

Questions Without Words

A Comparison Between Decision Making Under Risk and Movement Planning Under Risk

Laurence T. Maloney, Julia Trommershäuser, & Michael S. Landy

If you want answers without words, then ask questions without words.

—Augustine of Hippo (translation, Wills, 2001, p. 139)

We describe speeded movement tasks that are formally equivalent to decision making under risk. In these tasks, subjects attempt to touch reward regions on a touch screen and avoid nearby penalty regions, much as a golfer aims to reach the green while avoiding nearby sand traps. The subject is required to complete the movement within a short time and, like the golfer, cannot completely control the outcome of the planned action. In previous experimental work, we compared human performance to normative (optimal) models of decision making and, in marked contrast to the grossly suboptimal performance of human subjects in decision-making experiments, subjects' performance in these experiments was typically indistinguishable from optimal. We conjecture that the key difference between our tasks and ordinary decision making under risk is the source of uncertainty, implicit or explicit. In the movement tasks, the probability of each possible outcome is implicit in the subject's own motor uncertainty. In classical decision making, probabilities of outcomes are chosen by the experimenter and explicitly communicated to the subject. We present an experimental study testing this conjecture in which we introduced explicit probabilities into the movement task. Subjects knew that some regions (coded by color) were stochastic. If they touched a stochastic reward region, they knew that the chances of receiving the reward were 50 percent and similarly for a stochastic penalty region. Consistent with our conjecture, subjects' optimal performance was disrupted when they were confronted with explicit uncertainty about rewards and penalties.

Many everyday decisions are wordless. We slow down in navigating a narrow doorway. We speed up while crossing the street with an eye to oncoming traffic. We swerve to avoid a colleague on the stairs. These decisions likely depend on several factors, notably expectations of pain or embarrassment, but we would be hard pressed to justify our "choices" in words or explain exactly what those factors are. The range of similar "wordless" decisions is endless and much of our day is taken up with them. Yet, as important as these everyday decisions may be, it is not even clear that we are aware that we have made a particular decision or are aware of the factors that led to it. The reader may even hesitate to classify such wordless decisions with the

kind of decision making we engage in while attempting to fatten a stock portfolio or play out a hand of poker.

Moreover, even if we accept that these wordless decisions are decisions, we cannot be sure that we are at all good at making them. If we did not slow down in approaching the doorway, is there any appreciable chance that we would hit either side of it instead of passing through? Or are we being reckless in not slowing even more? The information needed to make this sort of decision well is, on the one hand, the gains and losses that are possible and, on the other, accurate estimates of the possible discrepancy between how we intend to move and how we actually do move. It will turn out to be the latter sort of information that distinguishes

ordinary decision making from the wordless decision making considered here.

In this chapter, we discuss results from experiments on human movement planning in risky environments in which explicit monetary rewards are assigned to the possible outcomes of a movement. We will show that these conditions create movement tasks that are formally equivalent to decision making under risk—if subjects can anticipate the stochastic uncertainty inherent in their movements. To anticipate our conclusion, we find that they can do so and that their performance in what we term *movement planning under risk* is remarkably good. These tasks form a promising alternative domain in which to study decision making, and they are distinguished by the fact that the uncertainties surrounding possible outcomes are intrinsic to the motor system and, so far as we can judge, difficult to articulate: These tasks are questions without words where stochastic information is available but not explicitly so.

We first review basic results on decision making under risk and then present experimental results concerning movement planning under risk. We end with a description of an experiment in which we ask subjects to carry out a movement task in an environment where the uncertainties surrounding possible outcomes are a combination of implicit motor uncertainty and explicit uncertainties imposed by the experimenter. As we will see, subjects who accurately compensate for their own implicit motor uncertainty fail to compensate correctly for added explicit uncertainty.

Decision Making Under Risk: A Choice Among Lotteries

Imagine that, when you turn this page, you will find a \$1,000 bill waiting for you. It's yours to keep and do with as you like. What alternatives spring to mind? A banquet? Part payment on a Hawaiian vacation? Opera tickets? Each possible outcome of the choice you are about to make is appealing, to some extent, but sadly they are mutually exclusive. You have only one \$1,000 bill, and you can only spend it once. Decision making is difficult in part because, by making one choice, we exclude other, desirable possibilities.

There is a further complication, though, that makes decisions even harder. You can choose how to dispose of your windfall, but you can't be completely sure of the outcome. You can plan to have a superb meal in a famous restaurant—and fall ill. You may schedule a vacation—and spend a week in the rain. Not every opera

performance is brilliant. What you choose in making a decision is rarely a certain outcome but typically a probability distribution across possible outcomes. If the possible outcomes associated with a particular choice are denoted O_1, \dots, O_n , the effect of any decision is to assign a probability p_i to each possible outcome O_i . The result is called a *lottery* and denoted $(p_1, O_1; p_2, O_2; \dots; p_n, O_n)$, where

$$\sum_{i=1}^n p_i = 1.$$

Decision making is, stripped to its essentials, a choice among lotteries. In choosing a plan of action, the decision maker, in effect, selects a particular lottery. In the previous example, we did not assume that the decision maker knows what the probabilities associated with any plan of action are. Decisions without knowledge of the probabilities associated with each outcome are referred to as *decision making under uncertainty*. When decision makers have access to the probabilities induced by each possible plan of action, they are engaged in *decision making under risk*. Our focus here is on the latter sort of decision. In Table 21.1, we enumerate four possible plans of action that assign probabilities to each of four possible monetary outcomes. If the decision maker selects Lottery 1 (L1), for example, there will be an 80% chance of winning \$100 and a 20% chance of losing \$100. In contrast, L3 guarantees a gain of \$50. The key problem for the decision maker is to select among lotteries such as those presented in Table 21.1.

For the decision maker who prefers more money to less, certain of these lotteries dominate others. L3 evidently dominates L4. Comparison of the probabilities

TABLE 21.1 Four Lotteries

	Possible Outcomes			
	\$100	−\$100	\$50	\$0
Probabilities				
Lottery 1	0.8	0.2	0	0
Lottery 2	0.5	0.5	0	0
Lottery 3	0	0	1	0
Lottery 4	0	0	0	1

Four possible monetary outcomes are listed in the first row. The remaining rows specify a lottery with the given probabilities assigned to the corresponding outcomes above. If the decision maker selects Lottery 4 (L4), for example, \$0 will be received with certainty. If the decision maker selects L3, \$50 is received with certainty. With L2, it is equally likely that the decision maker will win or lose \$100.

associated with each outcome in L1 and L2 shows that L1 guarantees a higher probability of gain and a lower probability of loss in every case. L1 dominates L2. That is, the decision maker intent on winning money should never select L2 if L1 is available or L4 if L3 is available. But the choice between L1 and L3 is not obvious: with L1, there is an evident tension between the high probability of winning \$100 with lottery L1 and the 20% chance of losing \$100. With L3, in contrast, the maximum possible gain is \$50, but it is also the minimum: L3 offers \$50 for sure. Which will we pick, given the choice? Which should we pick?

Decision Making Under Risk: Normative Theories

The classical decision-making literature distinguishes between descriptive and normative theories of decision making (Tversky & Kahneman, 1988). A descriptive theory of decision making attempts to predict what choice a decision maker would make if confronted with any collection of lotteries such as Table 21.1. Currently no descriptive theory is widely accepted (Birnbaum, 2004). We will briefly describe previous attempts to develop descriptive theories in the next section.

A normative theory is a rule that orders any set of lotteries from best to worst. We saw that we could order some of the alternatives in Table 21.1 by dominance ($L1 > L2$, $L3 > L4$). Any normative rule allows us to complete this ordering. The lottery highest in the resultant ordering is then the “best” lottery according to the normative rule. The two oldest normative rules are maximum expected value (MEV) and maximum expected utility (MEU).

