Trading off speed and accuracy in rapid, goal-directed movements

Mark Dean Department of Economics, New York University, New York, NY, USA

Shih-Wei Wu Department of Psychology, New York University, New York, NY, USA

Laurence T. Maloney Department of Psychology and Center for Neural Science, New York University, New York, NY, USA

Many studies have shown that humans face a trade-off between the speed and accuracy with which they can make movements. In this article, we asked whether humans choose movement time to maximize expected gain by taking into account their own speed–accuracy trade-off (SAT). We studied this question within the context of a rapid pointing task in which subjects received a reward for hitting a target on a monitor. The experimental design we used had two parts. First, we estimated individual trade-offs by motivating subjects to perform the pointing task under four different time constraints. Second, we tested whether subjects selected movement time optimally in an environment where they were rewarded for both speed and accuracy; the value of the target decreased linearly over time to zero. We ran two conditions in which the subjects faced different decay rates. Overall, the performance of 13 out of 16 subjects was indistinguishable from optimal. We concluded that in planning movements, humans take into account their own SAT to maximize expected gain.

Keywords: visuomotor control, movement planning, optimality, statistical decision theory, speed–accuracy trade-off, decision making

Citation: Dean, M., Wu, S.-W., & Maloney, L. T. (2007). Trading off speed and accuracy in rapid, goal-directed movements. *Journal of Vision, 7*(5):10, 1–12, http://journalofvision.org/7/5/10/, doi:10.1167/7.5.10.

Introduction

In executing any type of movement, there is typically a trade-off between the speed with which the movement is performed and the degree of precision with which it is made. Fitts (1954) first provided a quantitative description of the speed–accuracy relation in self-paced, cyclic tapping movements. The characterization, often referred to as Fitts' law, stated that movement time is a logarithmic function of task difficulty indexed by the ratio of movement amplitude and target width. Subsequent research that followed verified Fitts' formal description in a wide range of movement tasks (Crossman & Goodeve, 1983; Fitts & Peterson, 1964; for reviews, see Meyer, Smith, Kornblum, Abrams, & Wright, 1990; Plamondon & Alimi, 1997). It has also been shown, however, that Fitts' law fails in tasks where subjects are asked to move to the target at a specified time. Such tasks, often called temporally constrained tasks, differed from the spatially constrained tasks that Fitts' law described well. In temporally constrained tasks, studies have often observed a linear relation between spatial error and task difficulty. Attempts to explain the divergent findings led to the development of several important theories of speed–accuracy trade-off (SAT; Meyer, Abrams, Kornblum, Wright, & Smith, 1988; Plamondon & Alimi, 1997; Schmidt, Zelaznik, Hawkins, Frank, & Quinn, 1979) that incorporated the empirical regularities under different task constraints.

The relation between speed and accuracy is often critical to the final outcome of a movement, that is, its final position. Consequently, whatever the form of the SAT, it is critical to ask if agents take into account their own SAT when planning movements. A recent study by Augustyn and Rosenbaum (2005) showed that human subjects' choice of starting position of movement between two targets reflects knowledge of the SAT predicted by Fitts' law. In daily situations, many motor tasks require the agent to perform with high accuracy at high speed, in the sense that the task rewards both speed and accuracy. In such contexts, the agent immediately faces a decision problem that requires the knowledge and consideration of his or her own SAT: Should he or she move too fast, he or she would gain advantage in speed but sacrifice accuracy; should he or she slow down his movements, he or she would achieve higher accuracy but with loss of speed.

In this study, we examined whether humans can take into account their own SAT in tasks that reward both speed and accuracy and if they do so optimally. To model this decision, we extended a previous model proposed by Trommershäuser, Gepshtein, Maloney, Landy, and Banks (2005) and Trommershäuser, Maloney,

doi: 10.1167/7.5.10 Received November 15, 2006; published July 30, 2007 ISSN 1534-7362 © ARVO

and Landy (2003a, 2003b). In a series of studies, Trommershäuser et al. successfully modeled movement planning as the solution to an optimal control problem. Their model assumes that movement strategies are chosen to maximize expected gain given the costs and benefits explicitly implemented in the environment. When choosing a strategy, the movement planner takes into account his or her own intrinsic motor variability. Their models do well in predicting movement endpoints chosen by the subjects in rapid, goal-directed pointing tasks (but see Wu, Trommershäuser, Maloney, & Landy, 2006, on the limits to movement planning).

Those studies, however, did not model the SAT of a subject, a ubiquitous feature underlying almost every movement. The experiments of Trommershäuser et al. effectively fixed the length of time a subject had to perform the task by imposing a large "timeout" monetary penalty. As time taken to perform the task varied little from trial to trial after extensive training, their model treated motor variability as exogenously fixed for each subject (although motor variability was modeled as varying between subjects). This constraint is, however, artificial. In most tasks, a person gets to choose how long he or she takes over a movement and, hence, the accuracy of that movement. The degree of motor variability that a person faces becomes an endogenous choice variable. We therefore extended the model of Trommershäuser et al. by incorporating this choice into a new optimization model and compared its predictions to subjects' motor behavior.

Our experimental design had two sessions: a training session and an experimental session. In the training session, we ran a sequence of treatments in which subjects were rewarded for performing a pointing task within various time limits. In the experimental session, subjects performed the same pointing task, but in this case, the reward for successfully hitting the target decreased linearly with time after its presentation. We used data from both the training and experimental sessions to estimate the relationship between movement speed and accuracy for each subject. Armed with this estimated speed–accuracy relationship, we could calculate the optimal movement time for each subject in the experimental session. Having done so, we compared the predicted choice of movement time from the optimal model to the actual choice of movement time exhibited by the subjects.

Materials and methods

Overview

The environment in which we explored the trade-off between speed and accuracy was similar to that previously used to examine optimality in motor tasks (Trommershäuser et al., 2003a, 2003b). Subjects were presented with a circular visual stimulus on a touch screen, which they aimed to hit with their index finger. Hitting the target resulted in a monetary reward. Within this structure, each subject took part in two sessions, which we labeled the training session and the experimental session. The training session was designed to elicit the subject's SAT. Each subject was presented with four treatments in which the subject was asked to perform the pointing task within four different time limits. In each treatment, the subjects received a monetary reward if they hit the target within the time limit, but they received a large monetary penalty if they failed to make contact with the screen before the time limit expired.

