

Decision making, movement planning and statistical decision theory

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

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

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

We discuss behavioral studies directed at understanding how probability information is represented in motor and economic tasks. By formulating the behavioral tasks in the language of statistical decision theory, we can compare performance in equivalent tasks in different domains. Subjects in traditional economic decision-making tasks often misrepresent the probability of rare events and typically fail to maximize expected gain. By contrast, subjects in mathematically equivalent movement tasks often choose movement strategies that come close to maximizing expected gain. We discuss the implications of these different outcomes, noting the evident differences between the source of uncertainty and how information about uncertainty is acquired in motor and economic tasks.

Risky decisions and movement planning

Uncertainty plays a fundamental part in perception, cognition and motor control and a wide variety of biological tasks can be formulated in statistical terms. How the organism combines sensory information from many different sources ('cues') is currently an active area of research. Several groups have proposed [1,2] that perceptual estimation of properties of the environment can be framed within Bayesian decision theory, a special case of statistical decision theory [3]. Here, we show that framing behavioral tasks in the language of statistical decision theory enables a comparison of performance between motor tasks and decision making under risk.

Much research concerning decision making seeks to understand how subjects choose between discrete plans of action that have economic consequences [4]. A subject might be given a choice between a 10% chance of winning \$5000 (and otherwise winning nothing) and a 95% chance of winning \$300 (and otherwise winning nothing). These choices can be written in compact form as lotteries (L):

$$L_1 = [0.1, \$5000; 0.9, \$0] \text{ and } L_2 = [0.95, \$300; 0.05, \$0]$$

If subjects are given the probabilities then they are making 'decisions under risk', if not, they are making 'decisions under uncertainty' [5]. Here, we are concerned primarily with the former.

Of course, most subjects would prefer to receive \$5000 rather than \$300, or to receive \$300 rather than \$0. The key difficulty in making such decisions is that no plan of

action (lottery) available to the subject guarantees a specific outcome.

Here, we review recent experimental work in movement planning [6–9] in which humans perform speeded movements towards displays with regions that, if touched, lead to monetary rewards and penalties (Box 1). Our work shows that humans do very well in making these complex decisions in motor form. This outcome is particularly surprising because humans typically do not do well in equivalent economic decision-making tasks, as we describe next.

Sub-optimal economic decisions

Human performance in decision making under risk is markedly sub-optimal and fraught with cognitive biases [4] that result in serious deficits in performance. Patterned deviations from maximizing expected gain include a tendency to frame outcomes in terms of losses and gains with an exaggerated aversion to losses [10] and a tendency to exaggerate the weight given to low-probability outcomes [11,12]. The latter property parallels the human tendency to overestimate the relative frequencies of rare events [13,14]. This exaggeration of the frequency of low-frequency events is observed in many, but not all, decision-making studies [15]. These behaviors are typically modeled by Prospect Theory by introducing a probability weighting function and by assuming that subjects maximize a trade-off between losses and gains [10,12].

Motor tasks equivalent to decision making under risk

Recent work in motor control [9] formulates movement planning in terms of statistical decision theory, effectively converting the problem of movement planning to a decision among lotteries that is mathematically equivalent to decision making under risk. We can compare performance in economic decision-making tasks with performance in equivalent motor tasks and also study how organisms represent value and uncertainty and make decisions in very different domains [16–22].

In Figure 1a, we illustrate the task and show one of the target-penalty configurations used in Ref. [6]. The rules of the task are simple. The configuration appears on a display screen a short distance in front of the subject. The subject must reach out and hit somewhere on the display screen within a limited period of time. The subject knows that hits within the green circle result in a monetary payoff but hits within the red result in a loss of money. The amounts vary with experimental condition but in the example in Figure 1a they are 2.5 cents and 12.5 cents, respectively.

Corresponding author: Trommershäuser, J. (julia.trommershaeuser@psychol.uni-giessen.de).

Box 1. Constructing motor lotteries

In the main text, we describe a movement task equivalent to decision making under risk. On each trial in the task, subjects have to reach out and touch a computer screen within a limited period of time (e.g. 700 ms). Hits inside a green target region, displayed on a computer screen, yield a gain of 2.5 cents; accidental hits inside a nearby red penalty region incur losses of 12.5 cents. Movements that do not reach the screen within the time limit are heavily penalized (after training, they almost never occur). The subject is not completely under control of the movement outcome due to the short time limit [43]. Figure 1 shows simulated outcomes of attempts to hit the target center.