If the outcomes are numerical (e.g., money), then the expected value of a lottery $L = (p_1, O_1; p_2, O_2; \dots; p_n, O_n)$ is the sum of the values weighted by the corresponding probabilities.

$$EV(L) = \sum_{i=1}^n p_i O_i. \quad (1)$$

The lottery selected by the MEV rule is the lottery with the highest expected value (Arnauld & Nichole, 1662/1992). In Table 21.1, for example, the MEV lottery is L1 with expected value \$60.

Some decision makers may not agree with the recommendation of the MEV rule for Table 21.1 and instead prefer the sure win of L3 (with EV \$50) to L1 with EV \$60. If we were to multiply all of the outcomes by 100 so that

$$L1' = (0.8, \$10,000; 0.2, -\$10,000) \quad (2)$$

and

$$L3' = (1, \$5,000), \quad (3)$$

then almost all will choose L3' with EV \$5,000 rather than L1' with EV \$6,000. Decision makers are often risk averse in this way, especially when confronted with single decisions involving large values.

Daniel Bernoulli (1738/1954) proposed an alternative normative rule based on expected utility, intended to justify risk aversion. If any monetary outcome O_i is assigned a numerical *utility* denoted $U(O_i)$, then we can assign an expected utility (Bernoulli, 1738/1954) to each lottery

$$EU(L) = \sum_{i=1}^n p_i U(O_i + W), \quad (4)$$

where W is the total initial wealth of the decision maker. The decision maker who seeks to maximize expected utility chooses the action whose corresponding lottery has the MEU. When outcomes are numeric and the utility function is a linear transformation with positive slope, MEU includes MEV as a special case.

Bernoulli proposed MEU as both a normative and a descriptive theory, intended to explain risk aversion. If the utility function $U(O)$ is concave (the utility of each successive dollar is less than the preceding dollar), then risk aversion is a consequence of MEU. In addition, the concept of utility has the great advantage that we can potentially describe decision making among outcomes that are nonnumerical by assigning them numerical utilities.

Decision Making Under Risk: Descriptive Theories

Research in human decision making under risk during the past forty years is a catalog of the many, patterned failures of normative theories, notably MEU, to explain the decisions humans actually make (Bell, Raiffa, & Tversky, 1988; Kahneman, Slovic, & Tversky, 1982; Kahneman & Tversky, 2000). These failures include a tendency to frame outcomes in terms of losses and gains with an aversion to losses (Kahneman & Tversky, 1979) and to exaggerate small probabilities (Allais, 1953; Attneave, 1953; Lichtenstein, Slovic, Fischhoff, Layman & Coombs, 1978; Tversky & Kahneman, 1992) as illustrated by the two examples in Figure 21.1.

Figure 21.1A shows data from Attneave (1953). In this study, subjects were asked to estimate the frequency of occurrence of letters in English text. It is evident that

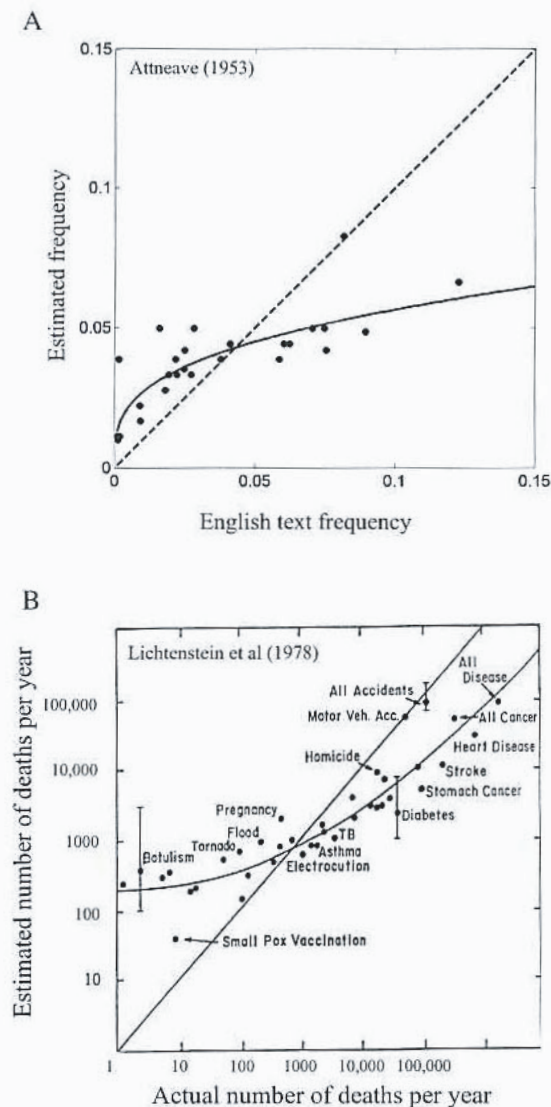


FIGURE 21.1 Estimated frequencies versus true frequencies. (A) A plot of estimated frequency of occurrence of letters in English text versus actual frequency (redrawn from Attneave, 1953). Subjects overestimated the frequency of letters that occur rarely relative to the frequency of letters that occur frequently. (B) A plot of estimated frequencies of lethal events versus actual frequencies (redrawn from Lichtenstein et al., 1978). Subjects markedly overestimated the frequency of rare events relative to more frequently occurring events.

they overestimated the frequency of letters that occur rarely compared with those that occur more frequently. In Figure 21.1B, we replot data from Lichtenstein et al. (1978). These data are the estimated frequencies of lethal events plotted versus the true frequencies. The data span

several orders of magnitude and the same pattern of exaggeration of small probabilities emerges. Estimates of the smaller probabilities are occasionally off by factors of 100 or more.

Studies that examine subjects' use of probabilities in decision making under risk draw similar conclusions (Gonzalez & Wu, 1999; Tversky & Kahneman, 1992; Wu & Gonzalez, 1998, 1999): subjects' use of probability or frequency information is markedly distorted and leads to suboptimal decisions, a phenomenon we will return to below.

There are other well-documented deviations from MEU 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 MEU rule in decision making tasks and that, in the terminology of Kahneman and Tversky, human decision makers are given to "cognitive illusions."

Movement Planning Under Risk: A Different Kind of Decision

The typical tasks found in the literature on decision making under risk are paper-and-pencil choices with no time limit on responding. Full information about outcomes and probabilities is explicitly specified, and there are usually only two or three possible choices (lotteries). In these tasks, both probabilities and values are selected by the experimenter and are communicated to the subject through numeric representations or by simple graphical devices. It is rare that any justification is given for why a particular probability should be attached to a particular outcome. These sorts of decisions are very far from the everyday, wordless decisions discussed in the introduction.