The experimental session was designed to test whether the subjects selected the right point on the SAT in environments where there were benefits associated with both quick and accurate movements. We therefore presented each subject with tasks in which the reward for successfully hitting the target decreased with the time taken. Unlike the training session, there was no time limit applied to the task. Instead, the subjects were free to choose how much time to take in their movement. The longer they took to make the movement, the more likely they were to hit the target but the lower the reward for doing so.

Each subject took part in two types of trial in the experimental session: "fast" decay trials in which the value of hitting the target decreased quickly and "slow" decay trials in which it decreased more slowly. Subjects were divided between two conditions, which varied only in the speed of the fast decay rate and the initial value in the fast decay trials. Condition A had a modest difference in speed between the fast and slow rates. The results from Condition A motivated us to run a second condition, Condition B, in which the fast decay rate was faster than that in Condition A, to see if subjects could perform as well on this more difficult task.

Apparatus

A touch monitor (Elo IntelliTouch 17-in. LCD monitor) was mounted vertically on a framework (Structural Framing System, McMaster Carr Inc.). This framing system was specifically selected to minimize the vibration of the setup caused by the speeded reaching movement to the monitor. A chin rest was used to control viewing distance, which was 30 cm in front of the monitor. The computer keyboard was mounted on the table and centered in front of the monitor. The experimental room was dimly lit. The experiment was run using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) on a Pentium 4 Dell OPTIPLEX GX280. For each subject, at the beginning of every experimental session, a calibration procedure for touch location was performed to optimize the accuracy of recorded endpoints.

Stimuli

Training session

The subjects were asked to hit a green target (radius = 7 mm) within a given time limit. Subjects started every trial by depressing the spacebar on the keyboard. Once the spacebar was held down, a central fixation cross appeared for 1.5 s and was followed by presentation of a blue square region (74 × 74 mm) delimiting the possible locations of the target. The fixation cross disappeared after the blue region was displaced. The blue region was always presented at a fixed location. Shortly after the blue region was presented, the target and a "time" bar, displayed horizontally on top of the blue region, were presented simultaneously. The length of the time bar indicated how much time remained before the time limit was reached. Movement onset was defined as the time the subject's finger left the spacebar. To prevent subjects from preplanning the movement, we perturbed both the timing of target onset and the target location. The target was presented 500 ms plus an amount drawn from a uniform distribution with a range of 100 ms after the onset of the blue region. Target location was perturbed in both the x and y directions with a range of ±23 mm relative to the screen center.

Experimental session

The setup was the same as in the training session, except that the time bar was replaced by a "money" bar to indicate the target value that decreased over time. See the Procedure section for details.

Procedure

The experiment comprised two conditions, A and B. Each participant in the experiment took part in one condition only. Each condition consisted of a training session, in which participants were faced with explicit time constraints, and an experimental session. The training sessions in the two conditions were identical. The experimental session consisted of a mixture of fast decay rate blocks and slow decay rate blocks, differentiated by the speed at which the monetary reward decayed. The slow decay times in the experimental sessions of both conditions were identical (1,000 ms to decay from an initial value of 100 points). The only differences between Conditions A and B were the reward profile and the decay times in the fast decay rate blocks in the experimental session. The fast time to decay was 770 ms in Condition A with an initial value of 100 points, whereas the fast time to decay in Condition B was 606 ms with an initial value of 200 points.

All subjects took part in a training session and an experimental session. A subject would participate in the two sessions on separate days but within 48 hr of each other.

In the training session, we implemented four time constraints for the task. The first four subjects had the following time limits: 400, 600, 800, and 1,000 ms. We found that with the 800-ms limit, all subjects had more than enough time to hit the target on essentially every trial: Their performance was at a ceiling. Accordingly, we reduced the time limits for the remaining subjects to 400, 535, 650, and 775 ms to improve estimation of the SAT. The timer started upon target presentation, and the time that the subject had left in a trial was indicated by a horizontal white bar (the "time bar") whose length decreased continuously over time. The bar was displayed above the target configuration. On every trial, subjects earned 100 points by hitting the target within the time limit but received a 700-point penalty for not hitting the screen within the time constraint. Hitting the screen but missing the target within the time limit earned the subject 0 points. At the end of the experiment, points were converted into money for the subject at a rate of 1 cent for every 20 points scored. Each constraint was run in separate blocks of trials. The entire session consisted of two runs, each consisting of four blocks of different constraints. The order of constraints was randomized for each run. Each block started with 20 practice trials with no monetary consequences, followed by 40 "live" trials. The subjects were notified that the order of blocks was randomized and that they had to make an effort to adjust speed based on the time constraint of a particular block. All subjects were advised to take breaks (3 min) between blocks and shorter breaks as needed (20 s) in between trials, especially within the block of shorter time constraints, to minimize the impact of fatigue. The training session took approximately 90 min to complete.

In the experimental session, the value of hitting the target was not fixed but decreased linearly over time to zero. There was no time limit to perform the task. We implemented the two conditions, A and B, with different sets of decay rates and assigned half of the subjects to each set.

In Condition A, target initial value was fixed at 100 points and decreased at two rates, fast and slow, run in two separate blocks. In the fast decay block, target value decreased to zero at 770 ms after its presentation, whereas target value decreased to zero at 1,000 ms after presentation in the slow decay block. In Condition B, we drastically increased the faster decay rate to 606 ms and also increased the initial value to 200 points. We increased the initial value so that the maximum expected gain in the slow and fast decay conditions would be roughly the same. Subjects started each block with 40 practice trials with no monetary consequences and performed 120 "live" trials. The order of fast and slow blocks was balanced across subjects. The experimental session took approximately 45 min to complete.

Figure 1 provides a graphical depiction of the different decay rates, whereas Figure 2 provides a visual depiction of the task (Panel A for the training session and Panel B for the experimental session).

Subjects and instructions

Sixteen subjects, unaware of the purpose of the experiment, participated. Among them, seven were male and nine were female. Nine were graduate students from the Department of Economics in New York University. The remaining subjects were graduate students from the Psychology Department or students from the Law School. All subjects except for one were right handed, and all had normal or corrected-to-normal vision. Informed consent was given by all subjects prior to the experiment. Subjects received US$36 (US$24 from the training session, US$12 from the experimental session) plus the additional bonus they earned through their performance. Total payment ranged from US$40 to US$55 across subjects.

