Risk sensitivity in a motor task with speed-accuracy trade-off — **Source link** ↗

Arne J. Nagengast, Daniel Braun, Daniel M. Wolpert

Institutions: University of Cambridge

Related papers:

- Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty.

- Optimal feedback control as a theory of motor coordination.

- Optimal Compensation for Changes in Task-Relevant Movement Variability

- Risk-sensitivity and the mean-variance trade-off: decision making in sensorimotor control

- Statistical decision theory and trade-offs in the control of motor response

Risk sensitivity in a motor task with speed-accuracy trade-off

Arne J. Nagengast,[1,2]*   Daniel A. Braun,[1]* and Daniel M. Wolpert[1]

[1]*Computational and Biological Learning Laboratory, Department of Engineering, and* [2]*Department of Experimental Psychology, University of Cambridge, Cambridge, United Kingdom*

Nagengast AJ, Braun DA, Wolpert DM. Risk sensitivity in a motor task with speed-accuracy trade-off. *J Neurophysiol* 105: 2668–2674, 2011. First published March 23, 2011; doi:10.1152/jn.00804.2010.—When a racing driver steers a car around a sharp bend, there is a trade-off between speed and accuracy, in that high speed can lead to a skid whereas a low speed increases lap time, both of which can adversely affect the driver's payoff function. While speed-accuracy trade-offs have been studied extensively, their susceptibility to risk sensitivity is much less understood, since most theories of motor control are risk neutral with respect to payoff, i.e., they only consider mean payoffs and ignore payoff variability. Here we investigate how individual risk attitudes impact a motor task that involves such a speed-accuracy trade-off. We designed an experiment where a target had to be hit and the reward (given in points) increased as a function of both subjects' endpoint accuracy and endpoint velocity. As faster movements lead to poorer endpoint accuracy, the variance of the reward increased for higher velocities. We tested subjects on two reward conditions that had the same mean reward but differed in the variance of the reward. A risk-neutral account predicts that subjects should only maximize the mean reward and hence perform identically in the two conditions. In contrast, we found that some (risk-averse) subjects chose to move with lower velocities and other (risk-seeking) subjects with higher velocities in the condition with higher reward variance (risk). This behavior is suboptimal with regard to maximizing the mean number of points but is in accordance with a risk-sensitive account of movement selection. Our study suggests that individual risk sensitivity is an important factor in motor tasks with speed-accuracy trade-offs.

risk aversion; risk seeking; motor control; decision making

THE SPEED-ACCURACY TRADE-OFF is a central theme in movement neuroscience and implies that higher movement speeds are generally accompanied by a loss in movement accuracy (Schmidt and Lee 1999). The relationship between speed and accuracy was famously investigated by Paul Fitts who devised a mathematical model (now known as Fitts's law) to explain quantitatively how movement time relates to movement amplitude and target size (Fitts 1954). In his original task, subjects moved a handheld stylus back and forth between two equally sized targets as rapidly as possible, such that their required movement time could be measured as a function of target width and inter-target distance. Fitts found a logarithmic law relating movement time to the quotient of target distance and target width, with closer and larger targets requiring less movement time. One component that can contribute to this speed accuracy trade-off is the presence of signal-dependent noise in the motor system (Harris and Wolpert 1998; Jones et al. 2002; van Beers et al. 2004; Faisal et al. 2008), because movements with larger amplitudes induce more noise into the nervous system in the presence of signal-dependent noise. In addition, faster movements allow less time for feedback control to correct for errors that arise during the movement.

Recently, a large number of phenomena in human movement control have been modeled based on optimality principles that can take the noise properties of the motor system into account (Todorov 2004; Trommershäuser et al. 2003a,b), thus explaining, for example, speed-accuracy trade-offs (Dean et al. 2007), variability patterns (Todorov and Jordan 2002), bimanual movement corrections (Diedrichsen 2007), adaptation to force-fields and visuomotor transforms (Izawa et al. 2008; Braun et al. 2009), and control of objects (Nagengast et al. 2009). Crucially, the optimality principles considered by these computational theories of motor control have exclusively focused on the optimization of expected movement costs as payoffs and have ignored the variability of movement cost, that is risk. Recently, violations of risk neutrality have been reported in a number of motor control tasks. Wu et al. (2009) showed, for example, that subjects exhibit risk-seeking behavior in a pointing task, because they systematically underweight small probabilities and overweight large probabilities. Similarly, Nagengast et al. (2010) showed that subjects exhibit risk-averse behavior in a motor task that required them to control a Brownian particle under different levels of noise, consistent with the notion that subjects trade off mean and variance of movement costs (Nagengast et al. 2011).

Here we investigate the hypothesis that individual risk sensitivity affects motor tasks that involve trade-offs between movement speed and accuracy. To test this hypothesis, we designed a continuous decision task similar to sporting activities such as golf or tennis in which more forceful movements offer potentially greater rewards, such as an ace, but at the same time are more variable due to signal-dependent noise. We tested subjects on two conditions that provided the same maximum possible reward and only differed in their riskiness, that is the variability of the reward. A risk-neutral account of motor control predicts that subjects should behave identically in the two conditions. In contrast, a risk-sensitive account predicts that risk-seeking individuals should move faster whereas risk-averse individuals should move more slowly in the risky condition compared with the less risky condition. The results of our experiment suggest that movement arises from an interplay between subjects own speed-accuracy trade-offs and their risk sensitivity.