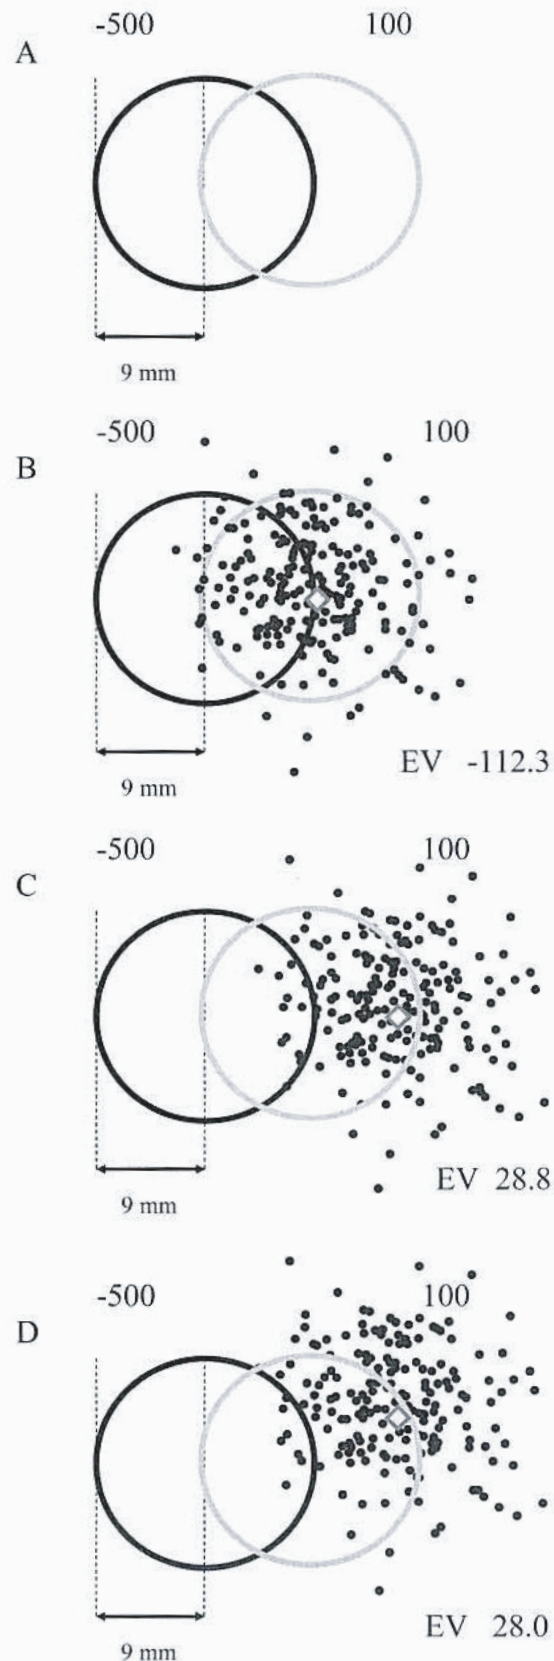
Here, we introduce a movement planning task that is formally equivalent to decision making under risk, and we describe how subjects perform in these sorts of tasks. As will become evident, one major difference between these tasks and more traditional examples of decision making under risk is that the source of uncertainty that determines the probabilities in each lottery is the subject's own motor variability. No explicit specification of probability or frequency is ever given to the subject. Information about probability or frequency is implicit in the task itself.

Trommershäuser, Maloney, and Landy (2003a, 2003b) asked subjects to make a rapid pointing move-

ment and touch a stimulus configuration on a touch screen with their right index finger. The touch screen was vertical, directly in front of the subject. A typical stimulus configuration from Trommershäuser et al. (2003a) is shown in Figure 21.2A. This stimulus configuration or its mirror image was presented at a random location within a specified target area on the screen. A trial started with a fixation cross. The subject was required to move the index finger of the right hand to the starting position (marked on the space bar of a keyboard). The trial began when the space bar was pressed. The subject was required to stay at this starting position until after the stimulus configuration appeared or the trial was aborted. Next, a blue frame was displayed delimiting the area within which the target could appear and preparing the subject to move. Five hundred milliseconds later the target and penalty circles were displayed. Subjects were required to touch the screen within 700 ms of the display of the circles or they would incur a “timeout” penalty of 700 points. If a subject hit within the green¹ target in time, 100 points were earned. If the subject accidentally hit within an overlapping red circle, points were lost. If the subject hit in the region common to the two circles, both the reward associated with the green and the penalty associated with the red were incurred. If the subject hit the screen within the time limit, but missed both circles, no points were awarded.

The penalty associated with the red circle varied with experimental condition. In one condition, the penalty associated with the red circle was zero (i.e., there were no consequences for hitting within the red circle). At the other extreme, the penalty for hitting within the red circle was 500 points, five times greater than the reward for hitting the green circle. Subjects were always aware of the current penalty associated with the red circle.

FIGURE 21.2 Possible movement strategies. (A) Stimulus configuration from Trommershäuser et al. (2003a). The reward circle is gray; the penalty circle is black. The reward and penalty associated with hitting within each region are shown. (B) The black dots represent simulated movement end points for a subject with motor variability $\sigma = 4.83$ mm who has adopted a movement strategy with mean end point marked by the diamond, at the center of the reward region. The expected value (see text) is shown and is negative. This maximizes the chances of hitting the reward region. (C) The mean end point is shifted horizontally by 7 mm and the expected value is now positive. (D) A shift upward from the axis of symmetry. The expected value is also positive but less than that of B.



They also knew that the summed points they earned over the course of the experiment would be converted into a proportional monetary bonus. At the end of every trial, the subject saw a summary of that trial (whether the subject had timed out, hit within the green and/or red, and how much was won or lost on that trial). The cumulative total of how much the subject had won or lost so far in the experiment was also displayed.

If subjects could perfectly control their movements, they could simply touch the portion of the green circle that did not overlap the red whenever touching the red region incurred any penalty. However, because of the time limit, any planned movement resulted in a movement endpoint on the screen with substantial scatter from trial to trial (Fitts, 1954; Fitts & Peterson, 1964; Meyer, Abrams, Kornblum, Wright, & Smith, 1988; Murata & Iwase, 2001; Plamondon & Alimi, 1997). In Figure 21.2B, we show a hypothetical distribution of endpoints. These points are distributed around a mean endpoint marked by a diamond which, in Figure 21.2B, is at the center of the green circle. In Figure 21.2C–21.2D, we illustrate endpoint distributions with different mean endpoints. We return to these illustrations below.

The uncertainty in the actual location of the endpoint on each trial is the key problem confronting the subject in these tasks, and if we want to model the subject's behavior, we must have an accurate model of the subject's movement uncertainty. In all of the experiments reported by Trommershäuser et al. (2003a, 2003b), the distributions of endpoints were not discriminable from isotropic Gaussian and the distribution for each subject could therefore be characterized by a single number, σ , the standard deviation of the Gaussian in both the horizontal and vertical directions.

In Figure 21.3A, we plot 1,080 endpoints for one typical naïve subject together with quantile–quantile plots of the deviations in the horizontal and vertical directions. The quantile–quantile plots in Figure 21.3B–21.3C compare the distributions to a Gaussian distribution. Estimated values of σ varied from subject to subject by as much as a factor of 2. Trommershäuser et al. (2003a, 2003b) verified that the value of σ did not vary appreciably across the conditions of each experiment and across the range of locations on the screen where stimuli could be presented. Thus, if the subject displaces the mean endpoint by a small amount, the distribution of endpoints is simply shifted by that amount.

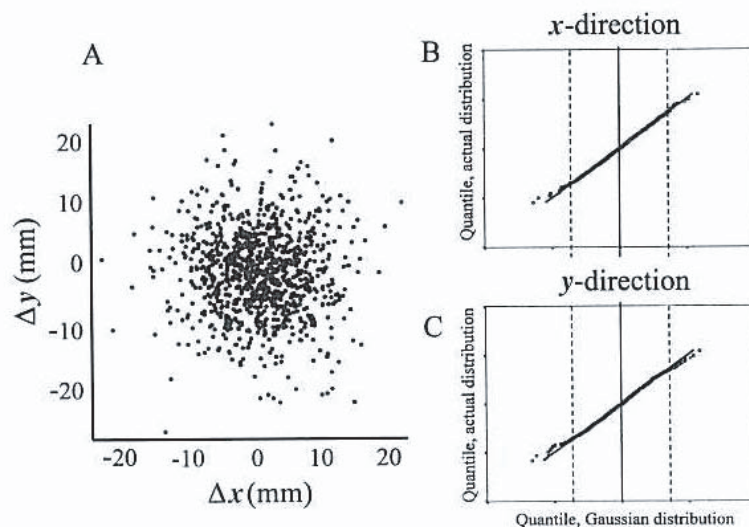


FIGURE 21.3 The distribution of endpoints. (a) Deviations Δx , Δy of 1,080 mean end points from the mean of the corresponding condition for a subject with motor variability $\sigma = 3.62$ mm. The distribution is close to isotropic and Gaussian. (b) Quantile–quantile plot (Gnanadesikan, 1997) of the deviations in the horizontal (x) direction compares the distribution of these deviations to the Gaussian distribution. (c) Quantile–quantile plot of the vertical (y) deviations. The linearity of these plots indicates that the distributions are close to Gaussian.