A movement that reaches the screen within the time limit can end in one of four possible regions: penalty only (region R_1 , gain $G_1 = -12.5$), target and penalty overlap (region R_2 , gain $G_2 = -10$), target only (region R_3 , gain $G_3 = 2.5$) or neither/background (region R_4 , gain $G_4 = 0$). On evaluating movement plans in this task, visuo-motor plans that lead to a touch on the screen within the time limit differ only to the extent that they affect the probability $P_s(R_i)$ of hitting each of the four regions R_i , where $i = 1, \dots, 4$. The combination of event

probabilities $P_s(R_i)$ resulting from a particular visuo-motor plan (aim point) s and associated gains G_i form a lottery:

$$L(s) = [P_s(R_1), G_1; P_s(R_2), G_2; P_s(R_3), G_3; P_s(R_4), G_4] \quad \text{[Equation I]}$$

An alternative visuo-motor plan s' corresponds to a second lottery:

$$L(s') = [P_{s'}(R_1), G_1; P_{s'}(R_2), G_2; P_{s'}(R_3), G_3; P_{s'}(R_4), G_4] \quad \text{[Equation II]}$$

Each lottery corresponds to an aim point. The lottery corresponding to the aim point in Figure 1a has an expected gain of -2.8 cents per trial (on average, the subject loses money), the expected gain associated with the aim point in Figure 1b is 0.78 cents per trial. Obviously, the aim point in Figure 1b offers higher expected gain. In planning movement in this task, subjects effectively choose between not just these two aim points but an infinite number of aim points (lotteries). They are engaged in a continuous decision-making task of extraordinary complexity, and it is a task that is performed every time they move.

Figure 1c is a plot the expected gain associated with every possible aim point with four of them highlighted. The aim point corresponding to the yellow diamond maximizes expected gain for the subject in this task.

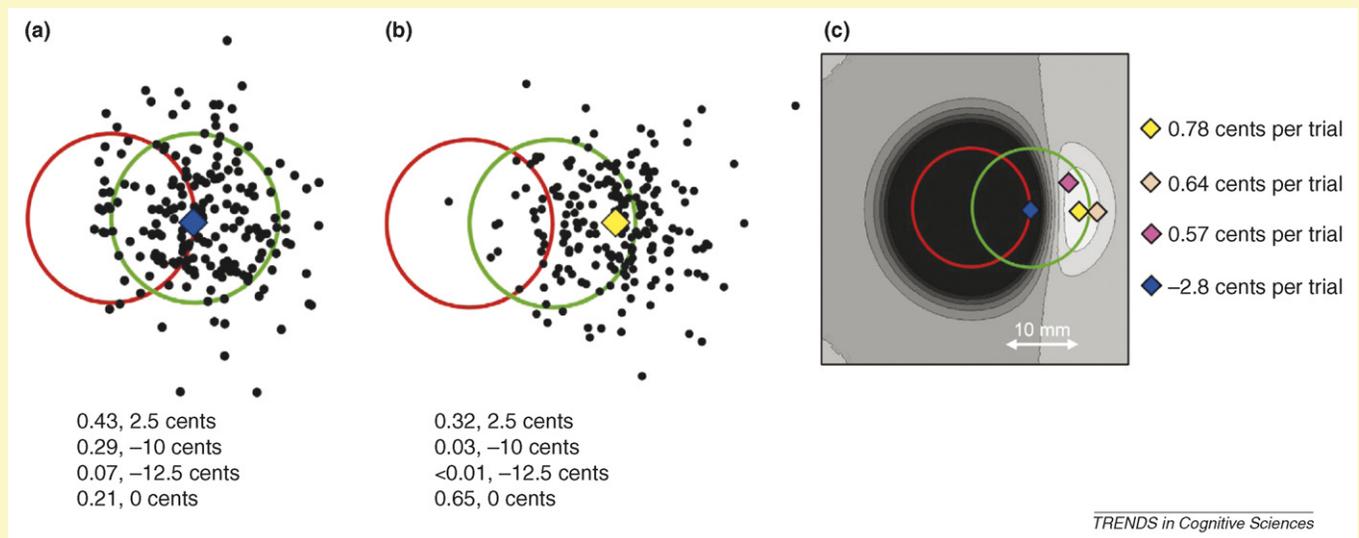


Figure 1. Equivalence of a movement task and decision making under risk. Subjects must touch a computer screen within a limited period of time (e.g. 700 ms). Subjects can win 2.5 cents by hitting inside the green circle, lose 12.5 cents by hitting inside the red circle, lose 10 cents by hitting where the green and red circle overlap or win nothing by hitting outside the stimulus configuration. Each possible aim point on the computer screen corresponds to a lottery. (a) Expected gain for a subject aiming at the center of the green target (aim point indicated by the blue diamond). Black points indicate simulated end points for a representative subject (with 5.6 mm end-point standard deviation); target and penalty circles have radii of 9 mm. This motor strategy yields an expected loss of 2.8 cents per trial. The numbers shown below the target configuration describe the lottery corresponding to this aim point (i.e. the probabilities for hitting inside each region and the associated gain). (b) Expected gain for a subject with the same motor uncertainty as in (a). Here, we simulate the same subject aiming towards the right of the target center (yellow diamond) to avoid accidental hits inside the penalty circle. This strategy results in an expected gain of 0.78 cents per trial and corresponds to the strategy (aim point) that maximizes expected gain. (c) Each possible aim point corresponds to a lottery and has a corresponding expected gain, shown by the grayscale background with four particular aim points highlighted.