METHODS

Experimental Setup

Eleven right-handed subjects (6 male, 5 female, aged 20–28) participated in the experiment after providing written informed con-

* A. J. Nagengast and D. A. Braun contributed equally to this work.

Address for reprint requests and other correspondence: A. J. Nagengast. Computational and Biological Learning Lab, Dept. of Engineering, Univ. of Cambridge, Cambridge CB2 1PZ, United Kingdom (e-mail: arne.nagengast@gmail.com).

sent. The experimental protocols were approved by the Cambridge Psychology Research Ethics Committee. Subjects were naive to the purpose of the experiment, and none of the subjects reported any sensory or motor deficits. While seated, subjects used their right hand to grasp the handle of a vBOT force generating robotic manipulandum, which could be moved in the horizontal plane (for details, see Howard et al. 2009). The position and velocity of the hand were computed online at 1,000 Hz. Subjects could not see their hand, but the position of their hand could be displayed in the horizontal plane using a reflected CRT monitor. Subjects viewed the reflected image of the monitor screen through field-sequential shuttered glasses (Crystal Eyes, Stereo-graphic) by looking down at a mirror. Each eye was presented with an appropriate planar view of the scene, and the shuttered glasses alternately blanked the view of each eye in synchrony with the left and right views so that subjects perceived a three-dimensional scene.

Experimental Protocol

The general structure of the task was similar to a golf putting action. Subjects attempted to hit the center of a 6-cm target line that was 20 cm from the start position of their hand (Fig. 1A). Crossing the target line caused a ball to start translating in the horizontal plane orthogonal to the line. The initial location of the ball was the point where the hand crossed the line, and the distance the ball translated was proportional to the hand's velocity orthogonal to the line as it was crossed. A three-dimensional landscape was displayed that extended from the target line and whose height and color indicated the level of

reward (points) for different final locations of the ball (see Fig. 1). The display and the dynamic movement of the ball made the task intuitive for the subjects so that they had a good understanding of the reward landscape as a function of position and velocity. During the movement of the vBOT handle, online feedback of the subjects' hand position was extinguished to prevent visually driven online corrections. Movements that took longer than 0.75 s to complete had to be repeated.

Training session. Subjects first completed a training session of 500 trials in which four different permissible velocity ranges (indicated by a visible desired range of the final ball location) were used. The four conditions were blocked (125 trials each), and their order was randomized for every subject. In the training session, the subjects' reward depended on their lateral accuracy only and a Gaussian reward function of deviation from the center of the target was used. The reward was 50 points for a direct hit and decayed according to the following equation in which x is the lateral error:

$$R(x) = \left[\frac{1}{\sqrt{(2\pi)}} e^{\left(-\frac{x^2}{2}\right)} - 0.5 \right] \times 100 \qquad (1)$$

The training session was thus used to establish the relationship between subject's hitting velocity v and the lateral accuracy σ_x measured as the spread in distance x of the ball from the center line

$$P(x|v) \sim \mathcal{N}\left[\mu_x = 0, \ \sigma_x^2 = (a + bv)^2\right], \qquad (2)$$