Model of optimal movement planning

Subjects' performance in the task was compared to an optimal movement-planning model based on statistical decision theory (Blackwell & Girshick, 1954; Berger, 1985; see Maloney, 2002, for discussion). The model is an extension of Trommershäuser et al. (2003a, 2003b). Here, we use "optimal" to refer to the maximization of expected monetary reward.

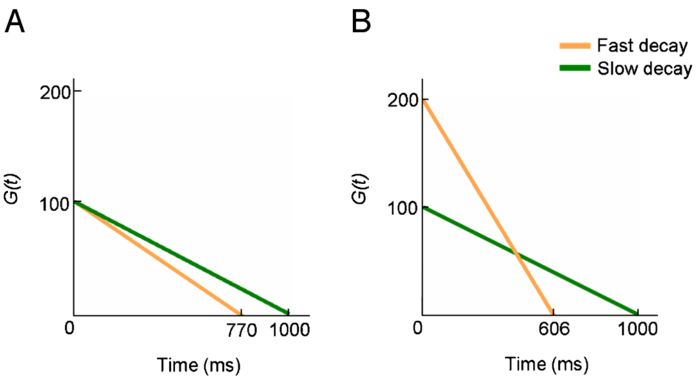

Figure 1. Summary of decay conditions. Target gain, G (in points), was plotted as a function of time (ms). (A) Decay conditions of Condition A. Target value started at 100 points but decreased in two different rates. In the slower decay rate blocks, the value decreased to zero at 1,000 ms. In the faster decay rate blocks, the value decreased to zero at 770 ms. Values started decreasing as soon as the target appeared. (B) Decay conditions of Condition B. The slow condition had the same decay rate as that in Condition A. However, in the fast decay rate blocks, target value initiated at 200 points but decreased 3.3 times faster than the slow decay rate blocks.

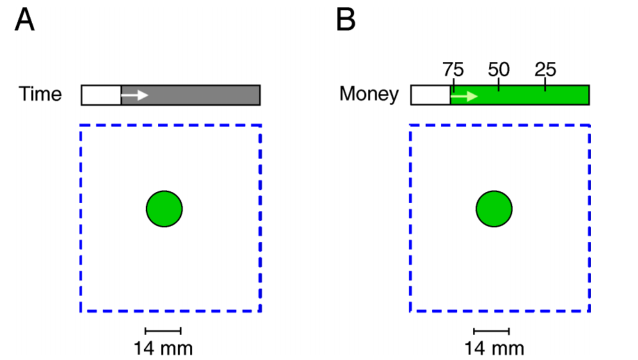

Figure 2. Stimulus configurations. (A) Stimulus configuration in the training session. Subjects saw the circular target (radius = 7 mm) within the blue rectangular region. Subjects had limited time to attempt the target and earned fixed monetary reward (100 points = 5 cents). The time bar above the configuration decreased in size continuously to indicate the amount of time left in a given trial. (B) Stimulus configuration in the experimental session. Unlike the training session, there was no time constraint in the experimental session. Instead, target value decreased rapidly over time. The time bar from the training session was replaced by a money bar to provide continuous feedback on value.

A general form of such an optimization problem is as follows:

$$\text{Choose } s \in S \text{ to maximize } \Gamma(s) = \int_G R(g) f(g|s) \mathrm{d}g, \quad (1)$$

where s is a motor strategy, S is the set of all possible strategies, G is a set of outcomes (i.e., realized movements), $R{:}G \rightarrow \Re$ are the monetary penalties or rewards associated with the events in G, and $f(.\,|\,s)$ is a probability distribution over the outcome space G conditional on choosing movement strategy s. The idea behind the model is that, because of motor uncertainty, when an agent selects a motor strategy, he or she is really selecting a probability distribution over realized movements. The model posits that the selected motor strategy should maximize expected gain conditional on the probability distribution generated by that strategy. The movement that actually takes place is then drawn from this conditional distribution.

In theory, S could be very large, containing a vast array of motor strategies described as a detailed sequence of motor commands. To make the model tractable, assumptions are used to reduce the strategy space to manageable proportions. In Trommershäuser et al., a strategy consisted of selecting a target point on the screen representing the target for the endpoint of the movement. Thus, a strategy could be represented as a tuple of $\bar{x}$ and $\bar{y}$ target coordinates. An innovation of this article is to extend the description of a strategy to include a time element. A strategy therefore now consists of a triple—an $\bar{x}$ and a $\bar{y}$ coordinate representing the target endpoint and a time $\bar{t}$

representing the target length of time to take over a movement.

For this experimental design, we can simplify the above general model significantly. First, assuming that spatial errors are symmetric, the optimal choice of x and y is trivial: One should always aim for the middle of the circle. Thus, we do not model the choice of x and y explicitly. Second, the reward that the agent receives depends only on whether or not the target is hit and the time at which the hit occurs. Third, although actual movement time will be stochastic, we show in the Results section that the probability of hitting the target depended only on planned movement time and not on actual movement time. Thus, we can rewrite Equation 1 as

Choose $\bar{t} \in (t^*, \infty)$ to maximize

$$\Gamma(\bar{t}) = \int_0^\infty (R(A,t)p(A|\bar{t}) + R(A^c,t)p(A^c|\bar{t}))f(t|\bar{t})\mathrm{d}t, \quad (2)$$

where A is the event that the target is hit, A^c is the event that the target is missed, t is actual movement time, $R(.,t)$ is the reward associated with a particular event occurring when actual movement time is t, $p(. \mid \bar{t})$ is the probability of an event conditional on planned movement time $\bar{t}$, and $f(. \mid \bar{t})$ is the distribution of actual movement time conditional on planned movement time. t^* is a lower bound on planned movement time.