TRENDS in Cognitive Sciences

Within the short time limit for movement execution, the movement cannot be completely controlled: even if the subject aims at the center of the green circle there is a real chance of missing it. And, if the subject aims too close to the center of the green circle, there is a risk of hitting inside the red. So, where should the subject aim?

In Box 1, we show how to interpret the subject's choice of aim point as a choice among lotteries and how to determine the aim point that maximizes expected gain. An economist would use the term 'utility', whereas a researcher in motor control would opt for 'cost' or 'biological cost', and a statistician would use 'loss' or 'loss function'. Adopting a field-neutral term, we refer to rewards and penalties associated with outcomes as 'gains', whether positive or negative.

We were surprised to discover that, in this decision task in motor form, participants typically chose visuo-motor plans that came close to maximizing expected gain. Figure 1b is a plot of the subjects' displacement in the horizontal direction away from the center of the green circle versus the displacement that would maximize expected gain, combining data across several experimental conditions varying the penalty amount and distance between target and penalty circles [6,7].

Additional research has extended this conclusion to tasks that involve precise timing and trade-off between movement time and reward [23,24] and to tasks involving rapid choices between possible movement targets [9]. In most cases, human subjects choose strategies that come close to maximizing expected gain in motor tasks with

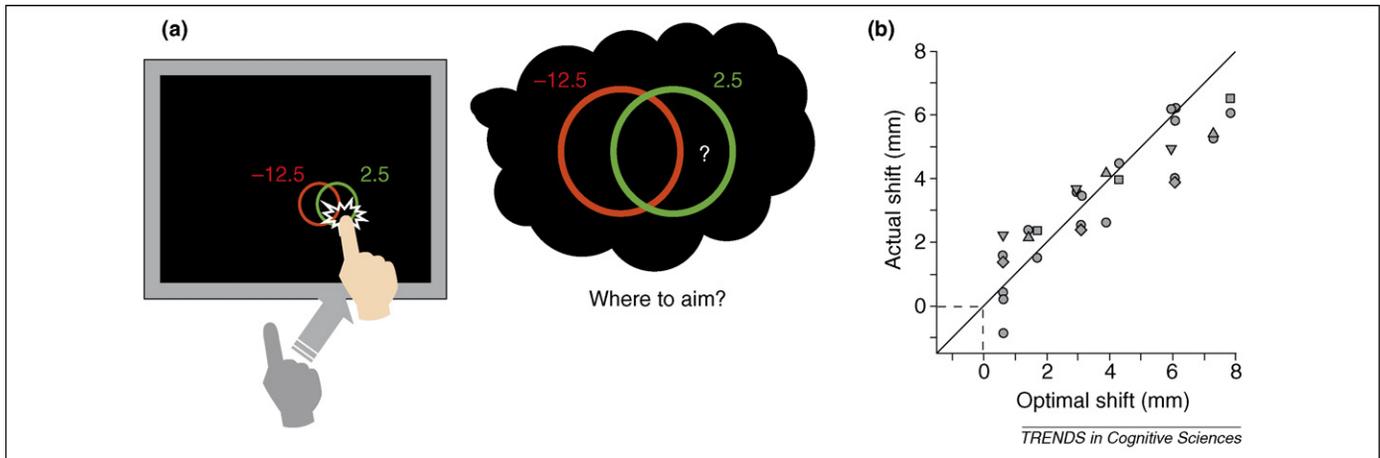


Figure 1. Reaching under risk. (a) An example of a stimulus configuration presented on a display screen. The subject must rapidly reach out and touch the screen. If the screen is hit within the green circle, 2.5 cents are awarded. If hit within the red circle, there is a penalty of 12.5 cents. The circles are small (9 mm radius) and the subject cannot completely control this rapid movement because timing constraints force it to be rapid. In **Box 1** we explain what the subject should do to maximize winnings. (b) A comparison of subjects' performance to the performance that would maximize expected gain. The shift of subjects' mean movement end points from the center of the green target region is plotted as a function of the shift of mean movement end point that would maximize expected gain for five different subjects (indicated by five different symbols) and the six different target-penalty configurations shown in **Figure 2a**. Replotted, with permission, from **Figure 5a** in Ref. [6].

changing stochastic variability [8,25] or combining noisy sensory input with prior information [26–29].

These results have implications for understanding movement planning and motor control. Typical computational approaches to modeling movement-planning take the form of an optimization problem, in which the cost function to be minimized is biomechanical and the optimization goal is to minimize some measure of stress on the muscles and joints. These models differ primarily in the choice of the cost function. Possible biomechanical cost functions include measures of joint mobility [30,31], muscle-tension changes [32], mean-squared rate of change of acceleration [33], mean torque change [34], total energy expenditure [35] and peak work [36]. These biomechanical models have successfully been applied to model the nearly straight paths and bell-shaped velocity profiles of arm movements and also capture the human ability to adapt to forces applied during movement execution [37]. We emphasize that these models cannot be used to predict subjects' performance in our movement-decision tasks in which performance also depends on externally imposed rewards and penalties. Moreover, subjects came close to maximizing expected gain with arbitrary, novel penalties and rewards imposed on outcomes by the experimenter.

