DECISION MAKING AND MOVEMENT PLANNING UNDER RISK

Julia Trommershäuser[1,2], Laurence T. Maloney[2], Michael S. Landy[2]
1: Department of Psychology, Giessen University, Otto-Behaghel-Str. 10F, 35394 Giessen, Germany;
2: Department of Psychology and Center for Neural Science, New York University; 6 Washington Place, New York, NY 10003, USA

A decision maker chooses among plans of action. Whether it is checking "yes" in response to a survey question or choosing when to start the swing of a baseball bat, the chosen plan of action may have serious consequences. While the decision maker can plan an action, s/he cannot always anticipate the precise outcome of the decision and its consequences.

If the possible, mutually-exclusive outcomes of a decision are denoted $O_1, \cdots, O_n$ then the effect of any plan is to assign a probability to each outcome. The result is a *lottery* $(p_1, O_1; p_2, O_2; \cdots; p_n, O_n)$ where p_i denotes the probability of outcome O_i and $\sum_{i=1}^{n} p_i = 1$ (i.e., the list of outcomes is exhaustive). Decision making is, pared down to its essentials, a choice among lotteries.

When the outcomes are framed in terms of money, we refer to them as gains, denoted G_i, and we can assign to each lottery $L = (p_1, G_1; p_2, G_2; \cdots; p_n, G_n)$, an expected gain (Arnauld & Nicole, 1662/1992),

$$EG(L) = \sum_{i=1}^{n} p_i G_i \ . \qquad (1)$$

The decision maker who seeks to maximize gain can evidently do so by selecting the action whose corresponding lottery has the maximum expected gain (MEG). Such a MEG rule is an example of a *normative rule*, intended to guide decision making (Bell, Raiffa & Tversky, 1988). Much research in human decision making under risk is a catalogue of the many, patterned failures of normative theories, including MEG, to explain the decisions humans actually make (Bell et al., 1988; Kahneman, Slovic & Tversky, 1982; Kahneman & Tversky, 2000).

Deviations of human decision makers from MEG include a tendency to exaggerate small probabilities (Allais, 1953; Attneave, 1953; Kahneman & Tversky, 1992; Lichtenstein et al., 1978), to convert gain into subjective utility (Bernoulli, 1738/1954; von Neumann & Morgenstern, 1947), and to frame outcomes in terms of losses and gains with an exaggerated aversion to losses (Kahneman & Tversky, 1979). There are other well-documented deviations from MEG predictions and the degree and pattern of deviations depends on many factors. How these factors interact and affect decision making is controversial. What is not in dispute is that it takes very little to lead a human decision maker to abandon an MEG rule in decision making tasks.

The typical tasks used in the literature on decision making under risk are paper-and-pencil choices that one can "meditate on" before responding. Probabilities and values are arbitrarily chosen by the experimenter and are represented numerically or by simple graphical devices. In contrast, Trommershäuser, Maloney and Landy (2003a,b) introduced an experimental paradigm in which subjects were asked to plan and execute rapid movements in "risky" environments. The task on each trial was formally equivalent to decision making under risk (Fig. 1), but information about probabilities was not communicated to the subject "in words." Despite the lack of explicit description of probability, subjects in these tasks consistently selected motor plans that came close to maximizing expected gain. The focus of the experiments reported here is to explore the link between performance in these tasks and human decision making.

If subjects were perfectly in control of their movements, they would simply touch the green circle, avoiding the red whenever it incurred any penalty. However, subjects were given a time limit, resulting in movement end points with

Contact information:
Julia Trommershäuser, Department of Psychology, Giessen University, Otto-Behaghel-Str. 10F, D-35394 Giessen, Germany; ph: 01149-641-26118; fax: 01149-641-26119;
email: Julia.Trommershaeuser@psychol.uni-giessen.de
Acknowledgement: We thank Paulina Trzcinka, Diana Pittig, and Florian Bayer for help with data collection.

substantial scatter (Fig. 1A). Thus, a choice of a particular movement strategy *s* effectively selected a probability distribution $p_s(x,y)$ of possible end points on the touch screen.