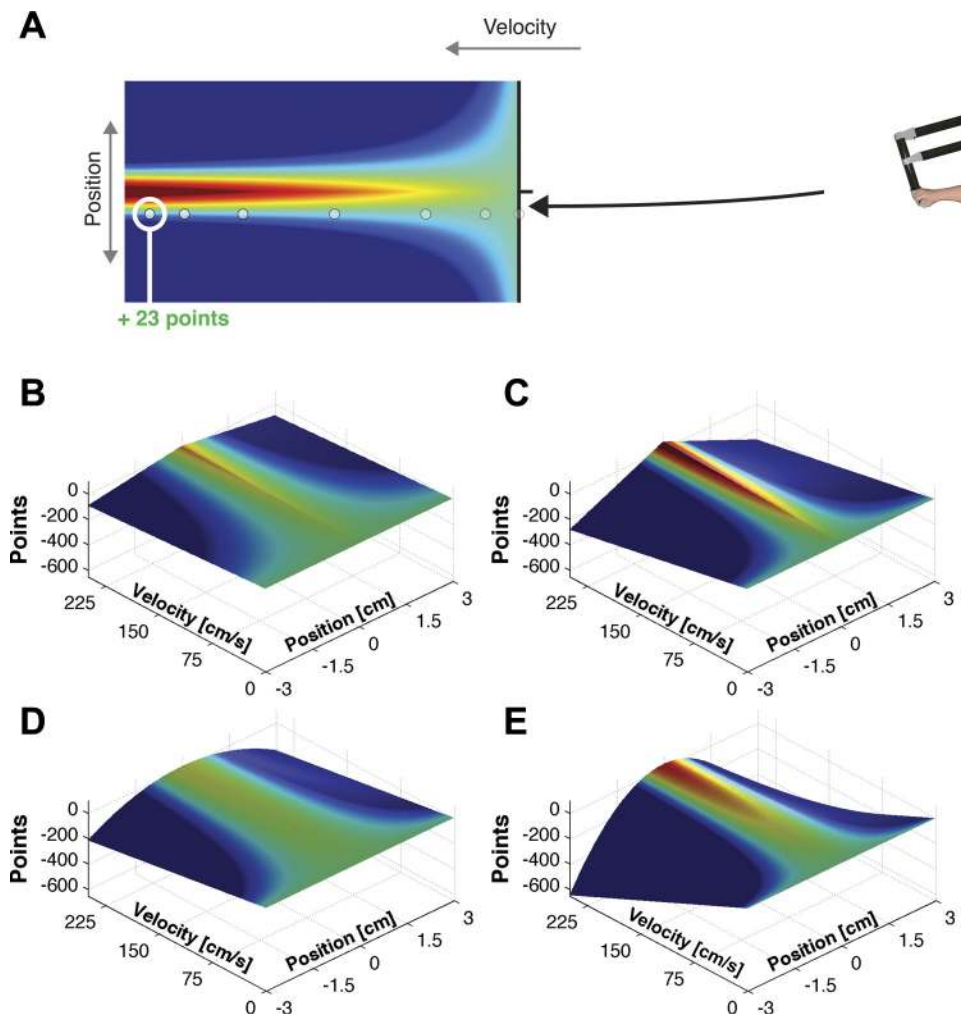where a and b are linear regression parameters for $\sigma_x = a + bv$ (see Fig. 2A).

Fig. 1. Schematic of the test session of the experiment. *A*: subjects held the handle of a robotic manipulandum and, in a limited time, attempted to hit the center of a visually presented target line and were free to decide with which velocity to execute the movement. Crossing the target line caused a ball to start moving in the horizontal plane orthogonal to the line. The initial location of the ball was the point where the hand crossed the line, and the distance the ball rolled was proportional to the hand's velocity orthogonal to the line as it was crossed. The payoff depended both on the final position and velocity of the subjects' hand, and this payoff was displayed to the subjects stereoscopically as a colored surface where the height (and color) determined the payoff. *B-E*: 4 reward functions used in the experiment as a function of final position and velocity of the subjects' hand when reaching the target line. In all 4 conditions, the reward increases linearly with the hand's velocity orthogonal to the target line but decreases with positional error (either linearly or quadratically depending on the condition). *B*: reward decreases linearly with absolute distance from the center of the target (R_{low}^1). *C*: same functional form as in *B*, but with the number of points scaled by a factor of 3 (R_{high}^1). *D*: reward decreases quadratically with distance from the center of the target (R_{low}^2). *E*: same functional form as in *D*, but with the number of points scaled by a factor of 3 (R_{high}^2).

A

B

Velocity [cm/s]   Distribution of Endpoints

229

171

128

81

C

D

Fig. 2. Speed-accuracy trade-off and reward variance. *A*: SD of lateral hand position at the target line vs. orthogonal velocity of all subjects during the training session. *B*: distribution of end-position changes with different velocities; single trial data and kernel density estimates for 4 different velocity bands of a representative subject. *C*: schematic of typical mean payoff vs. velocity for the linear reward function. The velocity that maximizes the number of points is the same in the high (gray line) and the low variance (black line) version for each class of reward function. *D*: payoff variance vs. velocity for the linear reward function. The payoff variance increases with velocity (due to signal-dependent noise leading to increased positional inaccuracy) and is 9 times higher in the high variance reward functions (gray) compared with the low variance (black) reward functions.
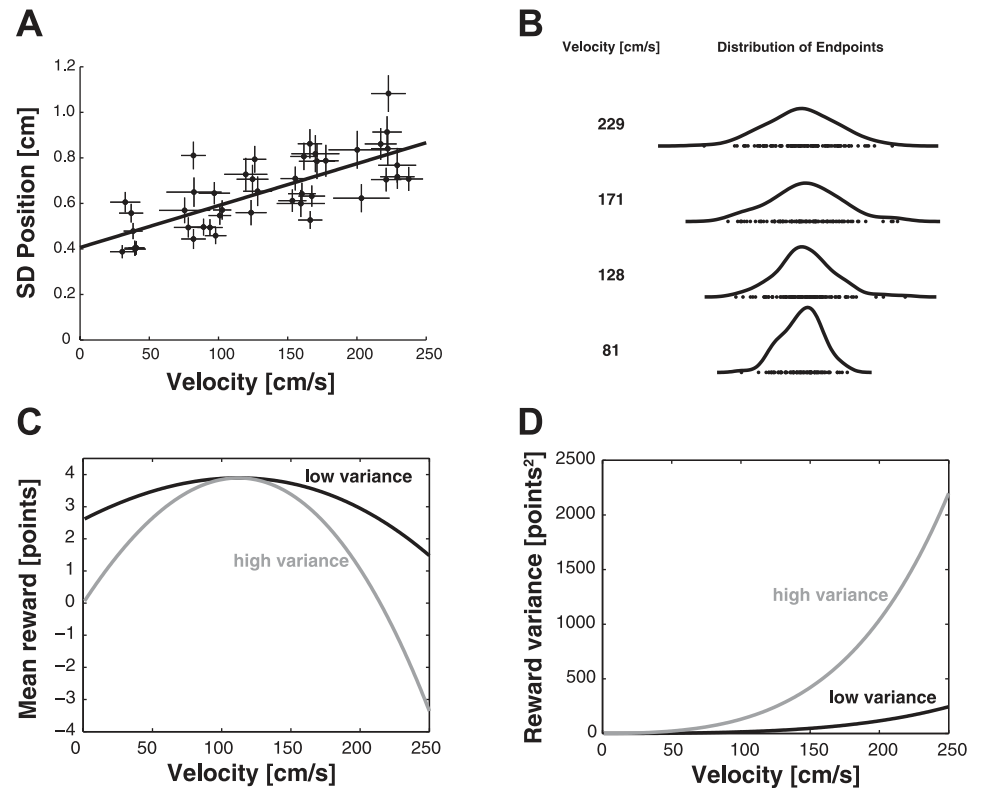
The purpose of the training session was threefold. First, it allowed subjects to become familiar with their variability in our task over a range of speeds. Second, it allowed us to ensure that the variance of position increased with velocity as assumed by our study. Third, it allowed us to estimate the regression coefficients for the linear fit of SD of lateral accuracy against speed for each subject (these values were used to adjust values d_1 and d_2 in *Eqs. 5* and *6*: see below).