Subjects do not always come close to maximizing expected gain in movement planning; for example when the number of penalty and reward regions is increased [38] and when the reward or penalty received is stochastic rather than determined by the outcome of the subject's movement [39]. Furthermore, when the penalty is so high that the aim point that maximizes expected gain lies outside of the target, results indicate that subjects prefer not to aim outside of the target that they are trying to hit [8]. Thus, although there is a collection of motor tasks (as described earlier) in which performance is remarkably good, we cannot simply claim that performance in any task with a speeded motor response will come close to maximizing expected gain. Further work is needed to delimit the range of movement-planning tasks in which subjects do well.

One evident question is: are subjects maximizing expected gain gradually by a process of trial and error?

Learning probabilities versus practicing the task

We were surprised to learn that subjects do not show trends that are consistent with a gradual approach to maximizing expected gain. The design of our studies [6–8] and related work [23] had a peculiar structure. Before the 'decision-making' phase of the experiment, subjects practiced the speeded motor task extensively by repeatedly touching single circular targets. During this initial training period, the experimenter monitored their motor performance until it stabilized and the experimenter could measure the residual motor variability of each subject.

Only after training did subjects learn about the gains and losses assigned to each region in the experimental condition. They were not explicitly told to take into account the spatial locations of reward and penalty regions and the magnitude of penalty and reward, but their highly efficient performance indicates that they did so from the first trial in which rewards and penalties were specified.

To summarize, in these experiments subjects were first trained to be 'motor experts' in a simple task in which they were instructed to touch targets on the screen. Only then were they given a task involving trade-offs between rewards and penalties. There were no obvious trends in the aim points of subjects [6,7] that would indicate that they were modifying their decision-making strategy as they gained experience with the decision-making task (**Figure 2a**).

To see how unusual this outcome is, consider applying a simple reinforcement-learning model according to which the aim point is adjusted gradually in response to rewards and penalties incurred [40–42]. In the absence of any reward or penalty, a learning model based on reward and penalty would predict that the subject should aim at the center of the green circle, just as in the training trials. The subject would then gradually alter the aim point in response to rewards and penalties incurred until the final aim point maximized expected gain (**Figure 2b**).

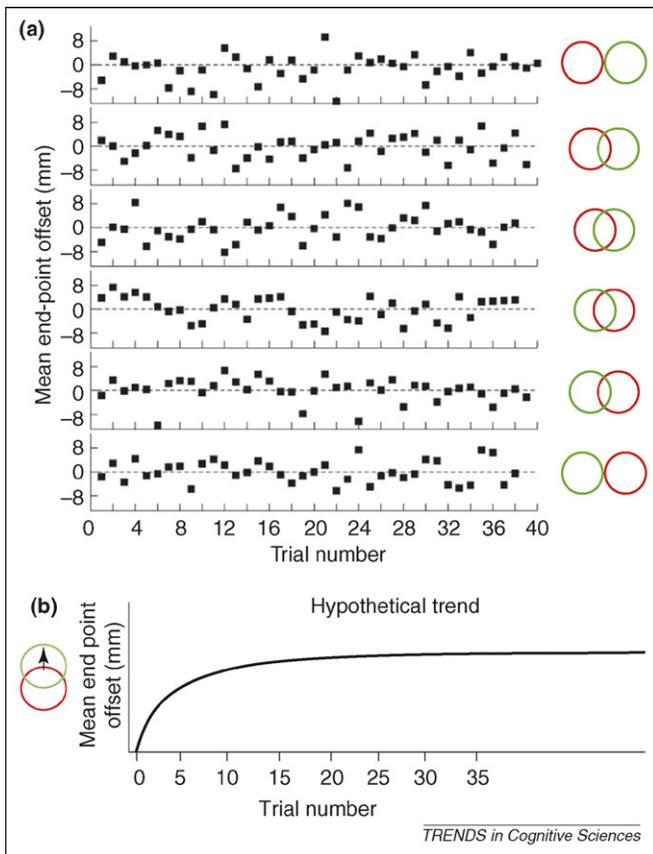


Figure 2. Absence of learning in reaching under risk. **(a)** Trial-by-trial deviation of movement end point (in the horizontal direction) from the mean movement end point in that condition as a function of trial number after introduction of rewards and penalties (reward: 2.5 cents; penalty -12.5 cents); the six different lines correspond to the six different spatial conditions of target and penalty off-set as shown on the right. Data replotted, with permission, from Figure 7 in Ref. [6]. **(b)** Trend of a hypothetical simple learning model in which a subject changes motor strategy gradually in response to rewards and penalties incurred. The subject initially aims at the center of the green circle. Before the subject's first trial in the decision-making phase of the experiment, the subject is instructed that red circles carry penalties and green circles carry rewards. Under this model, subjects would approach the aim point to maximize expected gain by slowly shifting the aim point away from the center of the green circle until the winnings match the maximum expected gain. However, the data shown in (a) do not support this learning model.