To be able to solve for the agent's optimal choice of $\bar{t}$ in the experimental session, we need to determine the nature of $p(A \mid \bar{t})$ and $f(t \mid \bar{t})$, both of which we do on an individual-by-individual basis. We estimated the function $p(A \mid \bar{t})$ using the data gathered from both the training and experimental sessions. To do so, we make two assumptions: First, we assume that $E(t \mid \bar{t}) = \bar{t}$ or that the expected length of time taken to perform a movement conditional on a planned movement time is equal to that planned movement time. Second, we assume that, for a set of trials in a given treatment, the subject chooses the same target time. We can then approximate the planned movement time for a treatment by taking the sample average of movement times within that treatment. Furthermore, we can approximate the probability of hitting the target associated with that planned movement time by calculating the proportion of hits within the same treatment. Thus, the training and experimental sessions provide six points on the SAT. We use these points to approximate the function $p(A \mid \bar{t})$.

Finally, we need to determine the nature of $f(t \mid \bar{t})$. We assume that the distribution takes the form of a truncated normal, with the truncation point set arbitrarily at 200 ms. As discussed above, we assume that the mean of the distribution is $\bar{t}$. Data from the training session suggest that the variance of the distribution of t increases linearly with $\bar{t}$ in most subjects. We therefore modeled the timing variance as a linearly increasing function of $\bar{t}$.

We are now in a position to calculate the optimal choice of $\bar{t}$ in the experimental session. In this session, the reward for missing the circle was always zero, whereas the reward schedule for hitting the circle, $R(A, t)$, varied between decay rate blocks. We therefore know everything we need to calculate the optimal movement plan for a particular agent in a particular treatment. As the optimization does not have an analytical solution given the functional form we have chosen for $f(t \mid \bar{t})$, we found the optimum by numerically computing $\Gamma(\bar{t})$.

Data analysis

For each trial, we recorded arrival time (time from target onset to arrival at the touch monitor),[1] the screen position (x, y) that was hit, and the score. Movement endpoints were recorded relative to the center of the target circle.

The effect of planned time and actual time on spatial accuracy

To determine the optimal arrival time for a subject, we needed to determine the exact nature of the SAT that the subjects faced. To begin with, we needed to determine whether it was planned or actual movement time that determined accuracy. To do so, we analyzed how spatial accuracy was affected by planned arrival time and actual arrival time separately. To do the former, and using the assumption that mean arrival time accurately reflects planned time within a treatment, we performed a regression analysis on spatial accuracy by mean arrival time. For each time constraint in the training session and decay rate in the experimental session, we computed the "absolute error" (the mean of the distance between movement endpoints[2] and target center) and mean arrival time. We then regressed mean distance by mean arrival time. To do the latter, for each individual, we regressed the distance from target center on actual time taken and a treatment dummy using all observations. By the assumption that planned time does not vary within a treatment, this gave us an estimate of the effect of actual time taken on accuracy, having controlled for planned movement time.

Estimating SAT

As explained in the Results section, we found that it is planned arrival time and not actual arrival time that determines accuracy. Accordingly, for each time constraint in the training session, we recorded the mean arrival time and the corresponding probability of hitting the target. Similarly, for each decay rate in the experimental session, we recorded the mean arrival time and the corresponding probability of hitting the target. This analysis gave us six data points (four from the training session, two from the experimental session) to characterize each subject's SAT. In analyzing the data from the training session, we assumed only that the observer executes the same movement plan when faced with the same time limit. Similarly, in analyzing the data from the experimental session, we assumed only that the subject executes the same plan when faced with the same decay rate. One concern with this assumption is that subjects may have learned over the course of the experiment, and hence, their plans may have changed for a given decay rate. However, as we found no systematic difference in mean movement time or accuracy between early and late trials, we discount this possibility.

We selected the following functional form (McElree & Carrasco, 1999) to characterize the trade-off between speed (mean arrival time) and accuracy (probability of hitting the target):

$$p(\bar{t}) = \beta(1 - e^{-(\bar{t}-\delta)/\lambda}), \tag{3}$$

where β, δ, and λ are estimated parameters. β captures the asymptotic level of p, δ captures the time point where p rises from zero, and λ describes the steepness of the trade-off function. We estimated the parameters β, δ, and λ using maximum likelihood for each subject.[3] Figure 3A gives an example of the SAT function from subject M.A.

Model prediction

Given each subject's estimated SAT, we computed the expected gain $\Gamma(\bar{t})$ for each agent and for each decay rate. The term $\bar{t}$ referred to planned movement time. We then searched $\Gamma(\bar{t})$ to find ($\bar{t}^{\mathrm{MEG}}$, MEG), where $\bar{t}^{\mathrm{MEG}}$ is the movement time that maximizes expected gain and MEG is the maximal value for $\Gamma(\bar{t})$. Figure 3B gives an example of the calculation of EG and MEG for subject M.A. under the slow decay rate.

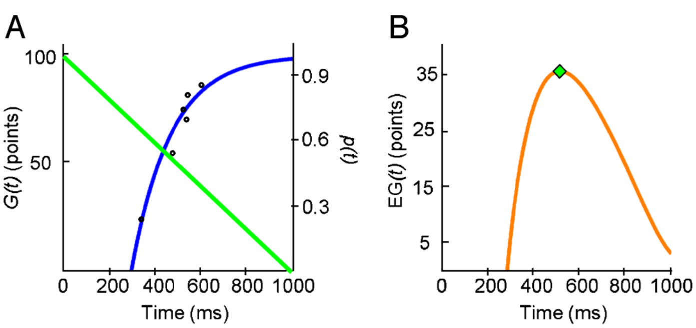

Figure 3. Calculation of maximum expected gain for subject M.A. in the slow decay rate block. (A) The two components of the MEG calculation. The probability of hitting the target was plotted as a function of mean arrival time. Six data points came from the training session (four points) and experimental session (two points). The blue curve was the estimated SAT. Green line: points gained for hitting the target plotted as a function of time in the slow decay rate block. (B) Expected gain, EG, plotted as a function of time. The orange line was EG($\bar{t}$). This line can be thought of as representing the product of the two lines shown in Panel A. The green diamond indicated the maximum.