The first four subjects were run with velocities of 75, 125, 175, and 225 cm/s. The range for the remaining subjects was slightly extended to 25, 90, 160, and 225 cm/s to get a more precise estimate of the relation between endpoint velocity and positional SD. As the data from the training session were used only to obtain a linear regression relation between SD of lateral accuracy and speed, the change in the speed range does not bias our analysis but potentially improves the accuracy of the regression fits. The tolerance on all these speeds was ±25 cm/s, and trials outside this range were repeated.

Test session. The training session was followed by the test session that consisted of a total of 200 trials and four different reward functions (see *Reward functions* for details). Blocks of 5 trials for a given reward landscape were interleaved pseudorandomly until subjects had completed 50 trials of each reward function. In the test session, the subjects' reward depended on both their lateral accuracy and endpoint velocity, which was visualized using the reward surfaces. Subjects were free to move at the velocity they desired and were only told to maximize the total number of points obtained during the test session. The number of points obtained in the current trial and a cumulative total were displayed in the corners of the screen.

Reward functions. The general setup of the task was such that the reward or loss on a given trial was a function of the subjects' lateral accuracy, which was then scaled by their endpoint velocity. Two different classes of reward function were used, the first, R^1, had a linear and the second, R^2, a quadratic relation between reward and positional lateral accuracy. Each class of function had a low and a high variance version for which the variance of the latter was nine times as high as for the low variance condition. The reward function

depended on the distance (x in cm) of the hand position relative to the center of the target line (lateral accuracy) and the orthogonal velocity to the line v (cm/s), giving the four reward functions

$$R^1_{high}(x, v) = 0.45 \times (0.5 - |x|)v \qquad (3)$$

$$R^2_{high}(x, v) = 0.3 \times (0.4 - x^2)v \qquad (4)$$

$$R^1_{low}(x, v) = \frac{R^1_{high}}{3} + d_1 \qquad (5)$$

$$R^2_{low}(x, v) = \frac{R^2_{high}}{3} + d_2 \qquad (6)$$

where d_1 and d_2 were adjusted for each subject so that the mean reward level at the optimum velocity (based on the linear regression fits in the training session) was the same for the low and the high variance condition. For both classes, the low variance reward was a simple linear transformation of the high variance reward. Since utilities are cardinal, that is they are only determined up to an affine transform, the two utility functions of the high variance and the low variance condition are equivalent. This means that for a risk-neutral decision maker who is trying to find the optimal hitting velocity v by maximizing the expected reward $E[R^i(v)] = \int_{-\infty}^{\infty} P(x|v)R^i(x, v)dx$, where we obtain $P(x|v)$ from the training session, the optimal velocity $\bar{v}^i = \arg\max_v E[R^i(v)]$ should be the same in the high variance and in the low variance condition. In contrast, risk-sensitive subjects should choose different hitting velocities for the high variance and the low variance condition, since they take the variance of the payoff into account.

Note that the reward functions themselves are a deterministic mapping between lateral accuracy and orthogonal velocity of the subject's hand on impact with the target line. The difference in reward variance arises from the subject's natural movement variability from trial-to-trial.