However, examination of the initial trials of the decision phase of the experiment indicates that subjects immediately changed their aim point from that used in training to that necessary to trade off the probabilities of hitting the reward and penalty regions (Figure 2a). This apparent lack of learning is of great interest in that it indicates that, although subjects certainly learned to carry out the motor task in the training phases of these experiments, and learned their own motor uncertainty, they seemed not to need further experience with the decision-making task to perform as well as they did, applying the knowledge of that motor uncertainty to new situations. The trends in performance found by repetition of economic-decision tasks seem absent in equivalent movement-planning tasks.

The contrast between success in 'movement planning under risk' and decision making under risk is heightened by the realization that, in decision making under risk, subjects are told the exact probabilities of outcomes and, thus, have perfect knowledge of how their choice changes the probability of attaining each outcome. The knowledge of probabilities in equivalent motor tasks is never communicated explicitly and, thus, can equal but never exceed the

knowledge available under decision making under risk. Yet the lack of learning in these motor tasks indicates that humans are able to estimate the probabilities of each outcome associated with any given aim point because of motor uncertainty and make use of this knowledge to improve their performance [8]. There is mounting evidence that decision makers behave differently if knowledge of probabilities is gained through 'experience' (Box 2). Our

Box 2. Decisions from experience

There are several factors that might have contributed to the remarkable performance of subjects in movement planning under risk [6–9,44]. In these experiments, the subject makes a long series of choices and, over the course of the experiment, their accumulated winnings increase. By contrast, subjects in economic decision-making experiments typically make a single 'one-shot' choice, choosing from a small set of lotteries. Indeed, when economic decision makers are faced with a series of decisions, they tend to move closer to maximum expected gain (e.g. Ref. [45]; and 'the house money effect' [46]).

Recent work indicates that subjects who are allowed to simulate a decision task learn from their experience [47], and, together with the studies just cited, it is likely that decision making improves with repetition. However, in the motor tasks discussed here, the learning phase does not involve explicit probabilities, values or trade-offs between risk and reward. In the experimental phase (see Figure 2 in the main text), they show no evidence of learning. This outcome indicates that they can explicitly transfer experience with motor uncertainty to the decision task (Figure 1), computing probabilities and planning movements on demand. Although subjects probably learn from experience in these motor tasks, experience does not involve simple practice or simulation of the actual decision task.

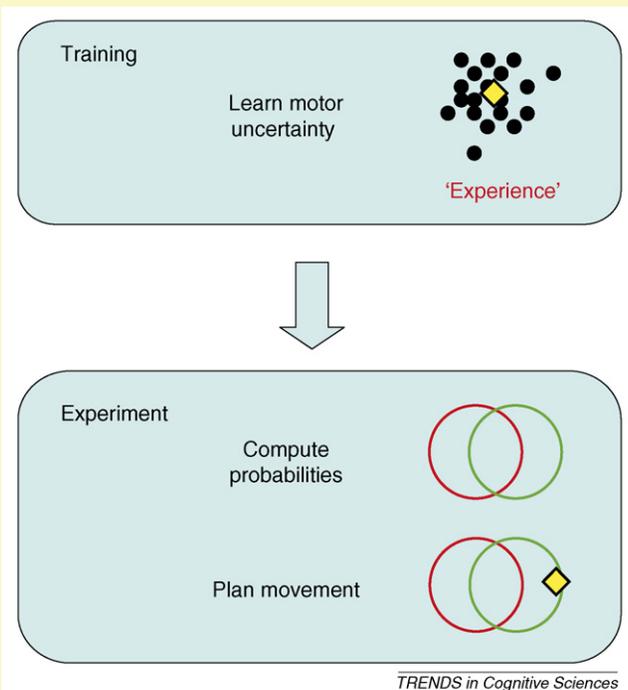


Figure 1. Motor decisions from experience. In the learning phase of the experiment, subjects learn to hit targets. Their performance improves until their movement variability has reached a plateau. During training, they have the opportunity to learn their own motor uncertainty but nothing about the training task requires that they do so. In the experimental phase, subjects plan movements that trade off the risk of incurring penalties against the possible reward of hitting targets. They show little evidence of learning and perform well in the task. This indicates that they can convert what they learned in the training phase into the information needed to plan effective movements under risk: the equivalent of estimating the probabilities of the various outcomes associated with any proposed aim point, followed by a computation of expected gain.

results add a new dimension to what kinds of experience lead to enhanced decision making.