We emphasize that we rated a subject's performance in terms of efficiency: how close each subject comes to predicted optimal performance based on his or her own SAT and the constraints of the experimental task. With faster movements (<250–300 ms), subjects' measured probabilities of hitting the target dropped to near 0, and given a target whose value dropped to 0 within 250 ms, we would expect subjects to earn very little. They would almost always miss with the occasional lucky hit. Yet, a subject in such a condition may still have a very high efficiency because efficiency is based on a comparison of what the subject wins in the task to the maximum possible winnings with optimal choice of movement time. Conversely, in an easy condition, a subject may win a considerable sum but have very low efficiency because, given his measured SAT, he or she could be expected to win much more with a different choice of movement time. Thus, there is no a priori relation between task difficulty and efficiency.

Efficiency

We defined efficiency as the actual average score a subject achieved divided by the MEG for that subject. We computed efficiency for each subject and each decay rate in the experimental session. We computed the 99% confidence interval of efficiency using bootstrap methods (Efron & Tibshirani, 1993) as follows. We assumed that the distribution of arrival time in each case is a truncated Gaussian ($\bar{t}$, $\sigma_{\bar{t}}$) and estimated the mean and standard deviation from the observed data. Knowing ($\bar{t}$, $\sigma_{\bar{t}}$) and the probability of hit estimated from the experiment, we simulated 10,000 runs of the experiment with each run consisting of 10,000 trials. For each run, we computed the average score of the simulated experiment. As a result, we obtained 10,000 average simulated scores. We computed 10,000 bootstrap estimates of the parameter set for the SAT function and performed EG computation to search for MEG. As a result, we obtained a distribution of MEG (10,000 replications per estimate). We randomly selected one average simulated score and MEG to compute efficiency and repeated this operation for 10,000 times to obtain the 99% confidence interval.

Results

The effect of planned time and actual time on spatial accuracy

Table 1 shows the results of the regression of accuracy on planned movement time, whereas Table 2 shows the

	Coefficient	*SE*	*p*	95% Confidence interval
1	−0.0384	0.0093	.02*	−0.0643, −0.0125
2	−0.0317	0.0092	.03*	−0.0573, −0.0062
3	−0.0026	0.0023	.31	−0.0089, 0.0036
4	−0.0318	0.0058	.01*	−0.0479, −0.0158
5	−0.0078	0.0040	.12	−0.0188, 0.0032
6	−0.0211	0.0025	.00*	−0.0280, −0.0141
7	−0.0244	0.0065	.02*	−0.0426, −0.0063
8	−0.0315	0.0040	.00*	−0.0426, −0.0204
9	−0.0225	0.0064	.03*	−0.0404, −0.0046
10	−0.0376	0.0048	.00*	−0.0511, −0.0242
11	−0.0609	0.0226	.06	−0.1238, 0.0019
12	−0.0201	0.0063	.03*	−0.0375, −0.0027
13	−0.0197	0.0093	.10	−0.0455, 0.0060
14	−0.0109	0.0054	.12	−0.0259, 0.0042
15	−0.1914	0.0618	.04*	−0.3629, −0.0198
16	−0.0078	0.0042	.13	−0.0194, 0.0038

Table 1. Ordinary least squares results for the regression of accuracy on planned movement time for each subject. The dependent variable was the mean distance of movement endpoint from the target within each treatment. This was regressed on the mean movement time within each treatment and a constant. The four treatments from the training session and the two treatments from experimental session provided six data points for each subject. The table reports the estimated coefficient on mean movement time, the standard error of the estimate, the probability level at which one can reject the hypothesis that the coefficient is equal to zero, and the 95% confidence interval for the coefficient. *$p < .05$.

results of the regression of accuracy on actual movement time, controlling for planned movement time. Although not unanimous, results of the regressions indicated that changes in planned movement time have an important effect on accuracy, whereas changes in actual movement time *conditional* on planned movement time did not. Ten of the 16 subjects showed a relationship between accuracy and planned movement time, which was significant at the 5% level. In comparison, only five subjects showed a statistically significant relationship between accuracy and actual movement time once planned movement time had been controlled for. This is true despite the fact that there were more observations (by two orders of magnitude) for the actual movement time regression than for the planned movement time regression (560 trials vs. 6 treatments). Furthermore, the average of the estimated coefficient on planned movement time across subjects was much larger than that for actual movement time (−0.035 vs. −0.005). Thus, we take this as evidence to support our model in which accuracy is determined only by planned movement time.[4]

Estimating SAT

All subjects' estimated SATs are provided in the supplementary document.

Model comparison

Condition A

In Figure 4A, we plotted mean arrival time against MEG timing in Condition A. If subjects were close to optimal in their timing, the points should lie on the 45° diagonal line. We regressed actual timing against MEG timing and found that the slope of the regression was significantly different from zero but not significantly different from 1 at the 5% level, whereas the intercept was not significantly different from zero at the same level. The first of these results indicated that MEG timing was significantly and positively related to actual timing, whereas the second indicated that we cannot reject the hypothesis that these points were distributed around the 45° line. Although most subjects were fairly close to being optimal, we observed that subjects tended to be slightly slower than predicted as most points lay above the line. This was particularly true for the fast decay rate, with the extent to which subjects were slower than the model prediction was more pronounced in the fast decay rate in six out of eight subjects. This suggested that, for most subjects, the change in planned timing between the slow and fast decay rates was smaller than predicted by the model. Figure 4B plotted the observed *change* in average arrival time between fast and slow decay rate trials against the change in MEG timing across subjects. Again, as we regressed actual on predicted change, we cannot reject the hypothesis that the slope was 1 and the intercept was 0 at the 5% level.

	Coefficient	*SE*	*p*	95% Confidence interval
1	−0.0038	0.0029	.18	−0.0095, 0.0018
2	−0.0137	0.0027	.00*	−0.0191, −0.0083
3	−0.0034	0.0018	.06	−0.0069, 0.0002
4	−0.0147	0.0028	.00*	−0.0202, −0.0092
5	−0.0036	0.0029	.21	−0.0092, 0.0021
6	−0.0035	0.0025	.17	−0.0084, 0.0015
7	−0.0074	0.0026	.01*	−0.0126, −0.0022
8	−0.0039	0.0042	.36	−0.0121, 0.0044
9	−0.0045	0.0028	.12	−0.0100, 0.0011
10	−0.0058	0.0032	.07	−0.0121, 0.0005
11	−0.0004	0.0065	.95	−0.0132, 0.0124
12	−0.0124	0.0037	.00*	−0.0197, −0.0052
13	−0.0044	0.0033	.18	−0.0108, 0.0021
14	−0.0102	0.0019	.00*	−0.0140, −0.0065
15	−0.0086	0.0092	.35	−0.0095, 0.0267
16	−0.0038	0.0029	.18	−0.0095, 0.0018

Table 2. Ordinary least squares results for the regression of accuracy on actual movement time controlling for planned movement time. The dependent variable was the distance of movement endpoint from the target for each trial. The dependent variables were the actual time taken for the movement, a set of dummies to indicate which treatment the observation came from, and a constant. Each regression had 560 observations. The reported statistics are as described for Table 1. *$p < .05$.