What should a subject do to maximize his or her winnings in the task just described? The subject can choose a particular movement strategy s , i.e. a plan of movement that is then executed. The relevant outcome of the planned movement in a given trial is the point where the subject touches the screen. Even if the subject executed the same plan over and over, the outcomes would not be the same as illustrated in Figure 21.2. By selecting a movement plan, the subject effectively selects an isotropic bivariate Gaussian density function $\phi_s(x, y; x_c, y_c, \sigma)$ of possible endpoints on the touch screen centered on the point (x_c, y_c) with standard deviation σ ,

$$\phi_s(x, y; x_c, y_c, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{(x-x_c)^2 + (y-y_c)^2}{2\sigma^2}}. \quad (5)$$

Subjects learn the task by performing more than 300 practice trials so that σ has stabilized before data collection starts. Therefore, the choice of movement strategy is the freedom to choose a mean endpoint (x_c, y_c) . In the following, we identify a movement strategy s with its mean movement endpoint (x_c, y_c) and we denote the distribution more compactly as $\phi(x, y; s, \sigma)$.

Under these experimental conditions, the subject's choice among possible movement strategies is precisely equivalent to a choice among lotteries. To see this, first consider the possible outcomes of their movement when there is one penalty and one reward circle present on the screen as in Figure 20.2A, with values of -500 and $+100$ points, respectively. A movement that hits the touch screen within the time limit could land in one of four regions: penalty only (region $R_{\bar{R}\bar{C}}$, value $V_{\bar{R}\bar{C}} = -500$), the penalty/reward overlap (region R_{RC} , value $V_{RC} = -400$), reward only (region $R_{\bar{R}C}$, value $V_{\bar{R}C} = 100$), or outside of both circles (Region $R_{\bar{R}\bar{C}}$, value $V_{\bar{R}\bar{C}} = 0$). The probability of each of these outcomes depends on the choice of mean endpoint. For example,

$$p_{RC}(s) = \int_{R_{RC}} \phi(x, y; s, \sigma) dx dy, \quad (6)$$

is the portion of the probability mass of the distribution that falls within the region R_{RC} when the mean endpoint is (x_c, y_c) .

The diamonds in Figure 21.2B–21.2D mark possible mean endpoints corresponding to different movement strategies s . For each movement strategy, we can compute the probability of each of the four outcomes, denoted $p_{RC}, \dots, p_{\bar{R}\bar{C}}$, by Monte Carlo integration. Any

such choice of movement strategy s corresponds to the lottery:

$$L(s) = (p_{RC}, V_{RC}; p_{\bar{R}\bar{C}}, V_{\bar{R}\bar{C}}; p_{\bar{R}C}, V_{\bar{R}C}; p_{RC}, V_{RC}). \quad (7)$$

The expected values of movement endpoints are given in Figure 21.2B–21.2D. In Figure 21.2B, for example, the mean endpoint is at the center of the green circle. Pursuing this strategy leads to the highest rate of hitting within the reward circle. However, if the penalty associated with the penalty circle is 500 as shown, the expected value of this strategy is markedly negative. In contrast, the strategy illustrated in Figure 21.2C has positive expected gain, higher than that corresponding to Figure 21.2B and 21.2D. However, we cannot be sure that any of the movement strategies illustrated in Figure 21.2B–21.2D maximizes expected value without evaluating the expected value of all possible movement endpoints.

Many other lotteries are available to the subject, each corresponding to a particular movement strategy or aim point and each with an associated lottery and expected value. In choosing among these possible motor strategies and all others, the subject effectively selects among the possible sets of probabilities associated with each outcome and an expected value associated with the associated lottery. This is illustrated in Figure 21.4, which shows the expected value corresponding to each possible mean endpoint as a surface plot (upper row) and as a contour plot (lower row) for three different penalty values (0, 100, and 500). The reward value is always 100. The maximum expected value (MEV) point is marked in each of the contour plots by a diamond. When the penalty is 0, the MEV point is at the center of the reward circle. As the penalty increases, the MEV point is displaced farther and farther from the center of the reward region. For a mean endpoint closer to the center of the reward region, the probability of reward is higher, but the probability of hitting within the penalty region is also higher. This increase in expected penalty more than cancels the increase in reward, i.e. the resulting expected value that corresponds to this mean endpoint is suboptimal. Alternatively, if the MEV point were farther from the center of the reward region, then the probability of hitting the penalty region would decrease but so would the probability of hitting the reward region. The MEV point strikes exactly the correct balance between risk and reward.

The MEV point depends on the geometry, number, and locations of the penalty and reward regions; the magnitudes of rewards and penalties; and the subject's

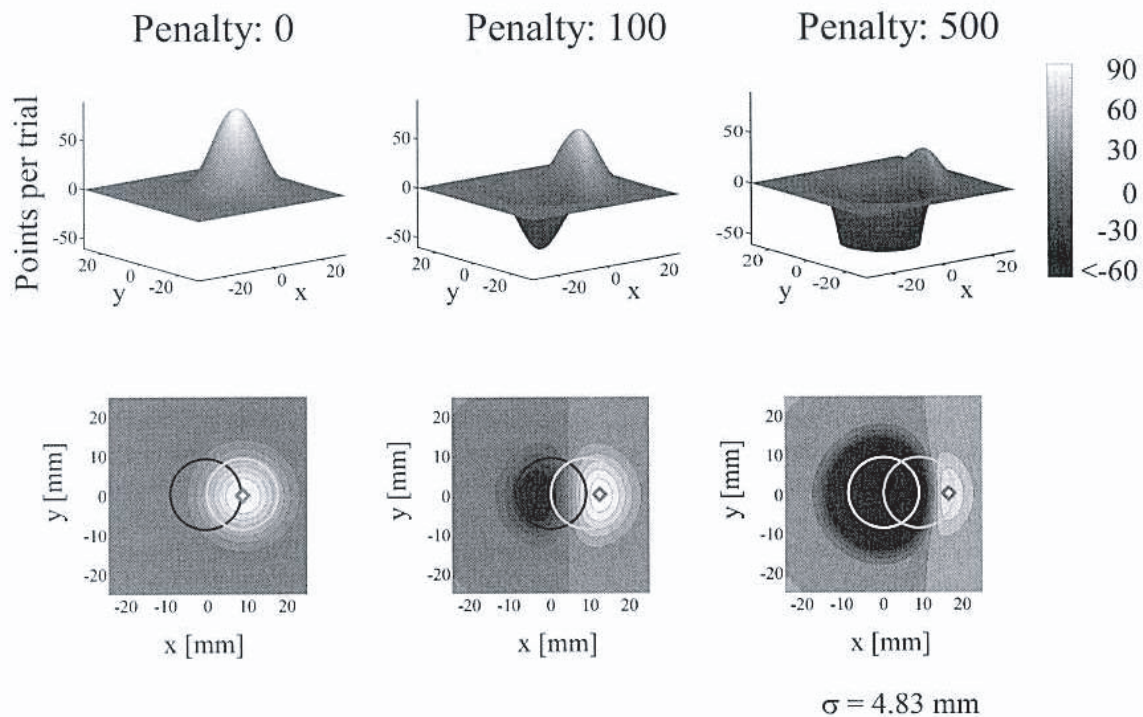


FIGURE 21.4 The expected value landscape. The expected value for each possible mean end point in the xy -plane is shown as surface plots (upper row) and corresponding contour plots (lower row) for three different penalty values (0, -100, -500). The maximum expected value (MEV) points in each contour plot are marked by diamonds. The MEV point is the center of the reward circle when the penalty is 0. It shifts away from the penalty circle as the penalty increases. Note that the penalty circle in the rightmost contour plot is coded as white rather than black so as to be visible.

own motor variability σ . These factors combine to create an infinite set of possible lotteries that the subject must choose among. Given the evident complexity of the decision making task implicit in Figure 21.2, it would be surprising if subjects chose strategies that maximized the expected value in these tasks. Moreover, the kinds of failures in decision making tasks that we discussed in the previous section should lead to particularly poor performance in these tasks.