There is growing interest in analyzing brain activity in response to manipulations of various components of

decision making under risk or uncertainty in human subjects (for more extensive reviews, see Refs [16,17,19,20]). The work described here effectively opens a second window for neural processing of uncertainty and value by allowing

Box 3. Statistical decision theory and sensory-motor control

Statistical decision theory [48] is a remarkably general framework for modeling tasks in cognition, perception and planning of movement [3]. In its simplest forms, it is the mathematical basis for signal-detection theory and common models of optimal visual classification [49]. The models of simple movement tasks considered here are examples of its application. Figure 1a illustrates its application to a more complex movement task that involves both visual and motor uncertainty. A dinner guest intends to pick up a salt shaker at the center of the table with his right hand. We follow this movement from initial planning to eventual social disaster (Figure 1b) or success (Figure 1c).

One possible plan of action is schematized as a solid line, sketching out the path of the hand that the guest plans to take. An actual movement plan would specify joint movements throughout the reach. His sensory information should take into account uncertainty in his estimates of object location in addition to his accuracy in movement. If his sensory information is poor under candlelight, he might do well to choose a path that gives the wine glass a wide berth and proceed slowly. But, if he moves too slowly, he will never get through his meal. The potential costs and benefits are measured in units of disgrace, esteem and dry-cleaning charges. Statistical decision theory enables us to determine the best possible choice of movement plan (i.e. the one that maximizes expected gain).

In detail, a movement strategy is a mapping from sensory input V to a movement plan $s(V)$ (Figure II). The expected gain associated with the choice of strategy $s(V)$ is given by:

$$EG(s) = \int \int \int g(t, w) p_T(t|s(v)) p_v(v|w) p_w(w) dv dt dw$$

[Equation III]

where w is the random state of the world (i.e. positions of arm, salt shaker, wine glass and so on) with prior distribution $p_w(w)$ based on past sensory information and knowledge of how a table is laid out, V is

current sensory information about the state of the world with likelihood distribution $p_v(v|w)$ and T is the stochastic movement trajectory resulting from the executed movement plan $s_T(V)$. The term $g(t, w)$ specifies the gain resulting from an actual trajectory t in the actual state of the world w . In the example given, it includes costs incurred by hitting objects while reaching through the dinner scene and possible rewards for successfully grasping the salt shaker. Equation III determines the movement strategy that maximizes expected gain.

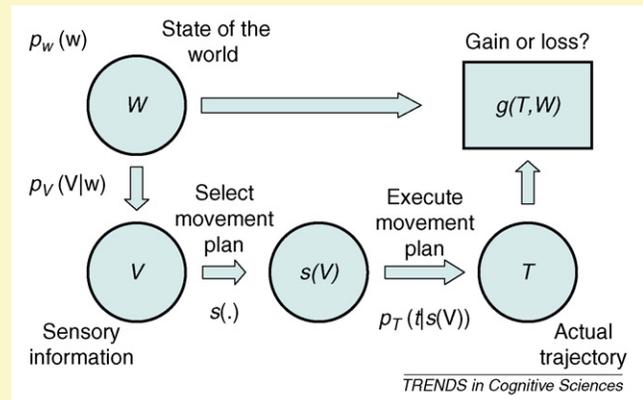


Figure II. Application of statistical decision theory to complex visuo-motor tasks. The goal is a mapping from sensory input V to a movement plan $s(V)$. Gains and losses $g(t, w)$ are determined by the actual trajectory t executed in the actual state of the world w . The movement plan that maximizes expected gain depends on both visual uncertainty and motor uncertainty. (Here, we follow the convention that random variables are in upper case (e.g. X), whereas the corresponding specific values that those variables can take on are in lower-case [e.g. $p(x)$]).