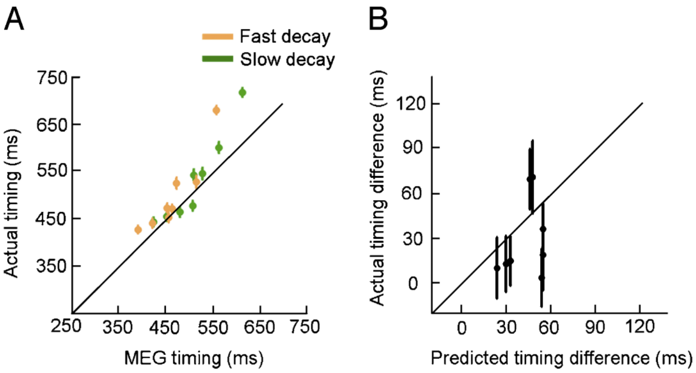

Figure 4. Model comparison for Condition A. (A) Model comparison on timing. Actual mean arrival time was plotted against MEG timing across eight subjects. The decay conditions were color coded. Orange indicated fast decay, and green indicated slow decay. (B) Model comparison on timing difference between the two decay conditions. Actual difference in mean arrival time between the two decay rates was plotted against the MEG timing difference.

Condition B

In Condition B, both the initial value and the slope of the value function increased in fast decay rate trials. The value started at 200 points and decreased 3.3 times faster in the fast rate trials than in the slow rate trials. With this combination of value and slope, MEG was approximately equal between the two rates based on simulations of expected gain obtained from all subjects in Condition A. One of the reasons for running this second condition is that the large difference in decay speeds led to a much larger difference in optimal arrival time between the two decay rates. As shown in Figure 5A, subjects' timing was close to the model prediction for both fast and slow decay rates. Regression analysis again showed that actual time was positively and significantly related to MEG time and that there was no difference in both the slope and the intercept from the diagonal line at the 5% level. As in Condition A, we observed that subjects tended to be slightly slower than predicted. However, subjects did speed up in the fast decay rate blocks in response to the much faster decay rate, as can be seen in Figure 5B. The tendency for slowness was little more marked for the extremely fast decay rates in Condition B than in the fast decay rate of Condition A.

Efficiency

Figure 6 shows subjects' efficiency along with the 99% confidence interval. In Condition A, seven of the eight subjects achieved efficiency, which was indistinguishable from 100% at the 1% confidence level for both decay rates. The point estimates for efficiency were also above 90% for all subjects except for one. It is clear that while subjects tended to be slower than the model prediction for the fast decay rate, this did not seriously affect their performance as most subjects' efficiency was not discernible from optimal.

In Condition B, 100% efficiency lay within the 99% confidence interval for all subjects for the slow decay rate. To our surprise, six of eight subjects were indistinguishable from 100% efficiency in the fast decay rate blocks. This suggested to us that subjects shifted their timing to achieve near-optimal performance even in much more difficult tasks where the optimal speed was close to the fastest they had performed in the training session.

Discussion

In recent years, researchers in areas as diverse as animal foraging (Stephans & Krebs, 1986), perception (Geisler, 1989; Knill & Richards, 1996; Rachlin, Battalio, Kagel, & Green, 1981), and neuroscience (Glimcher, 2002) have increasingly employed decision theoretical or game theoretical models of behavior associated with the literature in economics and mathematical statistics. Many of these models are effective models of optimal allocation of scarce resources (time, effort) and are readily recast as microeconomics. Winterhalder (1983), for example, argued that the economic analysis of optimal allocation in the face of budget constraints provides a framework to model and predict time and energy allocation during foraging behavior, which was absent in most optimization models based on energy rate maximization (see chapter 5

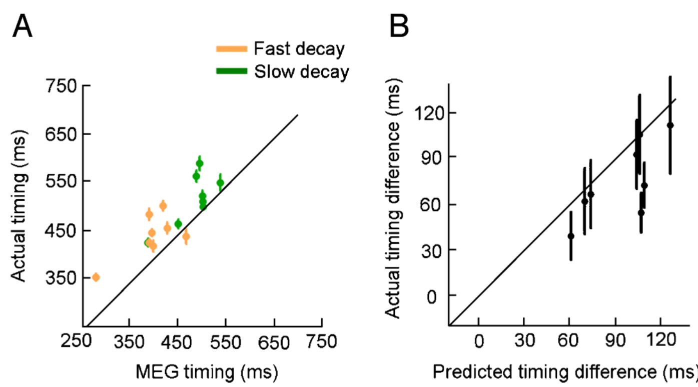

Figure 5. Model comparison for Condition B. (A) Model comparison on timing. Actual mean arrival time was plotted against MEG timing across eight subjects. The decay conditions were color coded. Orange indicated fast decay, and green indicated slow decay. (B) Model comparison on timing difference between the two decay rates. Actual difference in mean arrival time between the two rates was plotted against the MEG timing difference.

