

Humans rapidly estimate expected gain in movement planning

Julia Trommershäuser¹, Michael S. Landy², Laurence T. Maloney²

¹Giessen University, Department of Psychology, Otto-Behaghel-Str. 10F
35394 Giessen, Germany;

²New York University, Department of Psychology and Center for Neural Science,
6 Washington Place, New York, NY 10003, USA

Corresponding Author:

Dr. Julia Trommershäuser
Department of Psychology
Otto-Behaghel-Str. 10F
35394 Giessen, Germany
Email: Julia.Trommershaeuser@psychol.uni-giessen.de
Phone: +49 641 99 26118, Fax: +49 641 99 26119

Abbreviated Title: Rapid estimation of expected gain

Abstract (number of words): 150 (max. 150)

Total number of words (text): 4000 (max. 4000)

Number of Figures: 4

Number of Tables: 0

Number of Pages: 26

Key Words: Statistical decision theory, movement under risk, decision making, human psychophysics

Acknowledgments: We thank Tim Schönwetter for help with data collection. Support: DFG (Emmy-Noether-Programme, TR528 1-2), NIH EY08266.

Abstract

We studied human movement planning in tasks where subjects selected one of two goals that differed in expected gain. Each goal configuration consisted of a target circle and a nearby 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 keypress. 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 base the selection of their strategy on a rapid judgment about the expected gain assigned to each configuration.

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

When executing speeded arm movements under risk, humans select movement strategies that are nearly optimal (Trommershäuser, Maloney, & Landy, 2003a,b). In these studies, subjects pointed rapidly at stimulus configurations consisting of a small target and nearby penalty region. Reaches terminating within the target region yielded monetary reward, those ending in the penalty region could result in a loss. Target size and the distance between target and penalty regions were small (< 2 cm), similar in size to the subject's 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/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 screen (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., 2003a). The choice of strategy fixes the probability $P(R_i|S)$ of hitting each of the four regions R_i ($i=1,\dots,4$). In the decision making literature, this combination of event probabilities $P(R_i|S)$ and associated gains G_i is called a “lottery” $L(S)$,

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

An alternative movement strategy S' corresponds to a second lottery

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

As illustrated in Fig. 1A, every mean end point results in a lottery with a corresponding expected gain, i.e., the expected number of points a subject will earn, on average, having “aimed” at (\bar{x}, \bar{y}) . Fig. 1A indicates the expected gain of several mean end points, based on the measured end point variability of subject NK from the experiment reported below. The expected gain of movements aimed at the black triangle is less than that corresponding to 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 NK effectively selects among possible lotteries.

The results of our previous experiments indicate that subjects choose strategies maximizing expected gain (MEG), or nearly so. Efficiency was defined as the number of points won relative to the number expected using an optimal strategy. Subjects’ efficiencies were typically above 90% (Trommershäuser et al., 2003a,b). These efficiencies are higher than those derived from 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 based on an estimate of expected gain. We forced subjects to rapidly choose one of two configurations differing in expected gain, which requires the subject to make a quick estimate of the expected gain of each configuration. A configuration's expected gain is a complex function of motor noise, the pay-offs for each region, and the spatial arrangement of regions (Fig. 1B, see also "Model of optimal movement planning"). We asked whether subjects were able to rapidly estimate and compare the expected gains of two configurations.

Materials and Methods

Apparatus

The apparatus was similar to that described in Trommershäuser et al. (2003a,b). Subjects were seated in front of a monitor equipped with a touch screen (AccuTouch, Elo Touchsystems) 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. This start position was located 27 cm in front of and 34 cm below the screen center. The viewing distance was 52 cm. The experiment was run using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) on a Dell Optiplex PC (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 sessions of 50 minutes duration, run on six 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 stimulus configurations, either by rapidly pointing at one of two displayed configurations (Session 2/3, *selection movements*), or by button press (Session 4/5, *keypress 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 to 5) to make sure subjects assigned the correct color to each penalty value. We refer to movements directed at one of two configurations as selection movements (in contrast to pointing movements directed at a single configuration).

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

In the training session, there were two possible displacements of the penalty circle away from the target circle (“far” or $2R$ and “middle” or $1.5R$) and two penalty values (0

and -100 points). The session started with 30 warm-up trials with only a single green circle and a time limit of 1.5 s, followed by another 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 the target and penalty circles. The penalty was 0 points in the first block and -100 points in the second. Each block contained 9 repetitions of each of the 4 spatial configurations (middle and far, penalty displaced left and 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. The blocks consisted of eight repetitions of each of four stimulus configurations (leftward and rightward displacements of the penalty circle by R or $1.5R$, called near and middle, respectively). Successive blocks had penalty values of 0, -100, and -500 points, respectively.

The two following selection movement sessions (Sessions 2 and 3) consisted of two sessions of 342 trials (a total of 660 trials, split in two, plus 12 warm-up trials each). The 660 trials were 20 repetitions (completely randomized design) of each possible pair of stimulus configurations chosen from the eight displayed in Fig. 1B, excluding pairs of configurations with only 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. One of the configurations, chosen randomly, was displayed to the left of the screen center, the other one to the right. The amount of shift from screen center

was in the range 15-35 pixels (4.5-10.5 mm) 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 / 2.1-8.1 mm). These displacements were in opposite directions to avoid accidental hits into the “wrong” configuration. The time limit for the response was the same as in Session 1 (700 ms).

