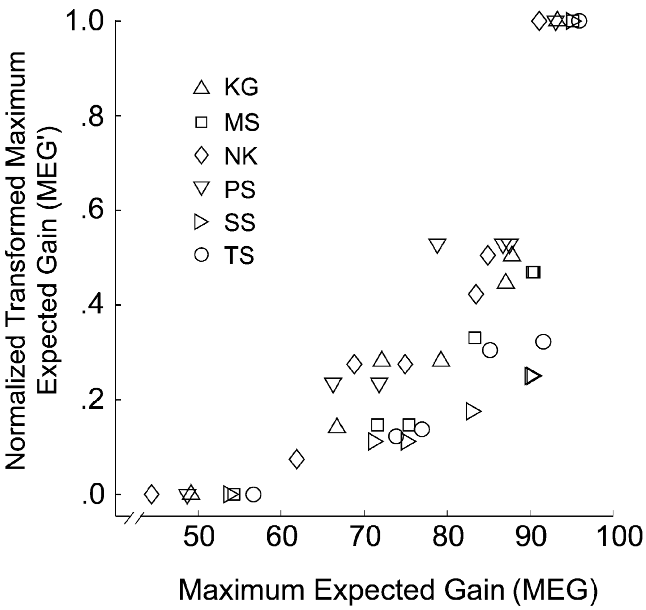

Fig. 3. **Transformed maximum expected gain (*MEG′*) as a function of *MEG*, for each configuration and subject in the selection-movement sessions. *MEG′* values for a given subject were scaled to range from 0 to 1 (equivalent to a different value of σ_G^2 for that subject) to emphasize the similarity of the form of the nonlinearity across subjects.**

The choice probabilities were fit by two models. In the simpler model, we assumed that selection was based on the actual monetary outcomes assigned to each configuration (i.e., $MEG' = MEG$), and only σ_G^2 was varied to fit the data (fit values for σ_G ranged from 13.10, for M.S., to 28.03, for P.S.). In the more elaborate model, σ_G^2 was fixed, *MEG′* of one configuration was fixed, and the other values of *MEG′* were varied to fit the data, subject to the constraint that they preserved the order of the corresponding *MEG* values. This resulted in six free parameters, as the two zero-penalty configurations shared the same value of *MEG*, and hence of *MEG′* as well. We rejected the hypothesis that ψ was linear ($p < .001, p_{\text{rep}} > .99$, for all subjects). Figure 3 shows that for selection movements, *MEG′* was a convex function of *MEG* and that the form was remarkably consistent across subjects. Figure 4 plots choice probability separately for each subject and configuration pair as a function of the difference in *MEG′* values for the two members of each pair. The quality of the fit is excellent.

CONCLUSION

We studied human movement planning under risk in a task in which selecting the movement goal involved an explicit choice between two possible goal configurations that differed in expected gain (i.e., the monetary rewards that could be expected when pointing at that configuration). Subjects preferred configurations with higher expected gain, regardless of whether they selected one of the configurations by pointing at it or by pressing a button. Movements directed at a single stimulus configuration exhibited the same movement dynamics as movements selecting one of two configurations. Selection movements did not differ from pointing movements and were executed with the same high efficiency. Our results suggest that movements under risk rely on rapid judgments about expected gain and that subjects base their judgments on internal estimates of expected gain that are a nonlinear function of actual expected gains.

We have argued that movement tasks are formally equivalent to decision making under risk. However, in marked contrast to the grossly suboptimal performance of human subjects in traditional economic decision-making experiments, our subjects' performance was often indistinguishable from optimal. Our results are consistent with the findings of Gigerenzer and Goldstein (1996) and of Weber, Shafir, and Blais (2004): Decision makers have difficulty reasoning with explicitly stated probabilities. Weber et al. found that experience-based choices do not exhibit the same suboptimality as pencil-and-paper tasks involving explicit probabilities. These results hint that the suboptimalities of human decision makers in the latter tasks are not characteristic of the large number of movement decisions people make each day.

Acknowledgments—We thank Tim Schönwetter for help with data collection. This work was supported by the National Institutes of Health (Grant NIH EY08266) and by the Deutsche Forschungsgemeinschaft (Emmy-Noether-Programme; Grant TR 528/1-2).

REFERENCES

Bell, D.E., Raiffa, H., & Tversky, A. (Eds.). (1988). *Decision making: Descriptive, normative and prescriptive interactions*. Cambridge, England: Cambridge University Press.

Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.

Eckstein, M.P., Beutter, B.R., & Stone, L.S. (2001). Quantifying the performance limits of human saccadic targeting during visual search. *Perception*, *30*, 1389–1401.

Fitts, P.M., & Petersen, J.R. (1964). Information capacity of discrete motor responses. *Journal of Experimental Psychology*, *67*, 103–112.

Gigerenzer, G., & Goldstein, D.G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.

Hick, W.E. (1949). Reaction time for the amendment of response. *Quarterly Journal of Experimental Psychology*, *1*, 175–179.

Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, *4*, 11–26.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, England: Cambridge University Press.