If, for example, subjects interpret the penalty and reward regions in terms of loss and gain on each trial, then the loss aversion documented by Kahneman, Tversky, and others should lead them to move their mean endpoint further from the penalty region than the MEV point. Moreover, when the penalty is large relative to the reward, the probability of hitting the penalty region is very small for mean endpoints near the MEV point. If the subject overestimates the magnitude of this probability (as subjects did in the experiments illustrated in Figure 21.1), the subject will also tend to move too far away from the penalty region to be optimal.

In summary, we have very little reason to expect that subjects will approach optimal performance in movement planning tasks of the kind just described. They are very complex in comparison to ordinary decision tasks. The subject has little time to decide. Known patterns of failure in decision making should lead to poor performance. Given these expectations, the results of the experiments of Trommershäuser et al. (2003a, 2003b), presented next, are remarkable.

Movement Planning Under Risk: Initial Experimental Results

Method and Procedure

In this section we describe the results of Experiment 2 in Trommershäuser et al. (2003b). This experiment included four one-penalty configurations consisting of a reward circle and a single overlapping penalty circle (Figure 21.5A) and four two-penalty configurations




FIGURE 21.5 Stimulus configurations. The stimulus configurations employed in Experiment 2 of Trommershäuser et al. (2003b). The reward region (gray circle) was always assigned a reward of 100 points. The penalty assigned to the penalty regions (black circles) was either 0 (no penalty) or 500. The different spatial configurations were interleaved but the penalty values remained constant within a block of trials. (A) The one-penalty configurations. (B) The two-penalty configurations.

consisting of a reward circle and two overlapping penalty circles (Figure 21.5B). Trials were blocked with 32 trials per block (four repetitions of each of the eight configurations in Figure 21.5). The reward value was always 100 points and the penalty value was either 0 or 500 points (varied between blocks). If subjects hit a region shared by two or more circles, the subject incurred all of the rewards and penalties associated with all circles touched. In particular, if the subject touched within the region shared by two penalty circles, as many as 1000 points were lost, 500 for each penalty circle.

Subjects completed 24 trials in total for each of the 16 conditions of the experiment (two penalties crossed with eight stimulus configurations), a total of 384 experimental trials. Subjects were well practiced, having completed a similar experiment involving only configurations similar to those in Figure 21.2A. Before that first experiment, subjects carried out several training sessions. The purpose of the training trials was to allow subjects to learn to respond within the timeout limit of 700 ms and to give them practice in the motor task. A subject did not begin the main experiment until movement variability σ had stabilized and the timeout rate was acceptably low (see Trommershäuser et al., 2003b, for details).

Results

Figure 21.6A shows the MEV points for one of the one-penalty configurations and a particular subject with $\sigma = 2.99$ mm, the least variable subject in this experiment. These points of maximum expected value are computed by first estimating the expected-value landscape numerically (as in Figure 21.4) for each subject (i.e., each value of σ) and stimulus configuration and picking the mean endpoint that maximizes expected value. When the penalty is 0, the MEV point is the center of the reward circle, marked by the open circles in Figure 21.6. When the penalty is 500, it is displaced from the center of the reward circle, marked by the xs. Since all of the remaining one-penalty configurations are rotations of the first (Figure 21.5A), the MEV points for all four configurations are also rotations of the first. Subjects with higher values of σ had MEV points in the penalty 500 conditions that fell on the same radii but were displaced further from the center of the reward circle. In Figure 21.6B, we show the MEV point for one of the two-penalty configurations (for the same subject).

In Figure 21.7A, we summarize the results of Experiment 2 for the one-penalty conditions in Trommershäuser et al. (2003b), plotting each subject's mean




FIGURE 21.6 Maximum expected value (MEV) predictions. (A) The predicted MEV point for a one-penalty configuration in Experiment 2 of Trommershäuser et al. (2003b), for subject S5 for whom $\sigma = 2.99$ mm. These predictions are obtained by computing the expected value landscape as in Figure 21.4 for each combination of values, spatial configuration, and subject's motor variability. The mean end point that maximizes expected value is the MEV point. (B) The predicted MEV point for a two-penalty configuration in Experiment 2 of Trommershäuser et al. (2003b), for subject S5 with motor variability $\sigma = 2.99$ mm.

endpoint in each condition of the experiment. The predicted MEV points when the penalty value was 500 are indicated by the xs. The mean endpoints of the subjects are shown as solid circles. The subjects' mean endpoints lie close to the predicted values and there is no patterned deviation across subjects. Figure 21.7B contains the corresponding results for the two-penalty configurations.

Trommershäuser et al. (2003b) also computed a measure of the efficiency of each subject's performance. For each subject and each condition, we computed the expected winnings of an ideal subject that used a movement strategy corresponding to the subject's MEV point and had the same movement variability as the actual subject. We computed the ratio between the subject's actual earnings and the expected earnings of the ideal and refer to this ratio (expressed as a percentage) as the subject's *efficiency*. Note, that, if the subject were 100% efficient, then we would expect the measured efficiencies of the subject to be both greater than and less than 100%, since the actual earnings on any one repetition of the experiment could be greater than or less than the expected earning (just as 100 tosses of a fair coin could result in more or less than 50 heads). These calculations of efficiency are specific to the subject. An ideal subject with a larger motor variability than a second ideal subject would be expected to earn

less. Efficiencies are noted in Figure 21.7. They are not significantly different from 100% except for Subject S1 in the two-penalty configurations. Thus, subjects efficiently solve movement planning problems that are equivalent to decision making under risk with an infinite number of four term lotteries (the one-penalty case) or an infinite number of seven- or eight-term lotteries (the two-penalty case). They solve these problems in under 700 ms with estimated efficiencies exceeding 90%.

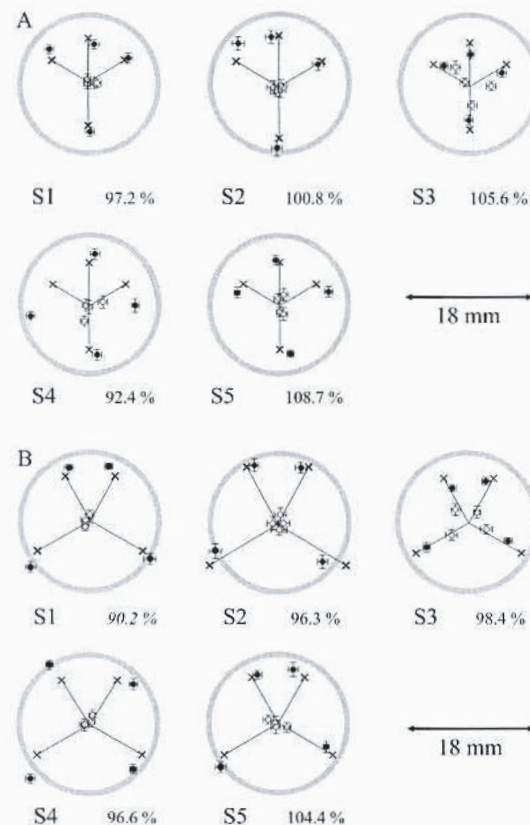


FIGURE 21.7 Results. (A) The predicted maximum expected value (MEV) points and actual mean end points for five subjects for Experiment 2 of Trommershäuser et al. (2003b). The one penalty configurations. (B) The two-penalty configurations. Open circles represent zero-penalty trials for which the MEV point is the center of the circle. Closed circles represent penalty -500 trials for which the MEV locations are indicated by the xs. The numbers to the right of each configuration indicate efficiency, computed as actual score divided by the MEV score (see Trommershäuser et al., 2003b, for details). Only one subject had an efficiency that was significantly different from 100% or optimal (*italicized*).