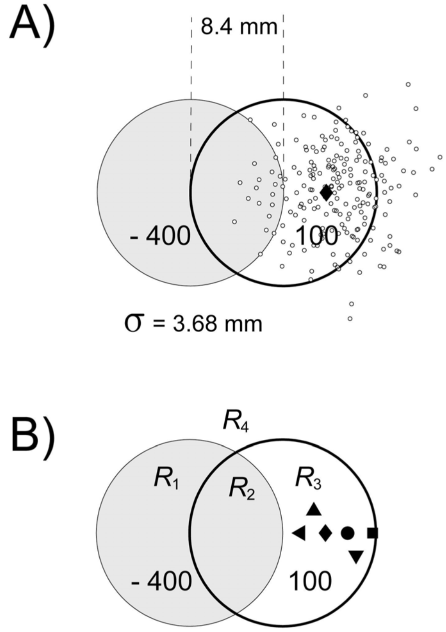

◄: (0.8%, -400; 25.3%, -300; 65.5%, 100; 8.4%, 0)
▲: (0.2%, -400; 13.1%, -300; 70.9%, 100; 15.8%, 0)
♦: (0.1%, -400; 10.8%, -300; 70.7%, 100; 18.4%, 0)
●: (0.0%, -400; 3.9%, -300; 62.4%, 100; 33.7%, 0)*
▼: (0.0%, -400; 2.1%, -300; 54.2%, 100; 43.7%, 0)
■: (0.0%, -400; 0.9%, -300; 43.1%, 100; 56.0%, 0)
⋮

Figure 1: Equivalence of the rapid pointing task (Expt. 1, certainty condition, see below) and choice among lotteries. A) Example of stimulus configuration along with a sub-optimal strategy (mean end point indicated by the diamond). In the experiment, the penalty region was a filled red circle, and the reward region a hollow green circle. 200 simulated responses are shown based on a bivariate Gaussian distribution and a motor uncertainty σ = 3.89 mm (corresponding to subject JM's estimated motor uncertainty). B) Six possible aim points are shown along with the equivalent lottery for each. For each lottery, the four possible outcomes are landing in the penalty (region R_1), the target/penalty overlap (R_2), the target (R_3), or neither (R_4). The corresponding probabilities are based on the given aim point and motor uncertainty. The asterisk marks the lottery with the highest expected gain.

The instructions to the subject were unusual for an experiment involving a motor task. Subjects were not instructed to hit the target nor were they told to avoid the penalty. Rather, they were instructed to earn as much money as possible, in any way they saw fit to do so. Probabilities were never mentioned. Yet, with these instructions, the task that they were asked to perform was equivalent to a choice among lotteries.

To see this, consider the possible outcomes when hits on the target and penalty yield gains of +100 and -400 points, respectively (Fig. 1B). A movement that hits the touch screen within the time limit could land in one of four regions: penalty only (Region R1, gain G1 = -400), target/penalty overlap (Region R2, gain G2 = -300), target only (Region R3, gain G3 = 100), or neither (Region R4, gain G4 = 0). The probability of each of these outcomes depends on the subject's choice of movement strategy *s*.

In the following, we identify a movement strategy *s* with a mean movement end point (i.e., an "aim point"). The diamond in Fig. 1A marks the mean of a Gaussian distribution of end points (with width σ) and, given this choice of movement strategy *s*, we can compute the probability of each of the four outcomes, denoted $p_1,\cdots,p_4$ (see Trommershäuser et al., 2003a,b, for details on how to compute p_i).

Therefore this choice of movement strategy *s* corresponds to the lottery

$$L(s) = (p_1, G_1;\ p_2, G_2;\ p_3, G_3;\ p_4, G_4) \quad (2)$$

while an alternative movement strategy s' (e.g., the small black circle in Fig. 1B) corresponds to a second lottery

$$L(s') = (p_1', G_1;\ p_2', G_2;\ p_3', G_3;\ p_4', G_4). \quad (3)$$

Fig. 1 is based on the measured end point distributions for subject JM from the first experiment reported below. For this configuration, the expected value of the lottery corresponding to movement strategy *s* is less than that corresponding to movement strategy s'. However, there are infinitely many other lotteries available to Subject JM, each corresponding to a particular motor strategy or aim point, and each with an associated expected value. Fig. 1B lists a subset of these strategies with associated probabilities. By choosing among these possible motor strategies, subject JM effectively selects among the possible sets of probabilities associated with each outcome.

Optimal Responses to Changes in Value

In one experiment of Trommershäuser et al. (2003b) the penalty for hitting the penalty was 0, -100 or -500 points, while the reward for hitting the target was constant (100 points). In another experiment, the value of the stimulus configuration was also varied, but this time 'spatially', by adding a second penalty region. We compared the performance of each subject to that of the movement strategy that maximized expected gain for that subject. We found that subjects' performance did not differ significantly from maximum expected gain in both experiments.