The final two keypress 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 the response was still 700 ms, even though subjects needed less time to complete these button press responses.

Procedure

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 she or he would later use 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 appeared. Subjects viewed the stimulus configuration for 400 ms after which a 1000 Hz tone indicated that the subject should start the movement.² A trial was aborted if the subject released the space bar earlier than 100 ms after presentation of the tone. Subjects were required to touch the screen within 700 ms after the tone to avoid a timeout penalty of 700 points. If the subject touched the screen within the area of the red or the green target, the region(s) hit “exploded” graphically and the subject received the points associated with each region. Then, the subject received feedback about the points scored in that trial, followed by the subject’s total accumulated points for that session.

In keypress movement trials (Sessions 4 and 5), subjects were instructed to select one of two configurations (as in Sessions 2 and 3), but made their choice by pressing one of two keys. They were told rewards and penalties would be based on performance in Sessions 2 and 3. Once the subject made a choice, the points for that trial were awarded by simulating the subject's movements based on that subject's performance in Sessions 2 and 3. The simulated end point for each subject 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 3.3-3.8 mm across subjects). If the simulated end point fell within the area of the red or the green target, the respective target(s) "exploded" graphically and the points associated with that region were awarded. Then, the subject received feedback exactly as in Sessions 2 and 3.

Subjects and Instructions

Six subjects participated in the experiment. The subjects were four male and two female students at the University of Giessen, and ranged in age from 22 to 29. All subjects but one were right-handed and all used their dominant hand to perform the experiment. All subjects had normal or corrected-to-normal vision. All subjects had given their informed consent before testing and were paid for their participation. All were unaware of the hypothesis under test. Subjects were informed of the payoffs and penalties before each block of trials. Subjects were told that the overall score over the six sessions would be converted into a bonus payment of 10 cents per 1000 points to motivate fast, accurate responses.

Model of Optimal Movement Planning and Choice Between Lotteries

In previous work, we have developed a model of optimal movement planning based on statistical decision theory (Trommershäuser et al., 2003a,b). 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 et al., 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,

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

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.

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/background (Region R_4 , gain $G_4 = 0$). The probability of hitting inside region R_i is defined by

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

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 ($i=1, \dots, 4$). 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 between one of infinitely many possible mean end points on the screen, the subject effectively selects between 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., 2003b for typical examples of distributions of expected gain). Our previous studies indicate that subjects choose mean end points corresponding to lotteries that nearly maximized 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 higher MEG configuration, we can conclude that they effectively have an internal ordering of the configurations based on an estimate of each configuration’s expected gain. We will compare the subjects’ preference for each configuration to that predicted by MEG. This prediction is different for each subject, and is computed using each subject’s end point variability σ . Fig. 1B shows the ordering by MEG, the corresponding optimal mean end points $(\bar{x}_{OPT}, \bar{y}_{OPT})$, and lotteries corresponding to maximum expected gain for one of the subjects (subject NK, with $\sigma = 3.81$ mm).

Data Analysis

For each trial we recorded *reaction time* (time between the tone and release of the space bar), *movement time* (time from movement onset until the screen was hit), the *screen position* that was hit, and the *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 the analysis. Data points that were further than 2 cm from any target center were classified as errors (e.g., knuckle hits) and were excluded from the analysis. Each subject contributed approximately 384 data points in Session 1, and approx. 660 data points each for Sessions 2/3 and Sessions 4/5 (20 repetitions of each pair of configurations). Movement end points were recorded relative to the center of the target circle. Movement end point data did not differ significantly for left-right symmetric configurations and were collapsed across symmetric conditions for data analysis.

Efficiency for different types of judgments. We asked whether efficiency in the motor task differed for pointing movements compared to selection movements. We computed efficiency for performance in Session 1 and in Sessions 2 and 3 (data collapsed across both sessions; see Results). 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 across all conditions. 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. (2003b) on how to compute this range of efficiencies.

Preference for configurations with higher MEG. We asked whether subjects preferred configurations with higher expected gain. For each configuration used in Sessions 2/3 and 4/5, we computed the MEG (based on each subject's variability and the optimal strategy). We emphasize that the subject is choosing between two configurations, each of which has an MEG that depends on the configuration and the subject's motor uncertainty.

To test whether subjects preferred the configuration with higher MEG, we computed the proportion of trials in which the configuration with higher MEG was chosen and tested whether that proportion was greater than 0.5 (binomial test, see Results). Next, for each pair of configurations that differed in MEG we computed the subjects' choice probability, i.e., the proportion of times the configuration in this pair with higher MEG was chosen. We tested whether the majority of pairs had choice probabilities above 0.5 using a one-tailed binomial test.

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 type of movements involves 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 (Session 2/3) than for pointing (Sessions 1) movements (Wilcoxon Signed Rank Test, $p < 0.01$ and $p\text{-rep} > 0.95$ or better for all subjects).³ 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 difference to increased practice, not to a difference in movement plan.