Trommershäuser et al. (2003b) considered the possibility that subjects only gradually learn the correct mean endpoint as a result of feedback. If this were the case, then we might expect to see trends in the choice of mean endpoint across the earliest trials in a condition. Subjects trained to hit within the center of the green circle during the preexperimental training phase might only gradually adjust their mean endpoints away from the penalty circle when the penalty is 500. Or, alternatively, a subject who is loss averse and given to exaggerating the (small) probability of hitting within the penalty circle might choose a mean endpoint that is initially too far from the penalty circle and only gradually adjust it toward the center of the reward region. In either case, we would expect to see a trend in mean aim point along the axis joining the center of the reward region and the centroid of the penalty region.

Trommershäuser et al. found no significant trends across subjects or conditions and concluded that subjects select movement strategies that maximize expected value almost immediately when confronted with a particular stimulus configuration. There is no evidence for learning in the data. Further, Trommershäuser, Gephstein, Maloney, Landy, and Banks (2005) have demonstrated that subjects can rapidly and successfully adapt to novel levels of movement variability imposed by the experimenter.

Explicit and Implicit Probabilities: An Experimental Comparison

Several factors may have contributed to the near-optimal performance of subjects in Trommershäuser et al. (2003a, 2003b). The movement planner makes a long series of choices and over the course of the experiment winnings increase. Decision makers faced with a series of decisions tend to move closer to MEV (Redelmeier & Tversky, 1992; Thaler, Tversky, Kahneman, & Schwartz, 1997; “the house money effect:” Thaler & Johnson, 1990). Further, the gain or loss associated with each trial is small. Studies of risky choice find that subjects are closer to maximizing expected value for small stakes (Camerer, 1992; Holt & Laury, 2002) and when subjects receive considerable feedback over the course of the experiment (Barron & Erev, 2003).

However, the most evident difference between movement planning under risk and ordinary decision making under risk is that, in the former, the subject is never given explicit information about probability distributions across outcomes associated with each possible

movement plan. He or she must in effect “know” the probability of hitting each region of the stimulus configuration in order to plan movements well. The results just presented suggest that human movement planning has access to such implicit probability information. We note that we do not claim that the subject has conscious access, but only that this information seems to be available for planning movement.

In the study reported next, we introduce an element of explicit probability information into the movement task of Trommershäuser et al. (2003a, 2003b). In explicit probability conditions, the outcome of each trial not only depended on where the subject touched the computer screen but also on an element of chance unrelated to motor performance.

Method

As in Trommershäuser et al. (2003a), the stimulus configuration consisted of a reward circle and a penalty circle (Figure 21.8A). The color of the penalty region varied between trials and indicated the penalty value for that trial (white: 0, pink: -200, red: -400 points). The reward circle was always green (drawn here as medium gray) and the reward value was always 100. The target and penalty regions had radii of 8.4 mm. The target region appeared in one of four possible positions, horizontally displaced from the penalty region by ± 1 or ± 2 multiples of the target radius (“near” and “far” in Figure 21.8A). The far configurations were included to keep subjects motivated through easily scored points, but were not included in the analysis. As in previous experiments, the stimulus configuration was displayed at a random location within a specified target region on each trial to prevent subjects from using preplanned strategies.

Explicit–Implicit Manipulation

The key manipulation concerned the certainty or lack of certainty of receiving a reward or penalty when a reward or penalty circle was touched. Either the reward or penalty circle could be stochastic. If a circle was stochastic, then the reward or penalty was obtained only 50% of the time when the circle was touched. The four probability conditions and the terms we use in describing them are shown in Figure 21.8B.² The *certainty condition* was similar to the experiment of Trommershäuser et al. (2003a).

The case Both 50% is of special interest. In this case, we scaled both the probability of getting a reward and the probability of incurring a penalty by 50%.

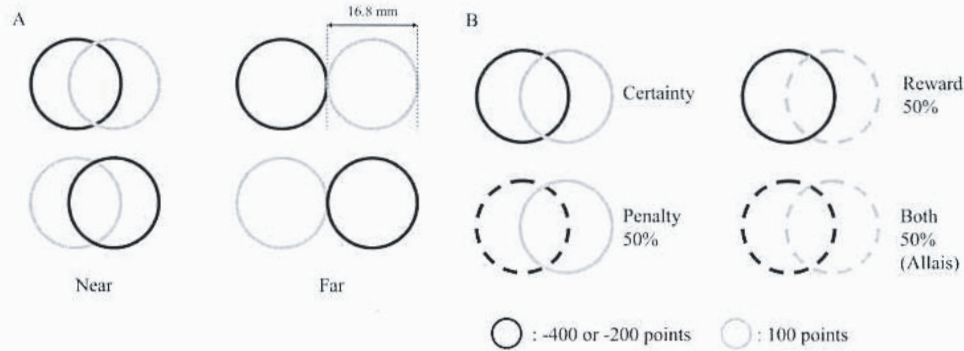


FIGURE 21.8 Experimental conditions. The configurations used in the final experiment. (A) The spatial conditions. (B) The probability conditions. In each condition, the subject received the reward or penalty associated with the region with either probability 1 (solid) or probability 0.5 (dashed). In the actual experiment, the penalty values were coded by color and the different probability conditions were constant across blocks and communicated to the subject at the start of each block.

The net result is to scale the expected value landscape by a factor of one-half. In particular, the location of the MEV point is unaffected. This condition is also of interest since not only should an ideal MEV mover choose the same mean aim point in both conditions, but also the performance of an MEU (utility-maximizing subject) should be invariant as well. To see this, we need only write down the conditions for the MEU subject to prefer lottery $L(s) = (p_1, O_1; \dots; p_n, O_n)$, over lottery $L(s') = (p'_1, O_1; \dots; p'_n, O_n)$,

$$\sum_{i=1}^n p_i U(O_i + W) > \sum_{i=1}^n p'_i U(O_i + W) \quad (8)$$

and note that multiplying both sides of the inequality by a positive factor to scale all of the probabilities does not affect the inequality. The ideal MEU subject will still prefer s even with all probabilities scaled by a common factor.

There are other connections between the conditions of the experiment that are summarized in Figure 21.9. Consider, for example, the Penalty 50% condition with penalty 400. The MEV subject should choose the same mean aim point as in the certainty condition with penalty 200 since the net effect of the stochastic penalty is to reduce the expected value of the penalty by a factor of 2. We refer to the conditions in which the




FIGURE 21.9 Equivalent conditions. (A) For the maximum expected value (MEV) movement planner, there are four pairs of probability conditions ("cases") where the expected value for any mean movement end point of the first configuration in the pair is either identical to that of the second or one-half of that in the second. The first condition in each case is explicit (either the reward or penalty or both is stochastic) and the second is implicit. Two cases are shown that are equivalent for an MEV movement planner but that might not be equivalent for an maximum expected utility (MEU) movement planner (with a non-linear utility function). Two cases are shown that would remain equivalent for an MEU movement planner as well as an MEV movement planner. For an MEV planner (A or B) or an MEU planner (B only), the predicted mean movement end points would be the same for the two configurations in each case.

probabilities are all implicit as *implicit conditions*, the conditions with explicit probabilities as *explicit conditions*. Figure 21.9 lists four equivalences between an implicit condition and an explicit condition that should hold for the MEV subject. We will refer to these pairs of equivalent conditions as “cases.” We emphasize that the computations involved in carrying out the MEV strategy with added explicit probabilities are very simple. The subject need only replace -400 by -200 , -200 by -100 , or 100 by 50 , depending on the pattern of explicit probabilities and payoff/penalty values in each session, to translate from the explicit condition of each case in Figure 21.9 to the equivalent implicit condition.