Optimal Responses to Changes in Uncertainty

In a second series of experiments, subjects' task-relevant movement variability was manipulated by randomly perturbing the visual feedback of the hand. Rewards and penalties were based on the perturbed, visually specified finger position (Trommershäuser, Gepshtein, Maloney, Landy, & Banks, under review). Subjects failed to counteract the visual perturbation of the finger position during the movement, but did rapidly acquire an estimate of their new variability. The change in variability caused a change in the probability distribution $p_s(x,y)$ of possible end points on the screen. Thus, the change in variability altered the probability of each of the four outcomes (p_1, …, p_4). We compared subjects' performance to the performance of an optimal movement planner maximizing expected gain, and found that subjects compensated optimally for externally imposed changes in their task-relevant variability. When exposed to novel stimulus configurations, subjects adjusted their aim points in the first trial without showing any further detectable trend across the following trials. Given the complexity of the decision making task implicit in Fig. 1B, the outcome of the experiments described here is remarkable.

We are left with a paradox. The movement planning task involves a speeded choice among infinitely many lotteries and yet, unlike performance in paper-and-pencil decision making tasks, subjects' performance is close to that required to maximize expected gain.

DECISION MAKING AND MOVEMENT UNDER RISK

Here, we report the results of three experiments that explored the link between movement planning under risk and human decision making.[1]

Expt. 1: Sub-optimal Performance due to 'Extrinsic' Uncertainties

In the first experiment reported here, we asked again what happens if the probabilities of scoring a reward or penalty change. But this time we manipulated the probabilities of incurring penalties and rewards by introducing stochastic rewards and penalties.

We modified the rapid pointing task (see Trommershäuser et al., 2003b, for a detailed description of apparatus and procedure) so that, in some conditions, the reward and/or penalty regions were stochastic: when the subject hit a reward or penalty region, s/he would receive the reward or penalty with probability 0.5. Penalty values were 0, -200 or -400 points; hitting the target always scored 100 points.

Subjects were tested on all four combinations of certain or stochastic rewards and penalties. Each condition was run in a separate session and the first and last sessions were always fully certain. The subject was explicitly told the probabilities before each session. Six naive subjects participated.

Five out of six subjects maximized expected gain in the certainty conditions, in both the first and the last session. However, five of the six subjects were markedly sub-optimal in one or more of the stochastic conditions (Fig. 2A). As shown in Fig. 2B, in the majority of cases,

[1] In all three experiments described here, pointing movements that hit the screen later than 700 ms after the start signal for the movement were penalized. All subjects performed an initial practice session to learn the time constraints of the task (270 trials, pointing at different configurations than in the experiment). Once subjects had adjusted their movements to meet the time constraints of the task, the rapid execution of the pointing movement led to substantial movement variability.

subjects shifted their movement end points in response to manipulations in explicit probabilities in the sub-optimal direction, i.e. closer towards the penalty when the chance of scoring a reward dropped to 50%, or further away from the penalty when the chance of scoring a penalty dropped to 50%. In the Both 50% condition, both the penalty and reward are stochastic, so that the optimal strategy is identical to that in the certainty condition (Allais, 1953). Yet, four out of six subjects changed their mean movement end points significantly in this condition. (Note that all subjects returned to their initial optimal strategy in session five when rewards and penalties were again scored with certainty).

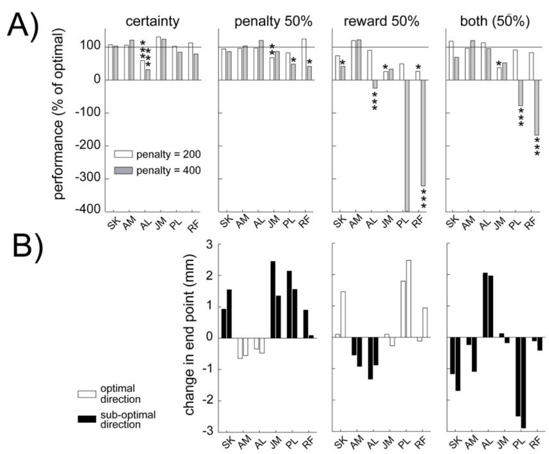

Figure 2: Performance of six subjects in rapid pointing task with stochastic rewards and penalties. A) Actual scores, normalized by the optimal score predicted by the MEG model. Model predictions were computed based on each subject's motor variance. Data are displayed for each subject individually. The solid horizontal line indicates perfect correspondence of model and experiment. B) Shift of mean movement end points with changes in extrinsic probability. Shifts are displayed as difference from the mean movement end point in the certainty condition. White bars indicate changes in end point into the optimal direction (with respect to end points in the certainty condition), black bar into the sub-optimal direction. Data are displayed for each subject individually and the penalty 200 (left bar) and 400 (right bar) conditions. Model predictions were computed based on each subject's motor variance.