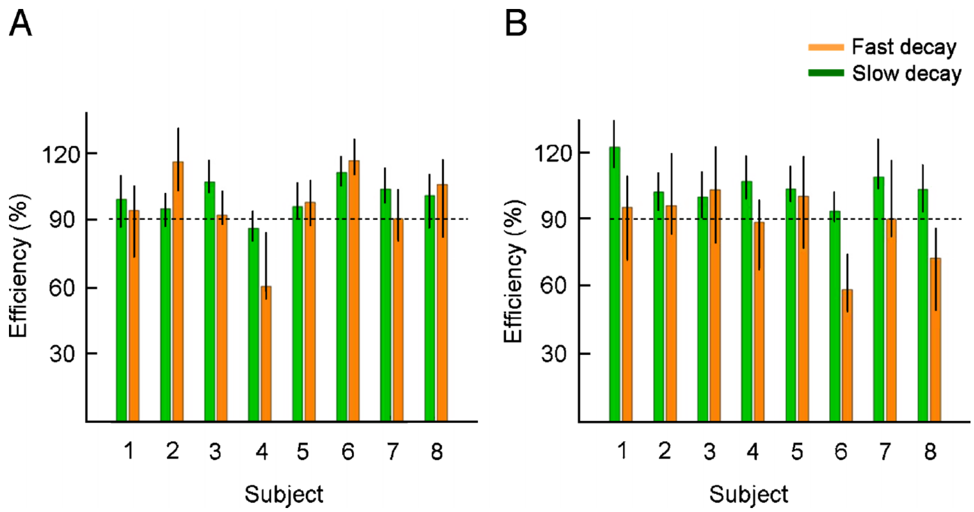

Figure 6. Efficiency. (A) Performance efficiency for subjects in Condition A. (B) Efficiency for subjects in Condition B. Efficiency is defined as the actual average score divided by the MEG. For each decay condition and subject, we computed efficiency and its 99% confidence interval using the bootstrap method (Efron & Tibshirani, 1993).

of Stephens & Krebs, 1986). In neuroscience, researchers seek to determine how the nervous system makes decisions and how trade-offs are represented and resolved (Dorris & Glimcher, 2004; Platt & Glimcher, 1999; Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003).

In this article, we employed similar methods to model and investigate the trade-off between speed and accuracy in simple economic tasks that involved pointing movements. In our experiment, each subject attempted to touch small targets whose value decreased linearly to zero over time ("time to decay"). The subjects earned a reward only if they hit the target, and the amount of reward depended on the duration of the movement. The faster the subject moved, the smaller the chances of hitting the target. The slower the subject moved, the less reward he or she would earn by hitting the target. The movement duration that maximized expected gain depended on the subjects' intrinsic SAT for this task as measured by their probability of hitting the target as a function of movement duration. In a preliminary part of the experiment, we estimated the SAT curve for each subject. With this information, we could predict the choice of movement time that would maximize expected gain for each subject for each choice of decay rate. We measured performance for each subject for two different decay rates. We compared human performance to an optimal movement planner that takes into account his or her own SAT and the reward function to select a movement strategy that maximizes expected reward.

Our work in the time domain is analogous to earlier work in the spatial domain (e.g., Körding & Wolpert, 2004; Trommershäuser et al., 2003a, 2003b; Wu et al., 2006) in the sense that all have expanded the use of statistical decision theory to assess the capability and the limitations of the human motor system. Our work is also a natural extension of previous work on SATs. SATs in human movement have been extensively studied (Meyer et al., 1988). In much of this work, the experimenter either varies speed as an independent variable and measures accuracy (Schmidt et al., 1979) or specifies accuracy and measures the maximum achievable speed of movement (Fitts, 1954; Fitts & Peterson, 1964). In this study, we go beyond this previous literature by challenging the subject to select the optimal trade-off between speed and accuracy in a simple economic task.

Subjects' choices of movement time for each decay rate came close to the optimal SAT for that rate. In Condition A of the experimental session, we tested eight subjects with decay times of 1,000 and 770 ms (decay time was constant in any one block of trials). Subjects selected movement durations that were close to optimal, although they seemed to move slightly slower than predicted, particularly under the fast decay rate. Although subjects were consistently slower than optimal, the consequences for their earnings were slight. We quantified each subject's efficiency as the ratio of actual earnings to the maximum expected gain possible given the subject's measured SAT, separately for each decay time. Efficiencies were high: ranging from 62% to 119% with a median of 101.5% (Figure 6A). We then ran a second condition (Condition B) with a separate group of subjects, in which we kept the slow decay time at 1,000 ms but decreased the faster decay time to 606 ms and doubled the initial value for fast decay trials. By doubling the initial value, we roughly equated the maximum expected gain in the fast and slow decay blocks of Condition B. We found that, while subjects continued to respond slightly slower than the optimal time predicted by our model, the faster decay rate did not lead to reduced efficiency. In fact, six of the eight subjects in the fast decay trials in Condition B achieved efficiency indistinguishable from 100%.

Our first conclusion is that subjects chose movement times that came close to maximizing expected gain for all choices of decay rates, in both the relatively benign conditions of Condition A and the more challenging conditions of Condition B. In more detailed analyses, we also found that planned movement time, not actual movement time, determined accuracy. To our knowledge, we are the first to demonstrate that planned movement time, not actual movement time, determines the trade-off between speed and accuracy.

Although most subjects' performance was close to optimal, subjects shared one clear pattern of deviation from optimal performance—the slight but systematic slowness in movement time compared to the optimal solution that we described above. We advance two possible conjectures for this discrepancy. The first conjecture in purely economic terms is that subjects were exhibiting risk aversion. Risk aversion refers to the willingness to trade off lower expected reward for an increase in the probability of a good outcome. It is well known that humans are typically risk averse (see, e.g., Holt & Laury, 2002). In our task, subjects had a higher probability of hitting the target (and hence a lower chance of winning nothing) when the movement was slower. Thus, the slight but consistent deviations from optimal timing that we found may suggest that our subjects were simply risk averse. They were sacrificing some of their expected reward to pay to increase the chances of winning something on each trial.

Alternatively, the same pattern of deviations could be explained in motor terms as a trade-off between monetary gain and biomechanical cost. Several groups of researchers have proposed that the motor system seeks to minimize specific biomechanical costs including muscle tension change (Dornay, Uno, Kawato & Suzuki, 1996), joint mobility (Soechting & Lacquaniti, 1981), mean torque change (Uno, Kawato, & Suzuki, 1989), and rate of change in acceleration (Flash & Hogan, 1985). Trommershäuser et al. (2003a, 2003b) proposed that individuals would trade off expected monetary gain for a decrease in biomechanical cost ("less gain, less pain"). If, in our experiment, the biomechanical costs incurred by faster movements were higher than those for slower movements and if the motor system were willing to "pay for" reduced biomechanical costs, we would expect to overestimate movement speed in our experiment.