Procedure

Each session began with a test to ensure the subject knew the meaning of each color-coded penalty circle. The subject received feedback and had to correctly identify each penalty type twice. After a subsequent calibration procedure, there was a short block of 12 warm-up trials with zero penalty. The score was then reset to zero and data collection began. The time course of each trial was as described for previous experiments and the same feedback was given after every trial except as described below. The experiment comprised five experimental sessions of 372 trials each, 12 warm-up trials (not included in analyses), and 10 blocks of 36 trials. Blocks alternated between blocks containing configurations with penalty values of 0 and 200, and blocks with penalty values of 0 and 400. A penalty 200 block consisted of six repetitions of penalty 200 and three repetitions of penalty zero for each of the four spatial configurations, and the penalty 400 block was organized similarly. Each session corresponded to a stochastic condition in the order: certainty, Penalty 50%, Reward 50%, Both 50%, certainty. The certainty condition was repeated at the end (in the fifth session) to make sure that subjects’ reaction to certain rewards and penalties remained stable across the course of the experiment.³ Trials in which the subject left the start position less than 100 ms after stimulus display or hit the screen after the time limit were excluded from the analysis. Each subject contributed approximately 1,800 data points; that is, 60 repetitions per condition (with data collapsed across spatially symmetric configurations; 120 repetitions in the certainty conditions).

Results

Mean movement endpoints for each condition were compared with optimal movement endpoints as predicted by the optimal movement planning model of Trommershäuser et al. (2003a) based on each subject’s estimated motor uncertainty σ . Subjects’ *efficiency* was computed as previously described: the ratio between a subject’s cumulative score in a condition and the corresponding expected optimal score predicted by the model. We used bootstrap methods to test whether each subject’s measured efficiency differed from the optimal performance possible for that subject. In Figure 21.10, we plot actual points won for each subject and for each of two penalties in the certainty conditions. Five out of six subjects’ scores were statistically indistinguishable from optimal (their efficiencies were indistinguishable from 100%) replicating previous findings (Trommershäuser et al., 2003a, 2003b). Only one subject (data indicated by the dashed circles)

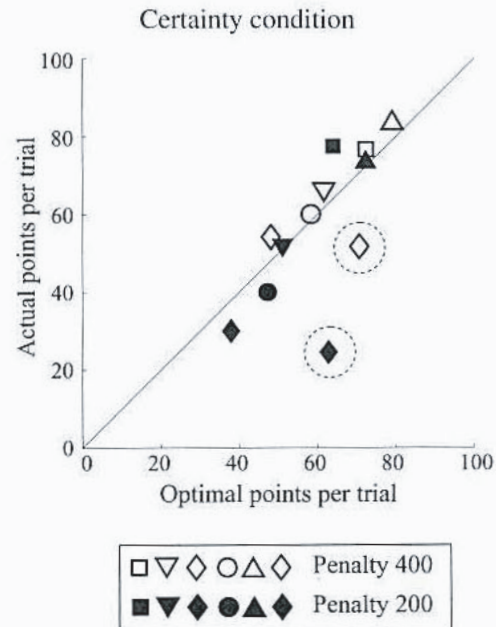


FIGURE 21.10 Results for the certainty conditions. A plot of points won versus the expected maximum expected value MEV points expected for each of the six subjects in the Certainty “near” condition. Different symbols correspond to individual subjects. Open symbols indicate the Penalty 400 condition, and filled symbols the Penalty 200 condition. Data points for which performance was significantly worse than predicted are indicated by the dashed circles.

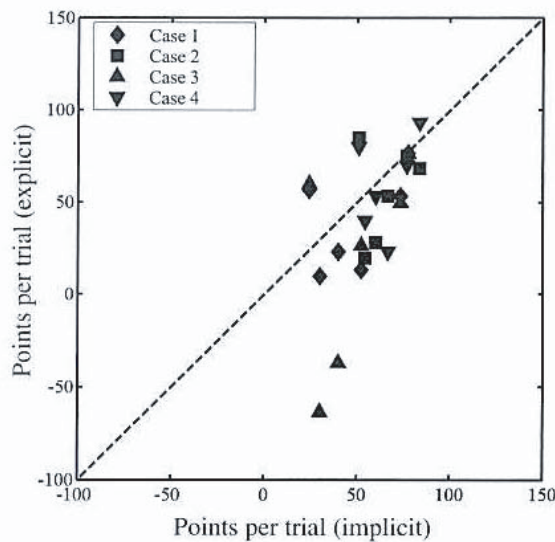


FIGURE 21.11 A plot of the outcomes of equivalent configurations. For each subject the mean number of points earned per trial for explicit probability conditions are plotted versus results for equivalent implicit probability conditions, for each of the cases in Figure 21.9. For the optimal maximum expected value observer, the expected values for the corresponding conditions (where one condition's results are doubled if necessary; see Figure 21.9) are identical. Thus, we would expect the plotted results to be distributed symmetrically around the 45-deg line. Instead, 19 out of 24 plotted points fall below the line. We reject the hypothesis that performance in explicit probability conditions equals that in equivalent implicit probability conditions (binomial test, $p = 0.003$).

differed significantly from optimal performance because the subject did not shift far enough away from the penalty region.

Next, we consider the pairs of conditions that should be equivalent for the MEV movement planner. We emphasize that the introduction of explicit probabilities imposed very little computational burden on subjects. One subject (JM) spontaneously reported after the experiment that she had followed the optimal strategy described above: a penalty of 400 that was incurred 50% of the time in the Penalty 50% condition should be treated exactly as the corresponding certainty condition with penalty 200, and so on. As we will see, however, her results are not consistent with the strategy she claimed to follow.

Figure 21.11 shows the winnings of each subject for each of the equivalent probability conditions shown in Figure 21.9 (doubling the winnings where one

condition is predicted to result in one-half the winnings of the other). In each case, the winnings from the implicit probability condition are plotted on the horizontal axis, the winnings from the equivalent explicit probability condition on the vertical axis. For the MEV movement planner, the plots should be randomly distributed above and below the 45-deg line. Instead, 19 out of 24 fall below, indicating that subjects tended to earn less in stochastic conditions than they did in equivalent certainty conditions. We can reject the hypothesis that subjects do equally well in the explicit as in the equivalent implicit probability conditions (binomial test, $p = .003$). The introduction of explicit probabilities tended to reduce subjects' winnings.

Subjects' performance dropped significantly below optimal when gains or losses were explicitly stochastic. This is also obvious in Fig. 21.12, which shows efficiencies of each subject in each condition. A subject following an MEV strategy would exhibit an expected efficiency of 100, marked by the horizontal line. For the certainty condition, one subject (AL) deviated significantly from MEV in both penalty conditions, as noted above. The estimated efficiencies for the other five subjects are distributed evenly around 100. In the three explicitly stochastic conditions, several other subjects exhibit large drops in efficiency, going as low as -400 (they are losing money at four times the rate they could have won money).

For five out of six subjects in the implicitly stochastic (certainty) conditions, earnings were indistinguishable from those expected from an optimal MEV strategy. Performance dropped significantly below optimal only when subjects were confronted with explicit uncertainty about whether they would incur a reward or penalty on a single trial. The introduction of explicit probabilities disrupted subjects' near-optimal performance even when the added cognitive demands were trivial.