We next asked whether movement dynamics differed for pointing and selection movements by comparing their movement times. Movement times differed significantly for pointing and selection movements for all subjects (Wilcoxon Signed Rank Test, $p < 0.05$ and $p\text{-rep} > 0.88$ or better for all subjects), with no consistent trend.

We next compared end points for pointing and selection movements. Most subjects moved their mean end point away from the penalty region with increasing penalty (Fig. 2) and with closer penalty region (data not shown), as we have found previously (Trommershäuser et al., 2003a,b; data displayed here only for the near condition). This shift was similar for movements directed at a single configuration compared to movements selecting one of two configurations ($r=0.439$, $p=0.032$ / $p\text{-rep}=0.91$, correlation computed across end points from Session 1 and Sessions 2 and 3 for all conditions in which the penalty value was < 0). 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 asked whether subjects based their judgment on an estimate of expected gain when selecting one of two configurations. We first determined the MEG for each configuration (based, as above, on the subject's variability σ , σ ranging from 3.3 to

3.8 mm across subjects). Subjects chose the configuration with higher MEG for the majority of selection and keypress trials. We calculated the fraction of configuration pairs for which the higher-MEG configuration was chosen more than half the time (pointing: 0.73-0.87 across subjects; selection: 0.66-0.91 across subjects). This fraction was significantly more than half the pairs (selection: .883, $p < .001$; keypress: .877, $p < .001$).

Finally, we asked whether subjects based their selection judgment on an estimate 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 configurations given estimates of the two values of MEG' perturbed by equal variance Gaussian noise ε . We introduce ε to model the subject's uncertainty in assessing MEG' rapidly and assume that, on each trial, the subject forms the decision variable

$$\Delta = \psi(MEG_2) - \psi(MEG_1) + \varepsilon, \quad (6)$$

where MEG_i ($i = 1, 2$) denotes the maximum expected gain of each configuration, ψ is an increasing transformation, and the subject's uncertainty $\varepsilon : Normal(0, \sigma_G^2)$ is modeled by additive Gaussian noise. We assume the subject chooses the second configuration if and only if $\Delta > 0$. We fit this model to the choice data for each subject and tested the hypothesis that the transformation ψ is the identity by a nested hypothesis test (Mood, Graybill & Boes, 1974, pp. 441ff). Choice probabilities in Fig. 4 were fit by two models. In the simpler model, we assume that selection is 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 between 13.10 for MS and 28.03 for PS). 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 6 free parameters, as the two zero-penalty configurations share the same value of *MEG*, and hence of *MEG'* as well. We reject the hypothesis that ψ is linear ($p < .001$ / $p\text{-rep} > 0.99$ for all subjects). Fig. 3 shows *MEG'* as a function of *MEG*. *MEG'* is a convex function of *MEG*; the form is remarkably consistent across subjects. Fig. 4 plots choice probability as a function of difference in *MEG'* for all subjects and configuration pairs. The quality of the fit is excellent.

Conclusion

We studied human movement planning under risk in a task in which selecting the movement goal involves 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, independent of whether they selected one of the configurations by pointing at the configuration 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 the expected gain associated with the configurations and that subjects base their judgment on an internal estimate of expected gain that is a non-linear function of actual monetary outcome.

We have argued that movement tasks are formally equivalent to decision making under risk. However, in marked contrast to the grossly sub-optimal 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. (2004) find that experience-based choices do not suffer from the same sub-optimal decisions as pencil and paper tasks involving explicit probabilities. These results hint that the sub-optimality of human decision makers in these tasks are not representative of performance in the large number of movement decisions we face each day.

References

- Bell, D. E., Raiffa, H., & Tversky, A. (Eds.) (1988). *Decision making: descriptive, normative and prescriptive interactions*. Cambridge, UK: Cambridge University Press.
- Brainard, D. H. (1997). The psychophysical 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. *J. Exp. Psychol., 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, UK: 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 Heyer, D. & Mausfeld, R. (Eds.), *Perception and the Physical World: Psychological and Philosophical Issues in Perception..* New York: Wiley, pp. 145-189.
- 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 Edition. 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 trade-offs in the control of motor response. *Spatial Vision*, *16*, 255-275.

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

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.

Figure captions

Fig. 1: Equivalence of the rapid pointing task and choice among lotteries. A) An example stimulus configuration is shown along with three possible visuo-motor strategies (i.e., mean end points) and the corresponding expected gain (points/trial) based on subject NK's end point variability ($\sigma = 3.81$ mm). B) Stimulus configurations used in Sessions 2-5 (selection and keypress movements). Subjects selected one of two configurations chosen randomly from pairs of the eight configurations shown here (or the spatially symmetric configurations where the target was on the left side of the penalty). The green target (indicated by a solid black line) and penalty areas were circular with radius $R = 28$ pixel / 8.4 mm. Penalty values were color coded; gray (bottom two configurations, indicated by a dashed line): penalty = 0 points; blue: penalty = -100 points (indicated by light gray); red: penalty = -500 points (indicated by dark gray). The target was displaced leftward or rightward of the penalty region. For each configuration we list the maximum expected gain (MEG) and the corresponding lottery (for subject NK). Configurations marked by an asterisk, were the ones used in Session 1 for movements directed at a single configuration (pointing movements).