Kahneman, D., & Tversky, A. (Eds.). (2000). *Choices, values & frames*. New York: Cambridge University Press.

Killeen, P.R. (2005). An alternative to null-hypothesis significance tests. *Psychological Science*, *16*, 345–353.

Maloney, L.T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical*

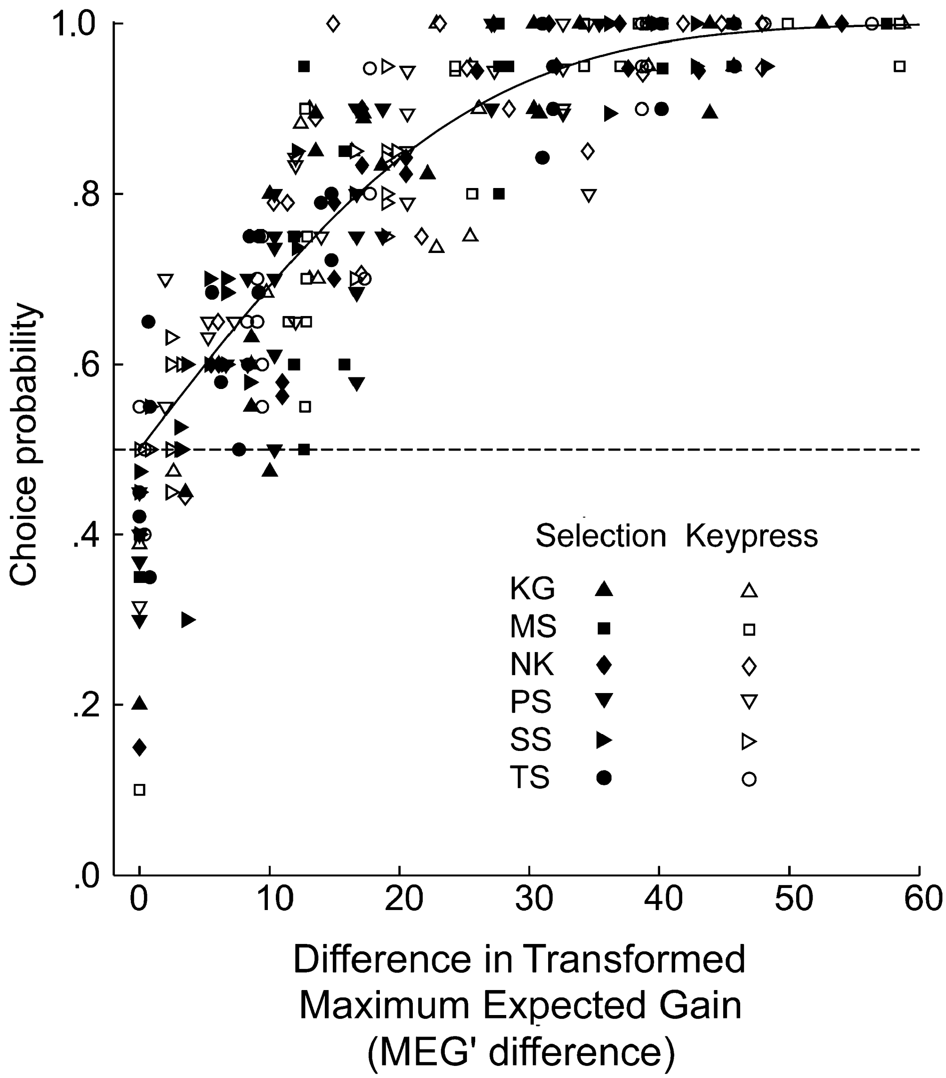

Fig. 4. Choice probability as a function of difference Δ in transformed maximum expected gain, *MEG′* (from Fig. 3; see also Equation 6). Data for each subject and for selection and key-press movements are graphed using different symbols. The solid curve is the prediction of the model.

world: Psychological and philosophical issues in perception (pp. 145–189). New York: Wiley.

Meyer, D.E., Abrahams, R.A., Kornblum, S., Wright, C.E., & Smith, J.E. (1988). Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review*, *95*, 340–370.

Mood, A., Graybill, F.A., & Boes, D.C. (1974). *Introduction to the theory of statistics* (3rd ed.). New York: McGraw-Hill.

Najemnik, J., & Geisler, W.S. (2005). Optimal eye movement strategies in visual search. *Nature*, *434*, 387–391.

Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.

Platt, M.L., & Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*, 233–238.

Roitman, J.D., & Shadlen, M.N. (2002). Response of neurons in the lateral intra-parietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, *22*, 9475–9489.

Sugrue, L.P, Corrado, G.S., & Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, *304*, 1782–1787.

Trommershäuser, J., Maloney, L.T., & Landy, M.S. (2003a). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A*, *20*, 1419–1432.

Trommershäuser, J., Maloney, L.T., & Landy, M.S. (2003b). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision*, *16*, 255–275.

Weber, E.U., Shafir, S., & Blais, A.-R. (2004). Predicting risk-sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.

(Received 11/11/05; Revision accepted 2/15/06; Final materials received 2/21/06)