Conclusion

We have introduced a class of movement tasks that are formally equivalent to decision making under risk and that model a very important class of everyday decisions "without words." In these decisions, there is a strong element of uncertainty in the outcome obtained as a consequence of any plan of action we choose, but this uncertainty originates in our own motor system. Our results suggest that, at least in some of these situations, we act as if we had good but "wordless" access to

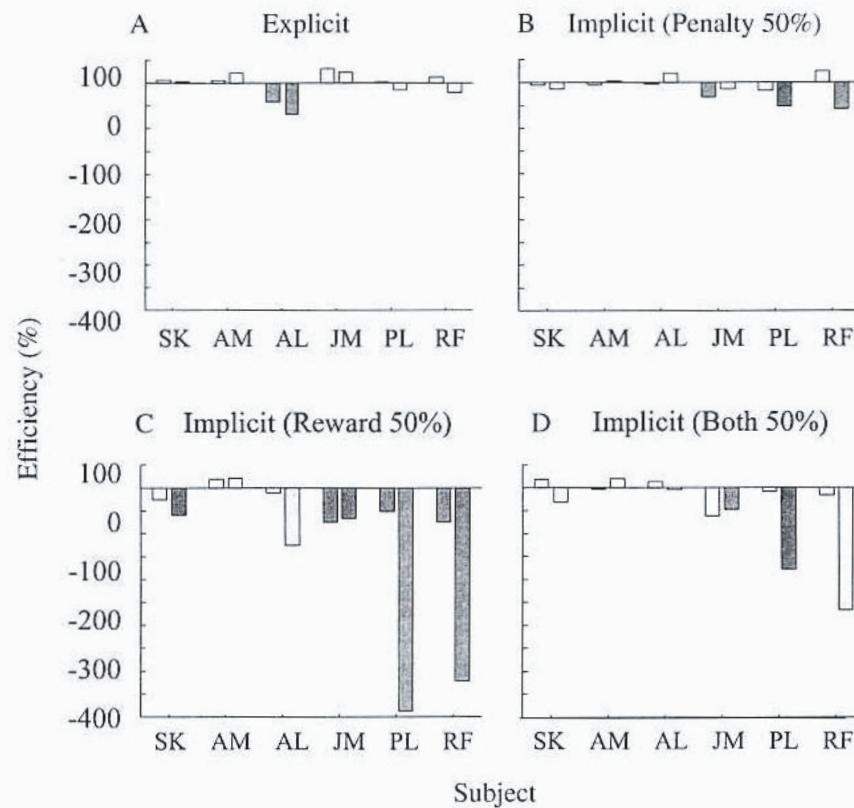


FIGURE 21.12 Efficiency is plotted for all subjects in all penalty and probability conditions. The expected efficiency for a maximum expected value movement planner is 100%. There are two bars for each subject, the left for the penalty 200 condition, the right for the penalty 400 condition. We tested each subject's estimated efficiency against optimal at the .05 level by a bootstrap method (Efron & Tibshirani, 1993). Gray bars indicate significantly suboptimal performance. (A) Certainty. In the certainty condition we could not reject the hypothesis of optimality for five out of six subjects (see also Figure 21.10). (B) Penalty 50%. (C) Reward 50%. (D). Both 50%. In the explicit probability conditions (B, C, D), all but one subject fell short of optimal in one or more conditions. In the five cases where the bar goes below 0, a subject could have won money on average, but instead lost money steadily.

estimates of the probabilities attached to the possible outcomes of any plan of action. Moreover, our use of these implicit probabilities is close to optimal.

Our results are consistent with the findings of Gigerenzer and Goldstein (1996) and Weber, Shafir, and Blais (2004; see also Hertwig, Barron, Weber, & Erev, 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 suboptimal decisions as pencil and paper tasks involving explicit probabilities. We note however that the absence of learning trends in Trommershäuser et al. (2003b) indicates that experience with a particular

stimulus configuration was not necessary. Subjects seemed to develop enough understanding of their own movement variability to allow them to adapt to novel configurations almost instantly.

The reader may still question whether the study of movement planning, as evinced by the tasks we have presented, has major implications for the classical study of decision making in economic tasks or whether it informs us about human cognitive abilities. Resolution of this issue will only follow considerable experimental work comparing human performance in both kinds of tasks. We will end, though, with two conjectures. First, we conjecture that the *cognitive illusions* that decision

makers experience in paper-and-pencil tasks are not representative of performance in the very large number of movement decisions we encounter in any single day. Second, we suggest that the human capacity for decision making bears the same relation to the economic tasks of classical decision making as human language competence bears to solving the Sunday crossword puzzle.

Acknowledgments

This research was funded in part by National Institutes of Health Grant EY08266, Human Frontier Science Program grant RG0109/1999-B, and by the Emmy-Noether-Programme of the German Science Foundation (DFG).

Notes

1. The stimulus configurations in all of the experiments discussed here were coded by colors, typically red ("penalty") and green ("reward"). We replace green by gray and red by black in the illustrations here.
2. The dashed lines used to represent probability condition in Figure 21.8B and following figures are for the convenience of the reader. The subject always saw solid circles with penalty coded by color and probability condition constant across each block of the experiment, communicated to the subject at the beginning of each block.
3. For one subject, mean movement endpoints differed significantly between the two certainty Sessions 1 and 5, indicating that his responses did not remain stable across sessions. His data were excluded from the analysis. Movement endpoints of the remaining six subjects were collapsed for Sessions 1 and 5.

References

- 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.
- Arnauld, A., & Nicole, P. (1662/1992). *La logique ou l'art de penser*. Paris: Gallimard.
- Atneave, F. (1953). Psychological probability as a function of experienced frequency. *Journal of Experimental Psychology*, 46, 81–86.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description based decisions. *Journal of Behavioral Decision Making*, 16, 215–233.
- Bell, D. E., Raiffa, H., & Tversky, A. (Eds.). (1988). *Decision making: Descriptive, normative and prescriptive interactions*. Cambridge: Cambridge University Press.
- Bernoulli, D. (1738/1954). Exposition of a new theory on the measurement of risk (L. Sommer, Trans.). *Econometrica*, 22, 23–36.
- Birnbaum, M. H. (2004). Human research and data collection via the Internet. *Annual Review of Psychology*, 55, 803–832.
- Camerer, C. F. (1992). Recent tests of generalizations of expected utility theory. In W. Edwards (Ed.), *Utility: Theories, measurement, and applications* (pp. 207–251). Norwell, MA: Kluwer.
- Efron, B., & Tibshirani, R. (1993). *An introduction to the bootstrap*. New York: Chapman-Hall.
- Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47, 381–391.
- , & Peterson, 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.
- Gnanadesikan, R. (1997). *Methods for statistical data analysis of multivariate observations* (2nd ed.). New York: Wiley.
- Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, 38, 129–166.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15, 534–539.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects in lottery choices. *American Economic Review*, 92, 1644–1645.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- , & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- , & Tversky, A. (Eds.). (2000). *Choices, values & frames*. New York: Cambridge University Press.
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Coombs, B. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory*, 4, 551–578.
- Meyer, D. E., Abrams, R. A., Kornblum, S., Wright, C. E., & Smith, J. E. K. (1988). Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review*, 95, 340–370.
- Murata, A., & Iwase, H. (2001). Extending Fitts' law to a three-dimensional pointing task. *Human Movement Science*, 20, 791–805.

- Plamondon, R., & Alimi, A. M. (1997). Speed/accuracy trade-offs in target-directed movements. *Behavioral Brain Sciences*, 20, 279–349.
- Redelmeier, D. A., & Tversky, A. (1992). On the framing of multiple prospects. *Psychological Science*, 3, 191–193.
- Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even—the effects of prior outcomes on risky choice. *Management Science*, 36, 643–660.
- , Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *Quarterly Journal of Economics*, 112, 647–661.
- Trommershäuser, J., Gepshtein, S., Maloney, L. T., Landy, M. S., & Banks, M. S. (2005). Optimal compensation for changes in task-relevant movement variability. *Journal of Neuroscience*, 25, 7169–7178.
- , Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision*, 16, 255–275.
- , Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and rapid, goal-directed movements. *Journal of the Optical Society A*, 20, 1419–1433.
- Tversky, A., & Kahneman, D. (1988). *Risk and rationality: Can normative and descriptive analysis be reconciled?* New York: Russell Sage.
- , & Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Risk and Uncertainty*, 5, 297–323.
- 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.
- Wills, G. (2001). *Saint Augustine's childhood: Confessions Book One*. New York: Viking.
- Wu, G., & Gonzalez, R. (1998). Common consequence effects in decision making under risk. *Journal of Risk and Uncertainty*, 16, 115–139.
- , & Gonzalez, R. (1999). Nonlinear decision weights in choice under uncertainty. *Management Science*, 45, 74–85.