In summary, associating extrinsic probabilities with rewards and penalties disrupted optimal performance in our movement planning tasks. While the results of our previous experiments (Trommershäuser et al., under review) suggest that human movement planners are capable of selecting a strategy that maximizes expected gain in the presence of *intrinsic uncertainties* (due to variability in sensory estimation and movement production), the introduction of *extrinsic uncertainties* (stochastic rewards and penalties) led to sub-optimal performance. These latter results are consistent with findings of Gigerenzer and Goldstein (1996) and of Weber, Shafir and Blais (2004) who report that decision makers have difficulty reasoning with explicitly stated probabilities.

Expt. 2: Stability of Performance after Prolonged Presentation Times

In Expt. 1, subjects responded optimally to changes of intrinsic uncertainty but were sub-optimal in response to changes of extrinsic uncertainty. Could this be directly related to the sub-optimal strategies found in traditional decision making tasks? To examine this, we varied the time between the target-penalty display and movement initiation to see if performance worsened in the motor decision making task as subjects had more time available to study the configuration of rewards and penalties.

In this experiment, we varied the time between stimulus display and the start signal for the pointing movement. The stimulus configuration was displayed, followed by a 400 Hz tone (0 ms, 400 ms or 1000 ms after stimulus display) indicating that subjects should initiate the movement. Subjects' movement end points and performance were unaffected by variation of stimulus presentation time. Human movement strategy was stable (and optimal) as the presentation time of the target-penalty display increased.

Expt. 3: Sub-optimal Performance due to Motor System Constraints

Next, we altered the pointing task in another way so as to make it more similar to traditional paper-and-pencil decision tasks. Five subjects were instructed to rapidly point at one of two simultaneously displayed stimulus configurations, each consisting of one target and one penalty region. In some trials, the two configurations differed in penalty value and/or spatial arrangement. Thus, the two configurations differed in the value of maximum expected gain. The stimulus was displayed, and 400 ms later a tone indicated the subject should initiate the movement. In this experiment, subjects' movement end points, for a given configuration, did not differ from the distribution of movement end points in trials with that configuration displayed alone (control experiment). In trials with two configurations with different maximum expected gain, four of five subjects pointed more frequently at the configuration with higher maximum expected gain (preferences of 81% to 55%). However, performance was sub-optimal for three out of five (right-handed) subjects due to a preference to point at the stimulus configuration presented in the right half of the screen (81% to 72% across the balanced design).

Thus, human movement strategies remained stable (and optimal) for these selection movements that indicated a choice among multiple configurations. But, the choice among configurations with different expected gain was sub-optimal due to constraints of the motor system.

CONCLUSION

In summary, we have argued here that movement under risk is formally equivalent to decision making under risk and uncertainty. Yet, many aspects concerning presentation, execution, completion and repetition of the task, the feedback about possible errors, as well as the implementation of the pay-off rule differ between the two approaches. Thus, cognitive, perceptual and motor constraints come into play when deciding on an action and subsequently planning and executing a goal-directed movement under risk.

REFERENCES

Allais, M. (1953). Le comportment de l'homme rationnel devant la risque: critique des postulats et axiomes de l'école Américaine. *Econometrica*, *21*, 503-546.

Arnauld, A., & Nicole, P. (1662/1992). *La Logique ou L'Art de Penser*. Paris: Gallimard

Attneave, F. (1953). Psychological probability as a function of experienced frequency. *Journal of Experimental Psychology*, *46*, 81-86.

Bell, D. E., Raiffa, H. & Tversky, A. [Eds.] (1988). *Decision Making: Descriptive, Normative and Prescriptive Interactions*. Cambridge, UK: Cambridge University Press.

Bernoulli, D. (1738/1954). Exposition of a new theory on the measurement of risk, translated by Louise Sommer. *Econometrica*, *22*, 23-36.

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

Kahneman, D., Slovic, P., & Tversky, A., Eds. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press.

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

Kahneman, D., & Tversky, A. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Risk and Uncertainty, 5*, 297-323.

Kahneman, D. & Tversky, A. [Eds.] (2000). *Choices, Values & Frames*. New York: Cambridge University Press.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M. & Combs, B. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory*, *4*, 551-578.

von Neumann, J., & Morgenstern, O. (1947). *Theory of Games and Economic Behavior.* Princeton, NJ: Princeton University Press.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and tradeoffs in motor response. *Spatial Vision*, *16*, 255-275.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and rapid, goal-directed movements. *Journal of the Optical Society A*, *20*, 1419-1433.

Trommershäuser, J., Gepshtein, S., Maloney, L.T., Landy, M. S., and Banks, M. S. Optimal compensation for changes in task-relevant movement variability, under review.

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.