Fig. 2: Motor responses for pointing vs. selection movements. The horizontal shift in end point away from the target center is plotted for the near location and penalty -100 and -500 conditions. Individual subjects are marked by different symbols. Error bars indicate \pm one standard error of the mean.

Fig. 3: MEG' as a function of MEG , for each configuration and subject (selection movements). MEG' values for a given subject were scaled to range from zero to one (equivalent to a different value of σ_G^2 for that subject) to emphasize the similarity of the form of the nonlinearity across subjects.

Fig. 4: Choice probability as a function of difference in MEG' (from Fig. 3). The solid curve is the prediction of the model.

Footnotes

¹ Decision making spans many fields and terminology differs from field to field. We prefer to use the terminology of gains/losses employed in statistical decision theory (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 that we use corresponds exactly to expected value in the psychological literature.

² The interval of 400 ms matches that for a representation of expected value to build up in area LIP as found in recordings from monkeys (Platt & Glimcher, 1999; Roitman & Shadlen, 2002; Sugrue, Corrado, & Newsome, 2004).

³ p-rep denotes the probability of replicating the result of the same sign (Killeen, 2005).

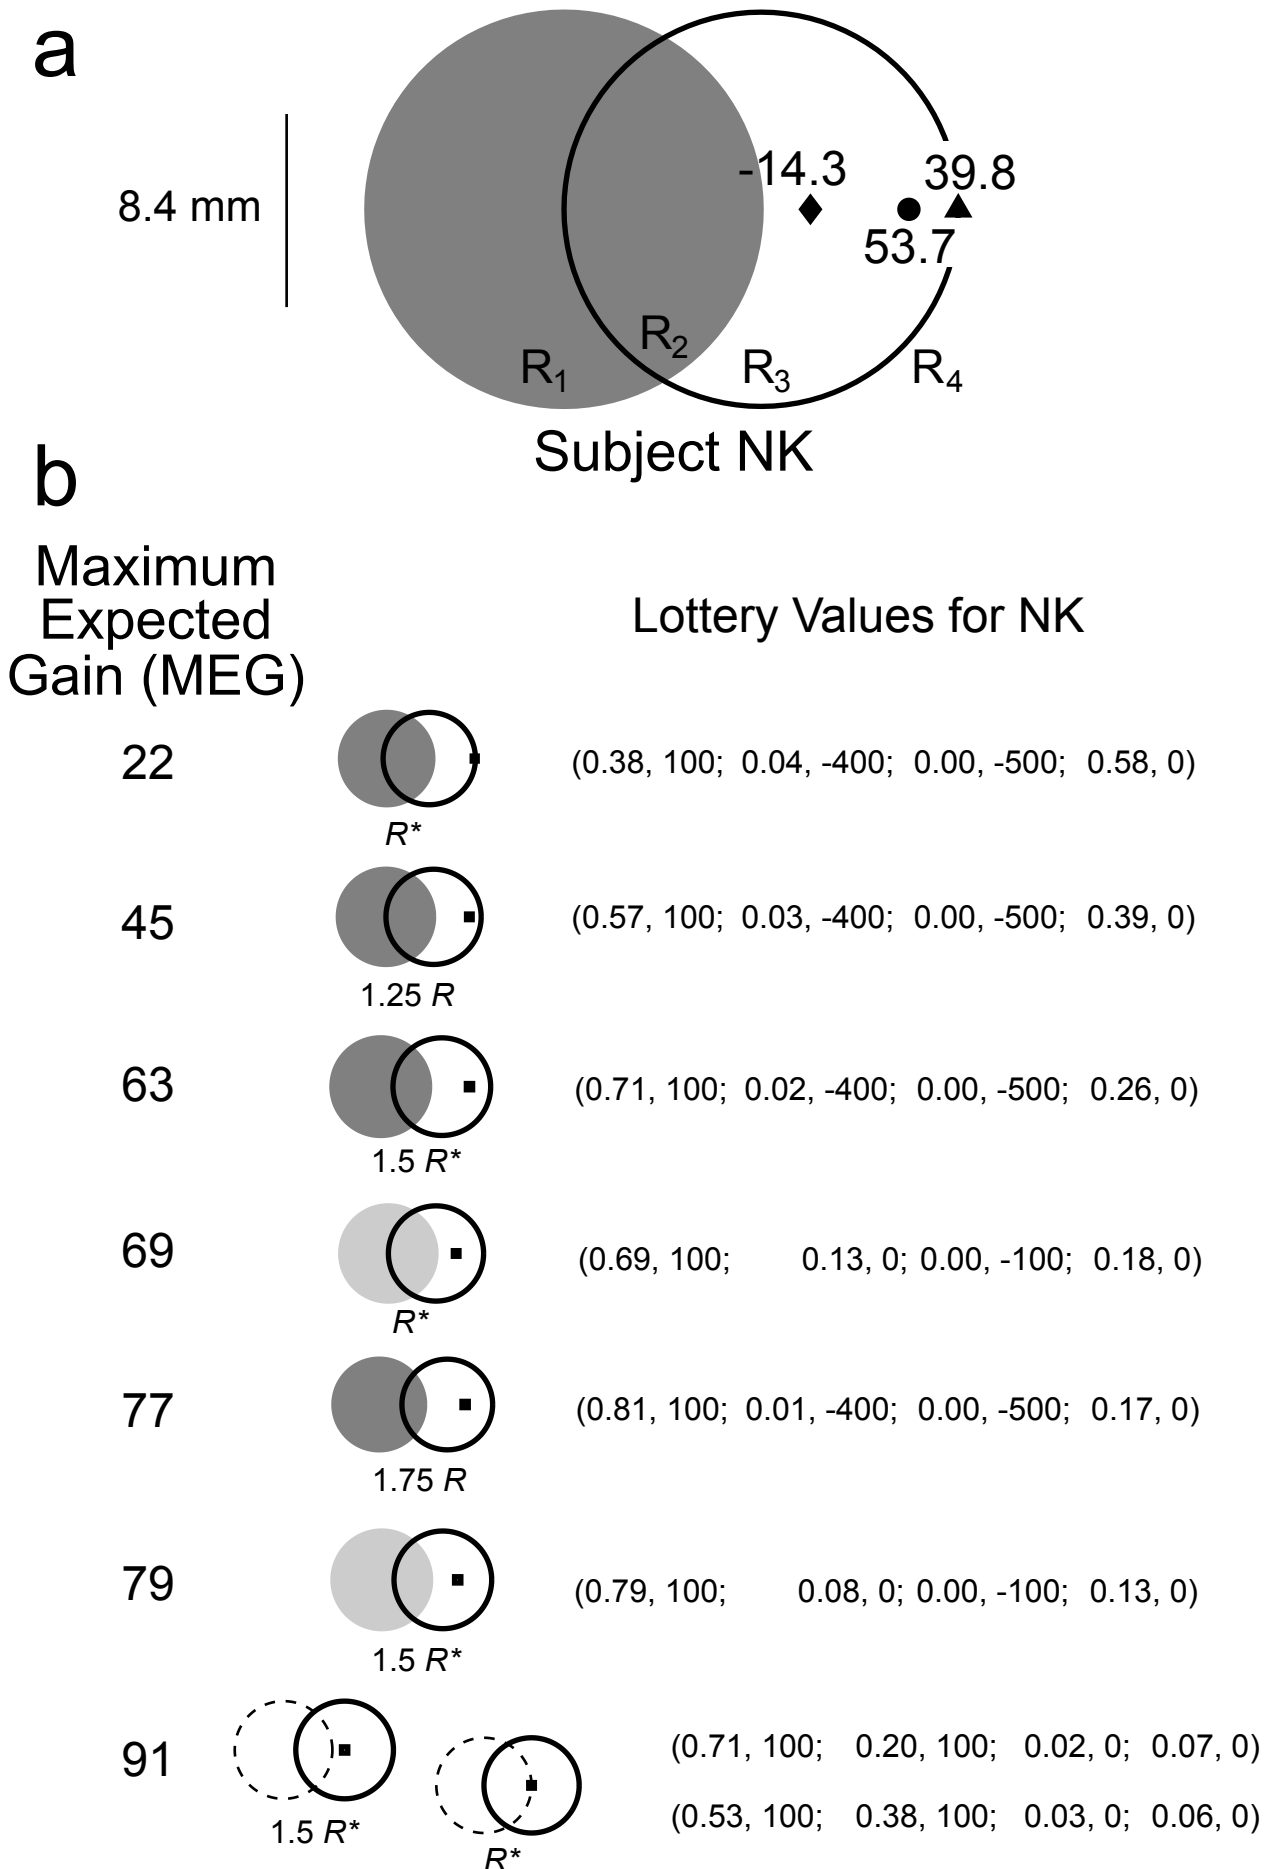


Fig. 1

Horizontal Shifts in End Points (mm)