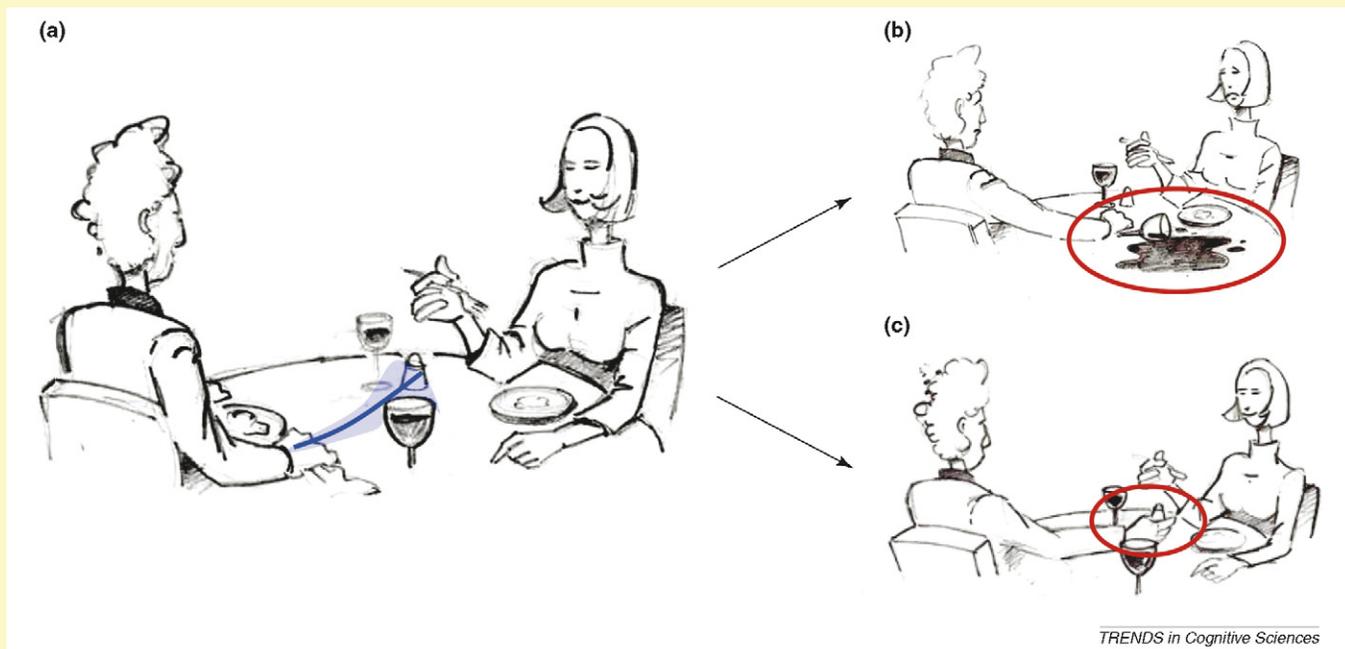


Figure I. Example of applying statistical decision theory to modeling goal-directed movement under visual and motor uncertainty. (a) A dinner guest intends to pick up the salt shaker at the center of the table with his right hand. An intended trajectory is shown along with a 'confidence interval' to indicate the range of other trajectories that might occur. (b) The actual executed movement might deviate from the intended and, instead of grasping the salt shaker, the guest might accidentally knock over his full wine glass. (c) If executed successfully, the dinner guest will pick up the salt shaker without experiencing social disaster. (Drawings by Andreas Olsson).

one to present exactly the same decision problems in different guises.

Statistical decision theory: future directions

The motor task we have considered is simple: a reaching movement to touch a target. Even this simple motor task corresponds to complicated choices among lotteries. We close by illustrating that the underlying statistical framework, statistical decision theory, can be used to model complex movement tasks shaped by externally imposed rewards and penalties in which visual uncertainty can play a larger part (Box 3).

By using the methods described here, visuo-motor and economic decision-making tasks can be translated into a common mathematical language. We can frame movement in economic terms or translate economic tasks into equivalent visuo-motor tasks. Given the societal consequences associated with failures of decision making in economic, military, medical and legal contexts, it is worth investigating decision tasks in domains in which humans seem to do very well.

Acknowledgements

We thank Marisa Carrasco, Nathaniel Daw, Karl Gegenfurtner and Paul Glimcher for helpful discussion and Andreas Olsson for the drawings. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, Emmy-Noether-Programm, grant TR 528/1–2; 1–3) and NIH EY08266.