Risk sensitivity. Risk sensitivity refers to sensitivity not only to the mean payoff but also to its variance. For a distribution of payoffs (R), the subjective value of the payoffs is $E(R) - \theta Var(R)$. The parameter

θ expresses the decision maker's risk attitude: risk-neutral decision makers are only sensitive to the expected payoff ($\theta = 0$), while risk-averse individuals discount payoff variability ($\theta > 0$) and risk seekers consider it a bonus ($\theta < 0$). Therefore, a risk-averse subject, by definition, will prefer to accept a lower mean reward if it decreases the variance of the reward and conversely for a risk seeker.

Figure 2, *C* and *D*, shows schematics of the predicted mean and variance of the payoff for the linear reward (R^1) for the high (gray lines) and low (black lines) variance conditions as a function of velocity. Given a subject's own noise properties, there is a velocity that maximizes the expected reward (peak of the parabola in Fig. 2*C*). As explained above, this optimal speed is the same for the low and high variance conditions. If a subject deviates from this peak, the mean reward will decrease. However, deviations that increase the speed will lead to higher variance (Fig. 2*D*) and be attractive to risk-seeking individuals. Conversely, risk-averse individuals can reduce their speed so as to reduce variance at the cost of a reduced mean. In general, the larger the absolute value of θ, the greater the deviation of the velocity from the peak.

To assess risk sensitivity from our data, we examined the change in chosen velocity between the high and low variance conditions. Whereas determining risk sensitivity from the velocity chosen in a particular condition relies on an accurate estimate of the speed-accuracy trade-off, the change in speed between low and high variance conditions is independent of the parameters of the speed-accuracy trade-off. Compared with the low variance condition, in the high variance condition the mean reward shape is scaled by a factor of 3 (with the offset set so that the maximum possible reward is unchanged), whereas the variance is scaled by a factor of $3^2 = 9$. Therefore, the new subjective value becomes $3E(R) - 9\theta Var(R) +$ offset. The optimum velocity therefore can be found by removing the offset and dividing by three to optimize $E(R) - 3\theta Var(R)$. This is equivalent to finding the optimum for the low variance task for a subject with a θ value three times larger, that is more risk seeking or more risk averse. Therefore, for a risk-averse individual they will tend to move in the direction of reducing their speed in the high variance condition. The converse is true for the risk-seeking individual. Therefore, we can classify behavior as risk averse and risk seeking if the speed decreases or increases, respectively, between the low and high variance conditions. The use of the linear and quadratic costs allows us to examine the speed changes in two different payoff regimes. We only classify subjects as risk seeking or risk averse if they show appropriate and significant changes in velocity between the low and high variance conditions of both regimes.

RESULTS

We designed a motor task that involved a speed-accuracy trade-off in which subjects could influence the variance of their own movement strategy. The general structure of the task was similar to a golf putting action. A three-dimensional landscape was stereoscopically displayed whose height indicated how good the subjects' drive was on a given trial (see Fig. 1). That is their payoff given in points depended both on how accurately (lateral accuracy) and how vigorously (velocity) a ball was struck (see Fig. 1*A*). A spot-on hit and a near-miss would result in a small positive and a small negative payoff when gently hitting the ball, whereas it would result in a large positive and a large negative payoff when striking it vigorously. As movements are contaminated with signal-dependent noise and therefore become more inaccurate the more forceful they are, subjects had to trade-off (to use the golf drive analogy) getting the ball further towards the hole at the risk of potentially putting it into a side bunker. This trade-off can be seen in Fig. 2*A* where we show that the endpoint spread of the ball is dependent on the hitting velocity.

The relationship was fit with a linear regression (SD = $0.4062 + 0.0018v$; $r^2 = 0.5499$; $P < 0.001$). We tested subjects on two classes of reward functions, one with a linear and the other with a quadratic change in reward with positional lateral accuracy. For a given class of reward function, there were two conditions that only differed in their riskiness (variance), that is one offered potentially greater rewards but also potential ruin while in the other less could be gained but also less could be lost. For a given relation between a subjects' endpoint velocity and positional SD, the mean reward is at a maximum for a particular velocity, while both lower and higher velocities offer on average less points (Fig. 2*C*). As the two reward functions of one class were simple linear transformations of one another, their maximum mean reward for a given subject was at the same velocity but with a ninefold difference in variance (Fig. 2*D*). A risk-neutral account predicts that subjects only maximize the mean reward and are indifferent to the greater reward fluctuation of the high variance condition, that is they should perform identically in the two conditions. In contrast, risk-seeking individuals are expected to move faster and risk-averse individuals more slowly in the high variance condition compared with the low variance condition. Therefore, subjects who had significantly higher speeds for the high variance condition of both linear and quadratic functions were classified as risk seeking. Conversely, subjects who had significantly lower speeds for the high variance condition of both linear and quadratic functions were classified as risk averse (see METHODS).