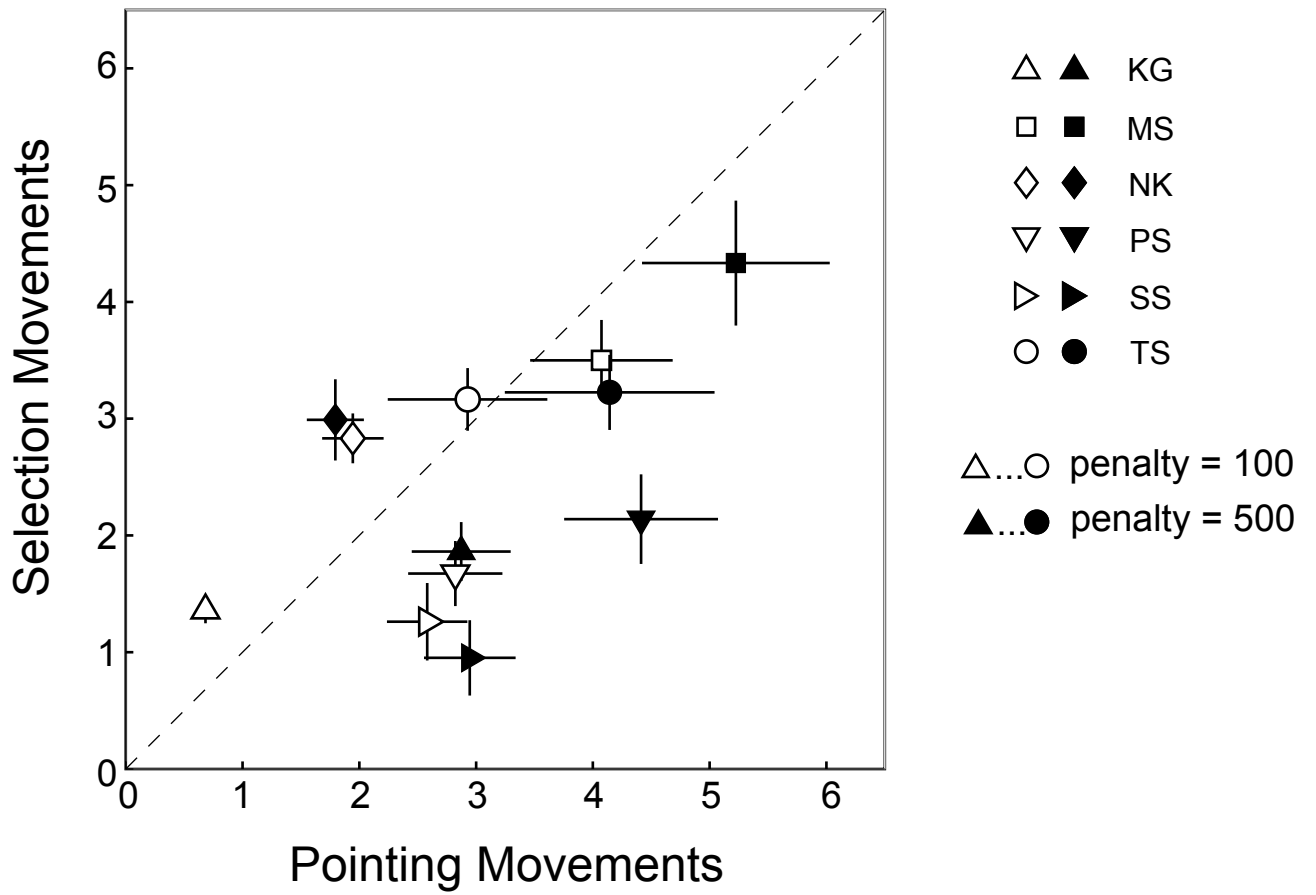


Fig. 2