We can compute a "biomechanical premium" directly from subjects' data. We found that subjects in Condition A were, on average, slower than predicted by 23 ms in the slow decay trials and 35 ms in the fast decay trials. The consequent reduction in expected gain (the "premium") averaged across subjects was 0.72 points (slow decay) and 1.21 points (fast decay), on average, per trial. For Condition B, the corresponding premiums averaged across subjects were 0.81 points (slow decay) and 5 points (fast decay) per trial. Across conditions and subjects, the cost of the premium was, on average, 1.9 points (0.1 cents) per trial. This suggested that the subjects were willing to pay the highest premium to "slack off" in the fast decay trials of Condition B.

Either the risk aversion or the biomechanical premium conjecture or both together could account for the systematic deviations from optimal performance we observed. Both possibilities deserve further investigation as both are part of a model of human movement planning, which proposes that even the simplest action is the result of a rapid, sophisticated, nearly optimal evaluation of, and trade-off among, the likely costs and benefits associated with the possible outcomes of a movement.

Acknowledgments

This study was supported by Grant EY08266 from the National Institutes of Health. We thank Frederick Hansen and Joseph Duke from the Technical Fabrication Facility, part of the New York University Physics Department, for designing and constructing the apparatus.

We thank Katja Doerschner for designing the icon.

Commercial relationships: none.
Corresponding author: Mark Dean.
Email: mark.dean@nyu.edu.
Address: Department of Economics, New York University, 19 West 4th Street, 5th Floor, New York, NY 10032.

Footnote

[1] We use the terms "arrival time" and "movement time" interchangeably in this article, although one might more correctly think of arrival time as being made up of a "reaction time" prior to movement onset followed by a "movement time." In fact, reaction time was almost constant across treatments for any given subject; hence, an alternative analysis based on postreaction movement time would yield the same result.

[2] As the "bias" in our data (difference between the mean of movement endpoints and target location) was small, our analysis would not be significantly altered by instead using the "variable error" or the standard deviation of movement endpoints within a treatment.

[3] We place restrictions on the parameter space such that $\beta \leq 1$, $\delta \geq 50$, and $\lambda \geq 50$ to ensure a reasonable shape for the SAT. All our estimates lay on the interior of the restricted parameter space.

[4] Note that the standard errors for the regression coefficients for planned arrival time were smaller than those for actual arrival time, and thus, it is unlikely that our results were driven by lack of variation in actual arrival time conditional on planned arrival time. In general, the

variance of average time taken across treatments was less than twice that of the variance of time taken within treatments.

References

Augustyn, J. S., & Rosenbaum, D. A. (2005). Metacognitive control of action: Preparation for aiming reflects knowledge of Fitts's law. *Psychonomic Bulletin & Review, 12,* 911–916. [PubMed]

Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis* (2nd ed.). New York: Springer.

Blackwell, D., & Girshick, M. A. (1954). *Theory of games and statistical decisions*. New York: Wiley.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436. [PubMed]

Crossman, E. R., & Goodeve, P. J. (1983). Feedback control of hand-movement and Fitts' law. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 35,* 251–278. [PubMed]

Dornay, M., Uno, Y., Kawato, M., & Suzuki, R. (1996). Minimum muscle-tension change trajectories predicted by using a 17-muscle model of the monkey's arm. *Journal of Motor Behavior, 2,* 83–100. [PubMed]

Dorris, M. C., & Glimcher, P. W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron, 44,* 365–378. [PubMed] [Article]

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. [PubMed]

Fitts, P. M., & Peterson, J. R. (1964). Information capacity of discrete motor responses. *Journal of Experimental Psychology, 67,* 103–112. [PubMed]

Flash, T., & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience, 5,* 1688–1703. [PubMed] [Article]

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review, 96,* 267–314. [PubMed]

Glimcher, P. (2002). Decisions, decisions, decisions: Choosing a biological science of choice. *Neuron, 36,* 323–332. [PubMed] [Article]

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review, 92,* 1644–1655.

Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. Cambridge, UK: Cambridge University Press.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature, 427,* 244–247. [PubMed]

Maloney, L. T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world: Psychological and philosophical issues in perception* (pp. 145–189). New York: Wiley.

McElree, B., & Carrasco, M. (1999). The temporal dynamics of visual search: Evidence for parallel processing in feature and conjunction searches. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1517–1539. [PubMed]

Meyer, D. E., Abrams, R. A., Kornblum, S., Wright, C. E., & Smith, J. E. (1988). Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review, 95,* 340–370. [PubMed]

Meyer, D. E., Smith, J. E. K., Kornblum, S., Abrams, R. A., & Wright, C. E. (1990). Speed-accuracy trade-offs in aimed movements: Toward a theory of rapid voluntary action. In M. Jeannerod (Ed.), *Attention and Performance XIII* (pp. 173–226). Hillsdale, NJ: Lawrence Erlbaum.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. [PubMed]

Plamondon, R., & Alimi, A. M. (1997). Speed/accuracy trade-offs in target-directed movements. *Behavioral and Brain Sciences, 20,* 279–303. [PubMed]

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature, 400,* 233–238. [PubMed]

Rachlin, H., Battalio, R., Kagel, J., & Green, L. (1981). Maximization theory in behavioral psychology. *Behavioral and Brain Sciences, 4,* 371–388.

Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science, 300,* 1755–1758. [PubMed]

Schmidt, R. A., Zelaznik, H., Hawkins, B., Frank, J. S., & Quinn, J. T. (1979). Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review, 47,* 415–451. [PubMed]

Soechting, J. F., & Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man. *Journal of Neuroscience, 1,* 710–720. [PubMed] [Article]

Stephans, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.

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. [PubMed] [Article]

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A, Optics, image science, and vision, 20,* 1419–1433. [PubMed]

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision, 16,* 255–275. [PubMed]

Uno, Y., Kawato, M., & Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. *Biological Cybernetics, 61,* 89–101. [PubMed]

Winterhalder, B. (1983). Opportunity-cost foraging models for stationary and mobile predators. *American Naturalist, 122,* 73–84.

Wu, S. W., Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2006). Limits to human movement planning in tasks with asymmetric gain landscapes. *Journal of Vision, 6*(1):5, 53–63, http://journalofvision.org/6/1/5/, doi:10.1167/6.1.5. [PubMed] [Article]