In our analysis, we found that the behavior of most subjects differed between the low and high variance conditions (Fig. 3). In 17 out of the 22 cases, the velocity of the two conditions was different (all $P < 0.05$, two-sample *t*-test with unequal variances for velocities of the low and the high variance condition). However, it could be that the differences in the two conditions were just due to chance and that they had nothing to do with the difference in variance. To check for consistency, we plotted (Fig. 4) the differences in the mean velocity $\pm$ their SE for the two conditions of a given class of reward functions against the difference in the other class. If differences in the two velocities were just due to noise, we would expect the plot to have no structure. In fact, behavior in the two classes of reward functions was highly correlated ($r^2 = 0.79$; $P < 0.005$), suggesting that the variance of the outcome influenced the subjects' motor strategies consistently depending on their risk attitude. The data suggest that three subjects were significantly risk seeking and three significantly risk averse.

To investigate potential nonstationarities in the subjects' performance, we performed two additional analyses. First, to test whether the variability (and hence shape of the curves in the Fig. 2, *C* and *D*) changed between the training and test session, we tested whether the coefficients of the regression of SD in position to velocity for each subject in the training session changed in the test session (for this session we binned the velocities to calculate the SD in position). We used a Chow test on the test and training data to test if the coefficients of the two regressions were different, and this showed that they were not significantly different for 10 out of 11 subjects (P values all > 0.05). One subject had a P value of 0.0449, but this subject was not classified as either risk seeking or risk averse. Second, we tested whether subjects differed in their reaction to a low or high score on the next trial by increasing or decreasing their velocity. We correlated the payoff on one trial
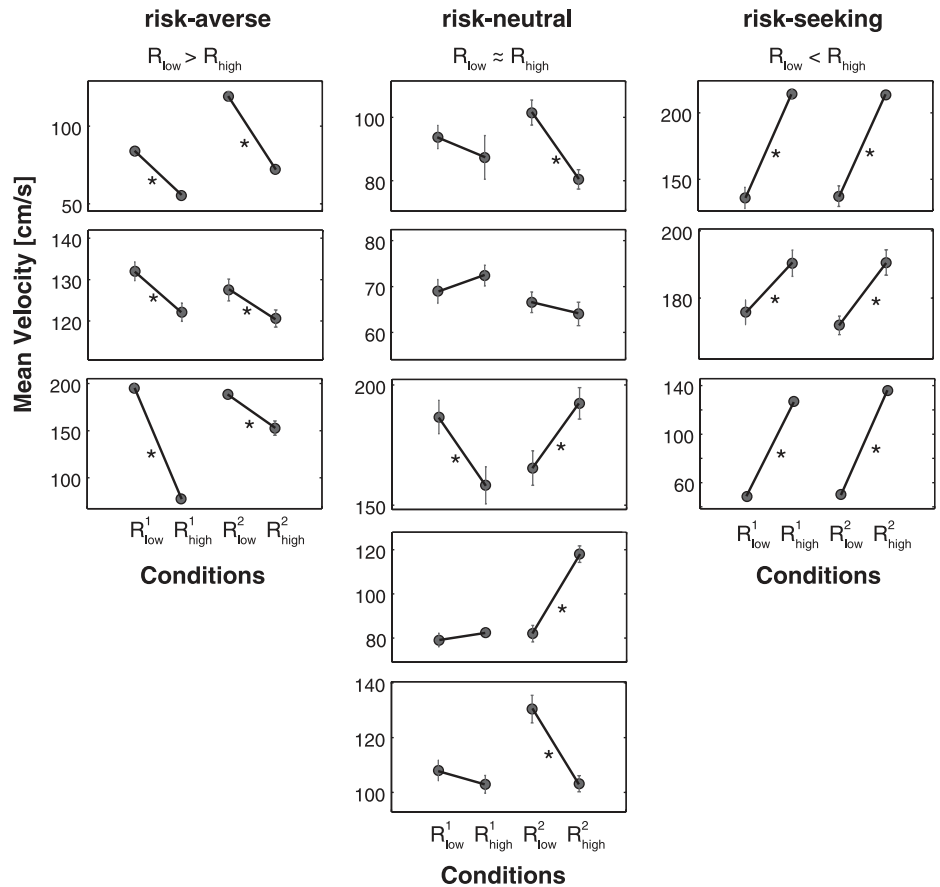
Fig. 3. Mean velocity in the 4 different reward functions for the 11 subjects. The mean velocity with 1 SE across trials for condition R^1_{low} (linear reward function - low reward variance), condition R^1_{high} (linear reward function - high reward variance), condition R^2_{low} (quadratic reward function - low reward variance), and condition R^2_{high} (quadratic reward function - high reward variance). Subjects at *left* column were classified as risk averse and those at *right* as risk seeking. Subjects at *middle* showed no systematic pattern of velocity changes and are classed as risk-neutral. *$P < 0.05$, significant differences between high and low variance conditions.

with the velocity on the next. For nine subjects, there was no significant correlation, and for two (one risk averse and one risk seeking), there was a small but positive correlation of 0.15 and 0.33, respectively.