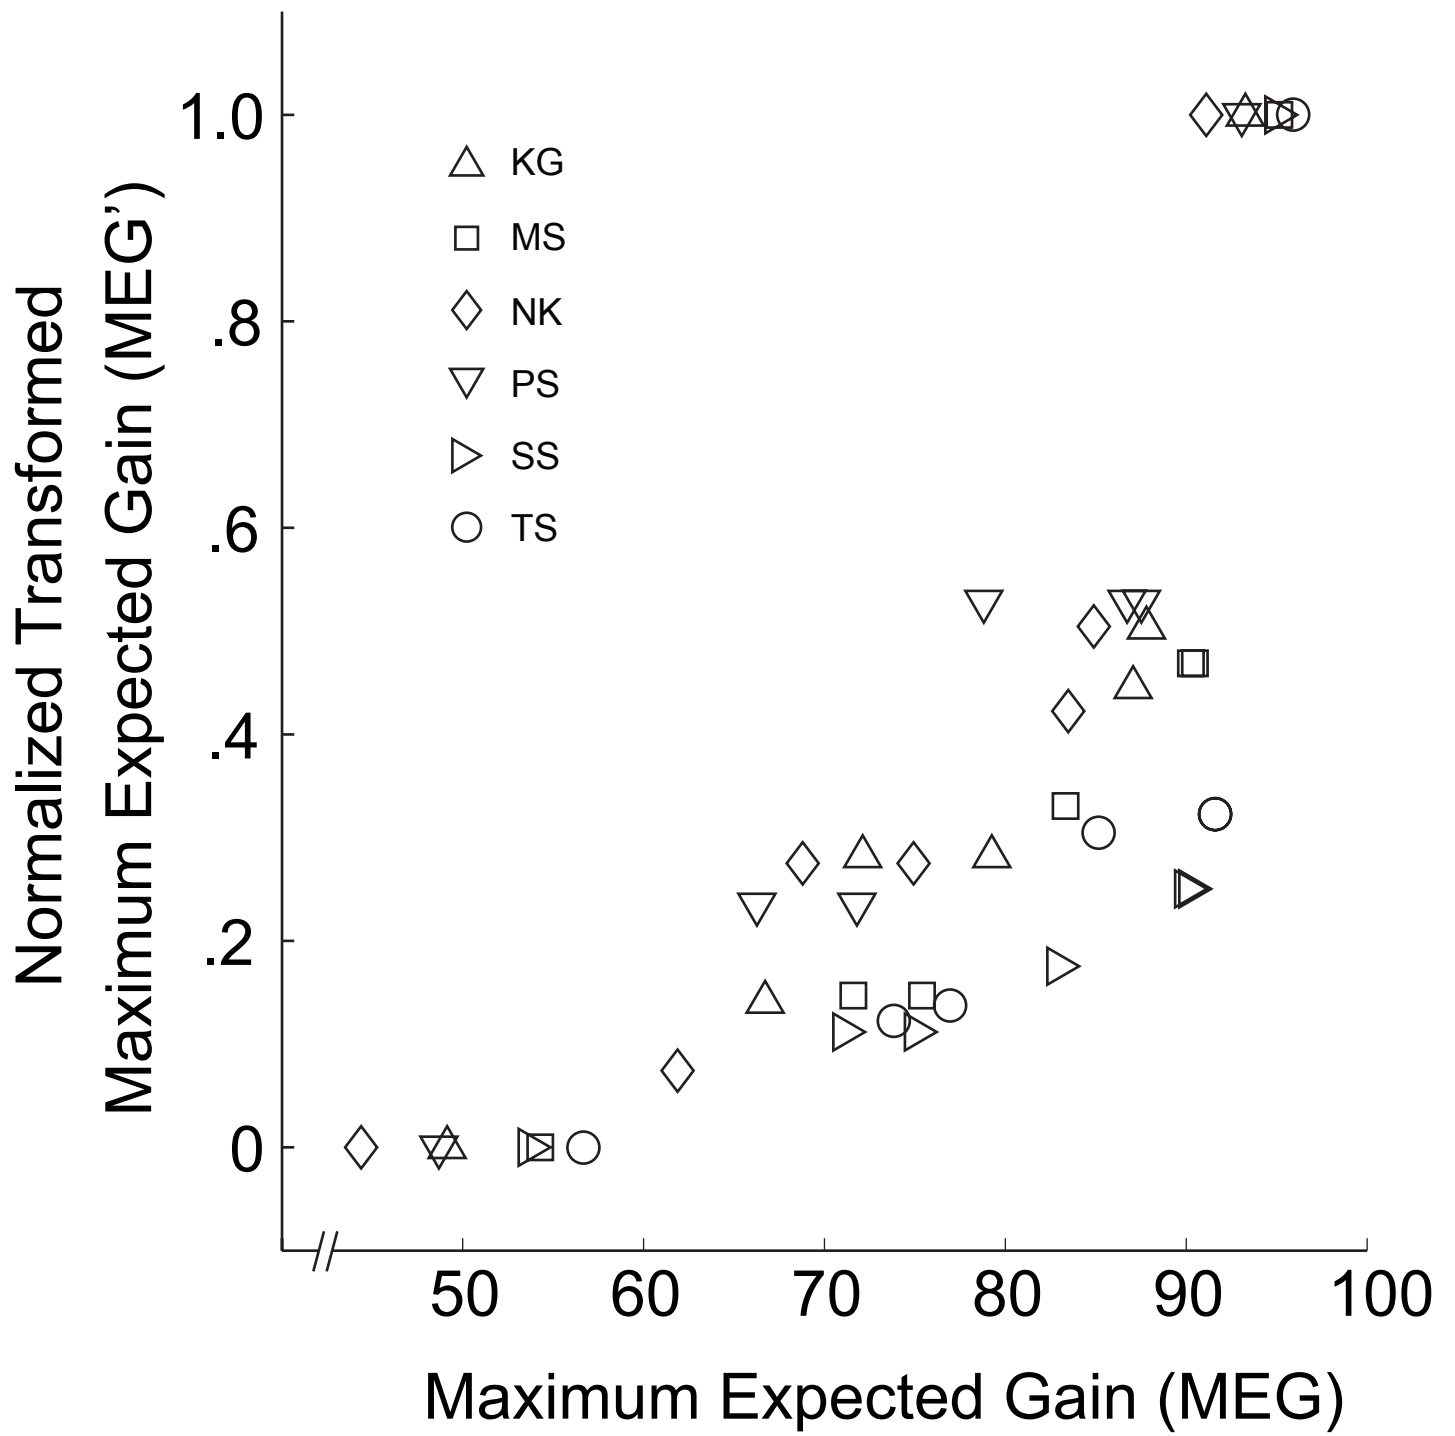


Fig. 3