References

- Knill, D.C. *et al.* (1996) Introduction: a Bayesian formulation of visual perception. In *Perception as Bayesian Inference* (Knill, D.C. and Richards, W., eds), pp. 1–21, Cambridge University Press
- Landy, M.S. *et al.* (1995) Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Res.* 35, 389–412
- Maloney, L.T. (2002) Statistical decision theory and biological vision. In *Perception and the Physical World: Psychological and Philosophical Issues in Perception* (Heyer, D. and Mausfeld, R., eds), pp. 145–189, Wiley
- Kahneman, D. and Tversky, A. (2000) *Choices, Values, and Frames*, Cambridge University Press
- Knight, F.H. (1921) The Place of Profit and Uncertainty in Economic Theory. In: *Risk, Uncertainty and Profit*, Houghton Mifflin
- Trommershäuser, J. *et al.* (2003) Statistical decision theory and the selection of rapid, goal-directed movements. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 20, 1419–1433
- Trommershäuser, J. *et al.* (2003) Statistical decision theory and trade-offs in the control of motor response. *Spat. Vis.* 16, 255–275
- Trommershäuser, J. *et al.* (2005) Optimal compensation for changes in task-relevant movement variability. *J. Neurosci.* 25, 7169–7178
- Trommershäuser, J. *et al.* (2006) Humans rapidly estimate expected gain in movement planning. *Psychol. Sci.* 17, 981–988
- Kahneman, D. and Tversky, A. (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47, 263–291
- Allais, M. (1953) Le comportement de l'homme rationnel devant la risque: critique des postulats et axiomes de l'école Américaine. *Econometrica* 21, 503–546; translated and reprinted in Allais, M. and Hagen, O. (Eds.) (1979) Expected Utility Hypotheses and the Allais Paradox. Dordrecht: D. Reidel; part I
- Tversky, A. and Kahneman, D. (1992) Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297–323
- Attneave, F. (1953) Psychological probability as a function of experienced frequency. *J. Exp. Psychol.* 46, 81–86
- Lichtenstein, S. *et al.* (1978) Judged frequency of lethal events. *J. Exp. Psychol. Hum. Learn.* 4, 551–578
- Sedlmeier, P. *et al.* (1998) Are judgments of the positional frequencies of letters systematically biased due to availability? *J. Exp. Psychol. Learn. Mem. Cog.* 24, 754–770
- Daw, N.D. and Doya, K. (2006) The computational neurobiology of learning and reward. *Curr. Opin. Neurobiol.* 16, 199–204
- Glimcher, P.W. and Rustichini, A. (2004) Neuroeconomics: the consilience of brain and decision. *Science* 306, 447–452
- Gold, J.I. and Shadlen, M.N. (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308
- Montague, P.R. *et al.* (2006) Imaging valuation models in human choice. *Annu. Rev. Neurosci.* 29, 417–448
- O'Doherty, J.P. (2004) Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Curr. Opin. Neurobiol.* 14, 769–776
- Sugrue, L.P. *et al.* (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787
- Sugrue, L.P. *et al.* (2005) Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* 6, 363–375
- Dean, M. *et al.* (2007) Trading off speed and accuracy in rapid, goal-directed movements. *J. Vis.* 7 (5), 10, 1–12
- Hudson, T.E. *et al.* Optimal compensation for temporal uncertainty in movement planning. *PLoS Comp. Biol.* (in press)
- Baddeley, R.J. *et al.* (2003) System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *J. Neurosci.* 23, 3066–3075
- Körding, K.P. and Wolpert, D.M. (2004) Bayesian integration in sensorimotor learning. *Nature* 427, 244–247
- Schlicht, E.J. and Schrater, P.R. (2007) Impact of coordinate transformation uncertainty on human sensorimotor control. *J. Neurophysiol.* 97, 4203–4214
- Tassinari, H. *et al.* (2006) Combining priors and noisy visual cues in a rapid pointing task. *J. Neurosci.* 26, 10154–10163
- Vaziri, S. *et al.* (2006) Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. *J. Neurosci.* 26, 4188–4197
- Kaminski, T. and Gentile, A.M. (1986) Joint control strategies and hand trajectories in multijoint pointing movements. *J. Mot. Behav.* 18, 261–278
- Soechting, J.F. and Lacquaniti, F. (1981) Invariant characteristics of a pointing movement in man. *J. Neurosci.* 1, 710–720
- Dornay, M. *et al.* (1996) Minimum muscle-tension change trajectories predicted by using a 17-muscle model of the monkey's arm. *J. Mot. Behav.* 28, 83–100
- Flash, T. and Hogan, N. (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* 5, 1688–1703
- Uno, Y. *et al.* (1989) Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. *Biol. Cybern.* 61, 89–101
- Alexander, R.M. (1997) A minimum energy cost hypothesis for human arm trajectories. *Biol. Cybern.* 76, 97–105
- Soechting, J.F. *et al.* (1995) Moving effortlessly in three dimensions: does Donders' law apply to arm movement? *J. Neurosci.* 15, 6271–6280
- Burdet, E. *et al.* (2001) The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414, 446–449
- Wu, S.W. *et al.* (2006) Limits to human movement planning in tasks with asymmetric gain landscapes. *J. Vis.* 6, 53–63
- Maloney, L.T. *et al.* (2007) Questions without words: a comparison between decision making under risk and movement planning under risk. In *Integrated Models of Cognitive Systems* (Gray, W., ed.), pp. 297–315, Oxford University Press
- Daw, N.D. *et al.* (2006) Cortical substrates for exploratory decisions in humans. *Nature* 441, 876–879
- Dayan, P. and Balleine, B.W. (2002) Reward, motivation, and reinforcement learning. *Neuron* 36, 285–298
- Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning, MIT Press
- Meyer, D.E. *et al.* (1988) Optimality in human motor performance: ideal control of rapid aimed movements. *Psychol. Rev.* 95, 340–370
- Seydell, A. *et al.* (2008) Learning stochastic reward distributions in a speeded pointing task. *J. Neurosci.* 28, 4356–4367

- 45 Redelmeier, D.A. and Tversky, A. (1992) On the framing of multiple prospects. *Psychol. Sci.* 3, 191–193
- 46 Thaler, R. and Johnson, E.J. (1990) Gambling with the house money and trying to break even: the effects of prior outcomes on risky choice. *Manage. Sci.* 36, 643–660
- 47 Hertwig, R. *et al.* (2004) Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* 15, 534–539
- 48 Blackwell, D. and Girshick, M.A. (1954) *Theory of Games and Statistical Decisions*, Wiley
- 49 Duda, R.O. *et al.* (2001) *Pattern Classification*, (2nd edn), Wiley