Only the orthogonal component of the hands velocity relative to the target line had an impact on the payoff. However, the same orthogonal component could be achieved with different tangential components. To test if the tangential component of the velocity differed between subjects with different risk sensitivities, we calculated the ratio of the magnitude of tangential velocity to the magnitude of orthogonal velocity for each trial for the three most risk-seeking and the three most risk-averse subjects. A comparison of the ratio between these two groups showed that the ratio was not significantly different (ratios for risk seeking: 0.0254, 0.0381, and 0.0726; risk averse: 0.0265, 0.0383, and 0.0624).

DISCUSSION

We examined whether subjects are sensitive to the variance of payoffs in a motor task that involved a speed-accuracy trade-off. Subjects had to hit the center of a target line with high speeds to accumulate points. In this task, the variability arose from their own motor noise in line with a speed-accuracy trade-off. We found that most subjects were sensitive to the variance of the payoffs, with some risk seeking and others risk averse. The nature of an individual's risk sensitivity was well maintained across two different reward functions (Fig. 4). Our results suggest that when people make movements that involve a speed-accuracy trade-off they take the variability of cost (or reward) into account and adjust their behavior depending on their risk attitude. The results are inconsistent with a risk-neutral account of motor control.


Fig. 4. Velocity differences between the 2 classes of reward functions. The difference in mean velocity between the high and the low variance conditions of the quadratic reward function vs. the linear reward function with 1 SE across trials. Data points for individuals whose behavior is consistent with risk aversion are at *bottom left*; those who are risk seeking are at *top right*. Three subjects were classified as risk seeking and 3 as risk averse. *$P < 0.05$, subjects who had significant changes in velocity for both tasks and were in the risk-seeking or risk-averse quadrants.

Most current computational accounts of motor control are risk neutral (Todorov 2004; Trommershäuser et al. 2003a,b) and only consider minimization of the expectation of a cost function, usually with terms for positional lateral accuracy and effort. Thus the variance of the cost does not influence these models when computing the optimal movement policy. Recently, this hypothesis of risk neutrality in motor control has been challenged. In the study of Wu et al. (2009), subjects made binary choices in a pointing task with monetary payoffs. Their choice patterns violated the expected utility hypothesis when confronted with choices of different mean and variance with respect to the payoffs. Wu et al. (2009) explained these deviations by a distortion of the estimated probabilities, as is typically done in prospect theory in economics (Tversky and Kahneman 1992). In the study of Nagengast et al. (2011), subjects also made binary choices in a pointing task but with a payoff that was given in effort (force). In their experiment, variance levels of different choices were directly manipulated. In each trial, subjects had to make a choice between a known effort on the one hand, and a risky effort on the other hand, that had a mean effort that could either be higher or lower than the certain effort. As the variance levels were fixed beforehand, this setup allowed a direct measurement of how subjects trade off mean effort against variance in effort. This trade-off between mean and variance can also be applied to normative theories of motor control, such as optimal feedback control theory (Whittle 1981; Nagengast et al. 2010).

Since the work of Fitts (1954), there has been a substantial body of work showing that many motor tasks can be understood as a simple speed-accuracy trade-off. However, the key result of the current study is that motor behaviors are determined not only by the speed-accuracy trade-off but also by the interplay of this trade-off and subjects' risk sensitivity. That is subjects exploit their own speed-accuracy profile, selecting a mean and variance of the payoffs (from those possible from their speed-accuracy trade-off) that is consistent given their risk-sensitivity preference. In contrast to our previous work (Nagengast et al. 2010), the noise in the current study is not imposed experimentally but derives from the subjects' own natural movement variability; these two sources of noise may not be treated similarly (Langer 1975). Previous studies have suggested that subjects maximize expected value in such a situation (Trommershäuser et al. 2003a,b), whereas the current study suggests risk sensitivity is also an important determinant of behavior. In another study, we showed that movements are chosen according to a mean variance trade-off in the payoff (Nagengast et al. 2011). In this previous study, we used a task in which subjects had a discrete choice between an option that requires movement with a medium, but certain, effort and an option that requires either a probabilistically selected low or high effort movement. As opposed to this previous study, the current study considers a more natural scenario in which subjects can choose from a continuum of velocities. Therefore, the current study is able to show that the movement arises from an interplay between subjects' own speed-accuracy trade-off and their risk sensitivity, and subjects can exploit their own speed-accuracy trade-off according to their risk-sensitivity preference.

Risk sensitivity has been previously studied in personality psychology and economic decision making. A key difference between personality questionnaire studies in psychology and our paradigm is that we take each individual's dexterity into account. Risk-assessment questionnaires sometimes involve questions such as "Do you enjoy rock climbing?" The problem with such questions is that the answer depends both on people's risk attitudes but also on their motor skills and previous experience. Our study avoids this possible confounding factor, as we assessed the risk attitude of individual subjects in a skill-independent way. Although the subjects' motor variability did determine their optimal movement speed, crucially between high and low variance conditions, the optimal speed stays the same if they are risk neutral. Thus the method provides behavioral tools of assessing risk attitude that are independent of the individual level of skill. However, previous studies have also shown that risk sensitivity in an individual is highly context dependent and can change across situations (see, for example, Hanoch et al. 2006; Weber et al. 2002; MacCrimmon and Wehrung 1986; Yates 1992). An individual's risk sensitivity in one task can differ from the risk sensitivity of the same individual in another task. Here, we could show that the same risk sensitivity was maintained in subjects for different payoff functions.