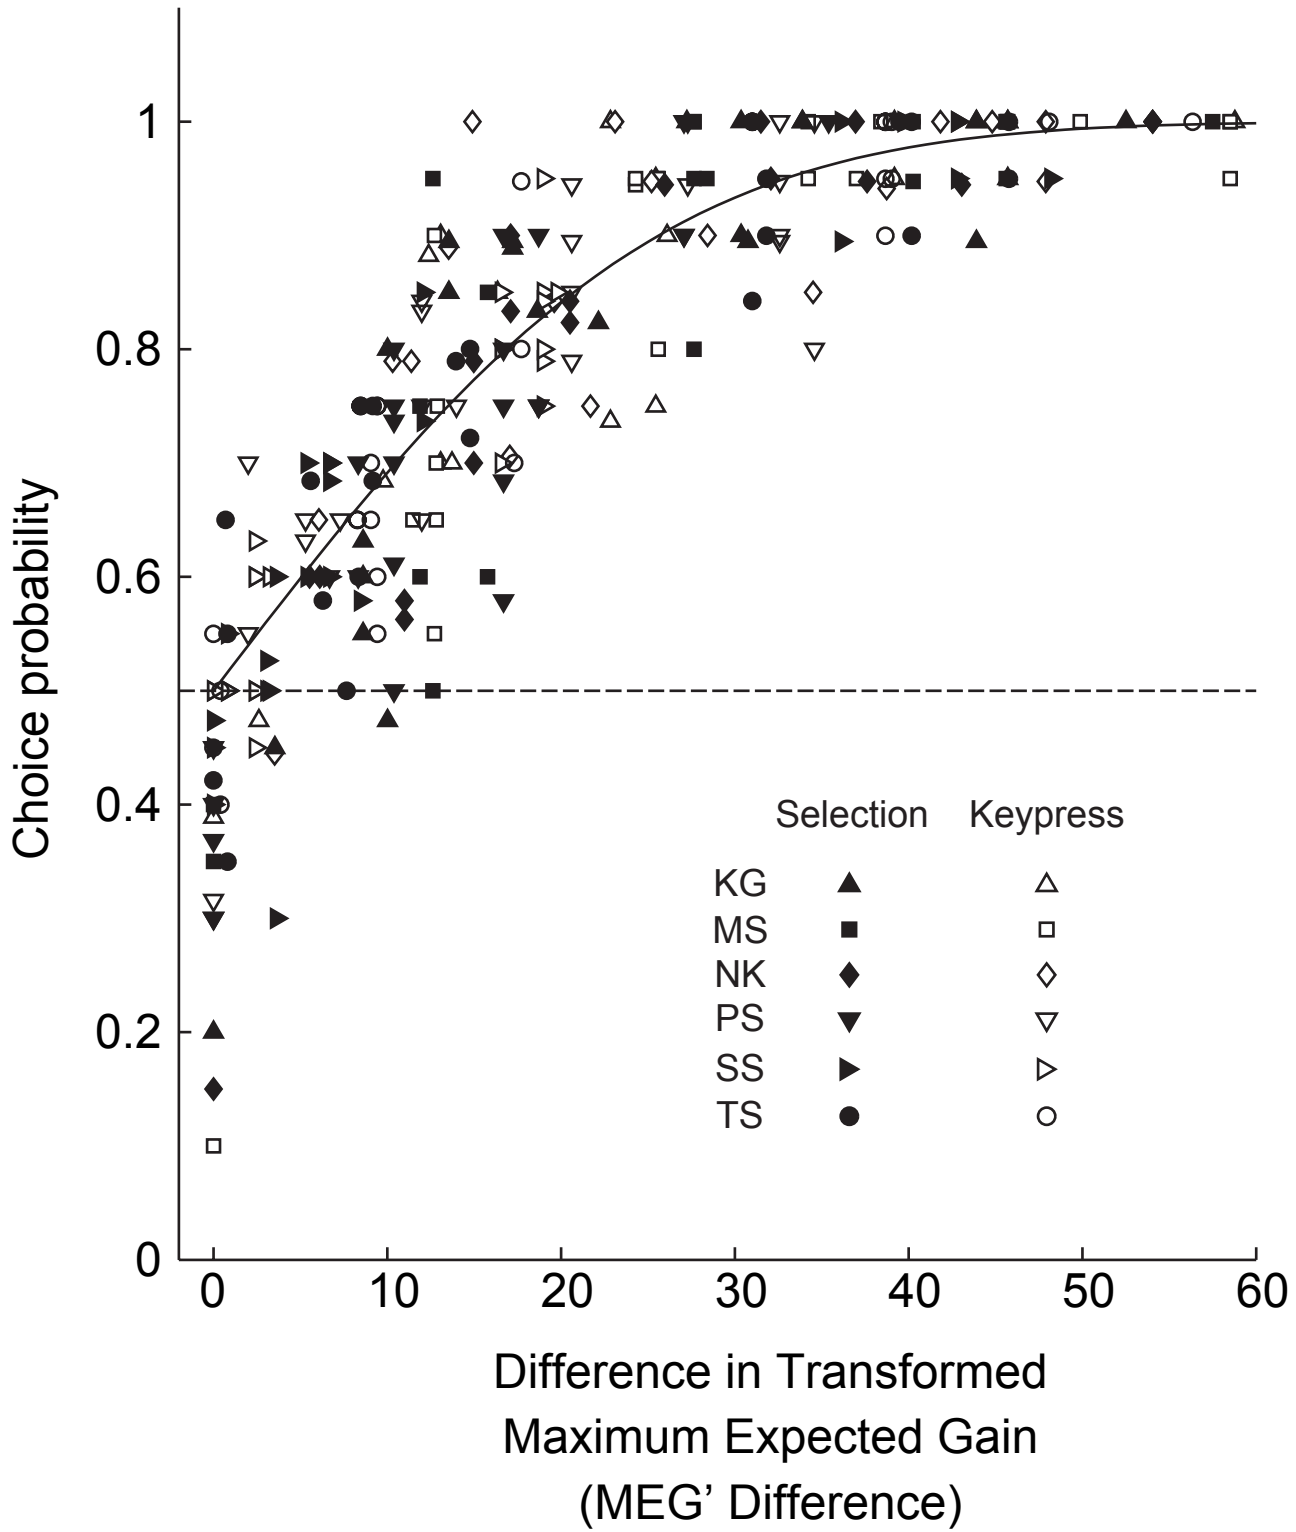


Fig. 4