The speed-accuracy trade-off is a fundamental property found in many motor control tasks (Schmidt and Lee 1999). In our task, we found a linear relationship between movement speed and endpoint spread when the payoffs were simply determined by lateral accuracy. This finding is in line with previous studies (Schmidt et al. 1979; Schmidt and Lee 1999) that have also found a linear speed-accuracy trade-off between movement speed and end-point spread. In these studies, subjects were required to make individual point-to-point aiming movements with predefined movement times from defined starting positions to final positions that could vary in distance. The difference to the logarithmic law found by Fitts could be due to the fact that the movements were ballistic rather than cyclic as in Fitts task (Schmidt and Lee 1999) or might be an approximation to the logarithmic relationship for our range of task parameters. Similar speed-accuracy trade-offs have also been found for eye movements (Abrams et al. 1989; Patla et al. 1985). It will therefore be interesting in the future to examine how risk-sensitivity impacts speed-accuracy trade-offs in different motor tasks.

REFERENCES

Abrams RA, Meyer DE, Kornblum S. Speed and accuracy of saccadic eye movements: characteristics of impulse variability in the oculomtor system. *J Exp Psychol* 15: 529–543, 1989.

Braun DA, Aertsen A, Wolpert DM, Mehring C. Learning optimal adaptation strategies in unpredictable motor tasks. *J Neurosci* 29: 6472–6478, 2009.

Dean M, Wu S, Maloney L. Trading off speed and accuracy in rapid, goal-directed movements. *J Vis* 7: 1–12, 2007.

Diedrichsen J. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr Biol* 17: 1675–1679, 2007.

Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci* 9: 292–303, 2008.

Fitts P. The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47: 381–391, 1954.

Hanoch Y, Johnson J, Wilke A. Domain specificity in experimental measures and participant recruitment: An application to risk-taking behavior. *Psychol Sci* 17: 300–304, 2006.

Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.

Howard IS, Ingram JN, Wolpert DM. A modular planar robotic manipulandum with end-point torque control. *J Neurosci Meth* 181: 199–211, 2009.

Izawa J, Rane T, Donchin O, Shadmehr R. Motor adaptation as a process of reoptimization. *J Neurosci* 28: 2883–2891, 2008.

Jones KE, Hamilton AF, Wolpert DM. Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88: 1533–1444, 2002.

Langer E. The illusion of control. *J Pers Soc Psychol* 32: 311–328, 1975.

MacCrimmon KR, Wehrung DA. *Taking Risks: The Management of Uncertainty.* New York: Free Press, 1986.

Nagengast AJ, Braun DA, Wolpert DM. Optimal control predicts human performance on objects with internal degrees of freedom. *PLoS Comput Biol* 5: e1000419, 2009.

Nagengast AJ, Braun DA, Wolpert DM. Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty. *PLoS Comput Biol* 6: e1000857, 2010.

Nagengast AJ, Braun DA, Wolpert DM. Risk-sensitivity and the mean-variance trade-off: decision making in sensorimotor control. *Proc Biol Sci* 2011 Jan 5. [Epub ahead of print.]

Patla AE, Frank JS, Allard F, Thomas E. Speed-accuracy characteristics of saccadic eye movements. *J Mot Behav* 17: 411–419, 1985.

Schmidt R, Lee T. *Motor Control and Learning. A Behavioral Emphasis* (3rd. ed.). Champaign, IL: Human Kinetics, 1999.

Schmidt RA, Zelaznik H, Hawkins B, Frank JS, Quinn JT. Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol Rev* 47: 415–451, 1979.

Todorov E. Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915, 2004.

Todorov E, Jordan MI. Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.

Trommershäuser J, Maloney LT, Landy MS. Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A* 20: 1419–1433, 2003a.

Trommershäuser J, Maloney LT, Landy MS. Statistical decision theory and trade-offs in the control of motor response. *Spat Vis* 16: 255–275, 2003b.

Tversky A, Kahneman D. Advances in prospect theory: cumulative representation of uncertainty. *J Risk Uncertainty* 5: 297–323, 1992.

van Beers RJ, Haggard P, Wolpert DM. The role of execution noise in movement variability. *J Neurophysiol* 91: 1050–1063, 2004.

Weber E, Blais A, Betz N. A domain-specific risk-attitude scale: measuring risk perceptions and risk behaviors. *J Behav Dec Making* 15: 263–290, 2002.

Whittle P. Risk-sensitive linear/quadratic/Gaussian control. *Adv Appl Probab* 13: 764–777, 1981.

Wu SW, Delgado MR, Maloney LT. Economic decision-making compared with an equivalent motor task. *Proc Natl Acad Sci USA* 106: 6088–6093, 2009.

Yates JF. *Risk-Taking Behavior.* New York: Wiley, 1992.