Section: Behavioral/Systems/Cognitive Neuroscience
Senior Editor: Dr. John H. R. Maunsell

Trading Off Speed and Accuracy in Rapid, Goal Directed Movements

Abbreviated Title: Optimality and the Speed-Accuracy Trade Off

Mark Dean[1], Shih-Wei Wu[2], Laurence T. Maloney[2,3]

[1]New York University, Department of Economics, 269 Mercer Street, New York, NY 10003, USA;

[2]New York University, Department of Psychology, 6 Washington Place, New York, NY 10003, USA;

[3] New York University, Center for Neural Science, 6 Washington Place, New York, NY 10003, USA;

Under Review, J. Neuroscience, 7/25/2006

Corresponding Author:
Mark Dean
Department of Economics, New York University
269 Mercer Street,
New York, NY 10003, USA
Email: mark.dean@nyu.edu
Phone: +1 212 998 8038, Fax: +1 212 995 3932

Number of Figures: 6
Number of Tables: 2
Number of Pages: 24
Key Words: Visuo-motor control, movement planning, optimality, statistical decision theory, speed/accuracy trade off, decision making.

Acknowledgments: Supported by Grant EY08266 from the National Institute of Health. We thank Frederick Hansen and Joseph Duke from the Technical Fabrication Facility, part of the NYU Physics department, for designing and constructing the apparatus.

Abstract

Many studies have shown that humans face a trade-off between the speed and accuracy with which they can make movements. Previous studies of optimal movement planning have not modeled this decision. In this paper we asked whether humans choose movement time in order to maximize expected gain by taking into account their own speed-accuracy trade-off. 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 times optimally in an environment where they are rewarded for both speed and accuracy; the value of the target decreased linearly over time to zero. We ran two experiments. In the first, each subject faced two different decay rates. Seven out of eight subjects' efficiency could not be differentiated from optimal, though movement time tended to be slightly slower than our model predicted. In the second experiment, we increased the faster decay rate of the first experiment so that being slow became more costly. While we observed the same asymmetry in timing across subjects, six out of eight subjects' performance was close to optimal in the fast decay condition. We concluded that in planning movements, humans take into account their own speed-accuracy trade-off to maximize expected gain.

The aim of this paper is to determine whether subjects trade off speed against accuracy in an optimal way when performing rapid, goal directed movements. 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. Characterizing the fundamental properties of the speed-accuracy trade-off has been a focus of research on human motor performance since Fitts' (1954) original proposal of a logarithmic speed-accuracy trade-off function (Fitts, 1954; Fitts and Peterson, 1964; Meyer et al., 1982; Meyer et al., 1988; Schmidt et al., 1979).

In a series of experiments, Trommershäuser et al. (2003ab, 2005) have 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 extrinsic costs and benefits inherent in the environment. When choosing a strategy, the movement planner takes into account their own intrinsic motor variability. These models do well in predicting movements made by subjects in rapid, goal directed pointing tasks.

Previous studies of optimal movement planning have not explicitly modeled the speed-accuracy trade-off of a subject. The experiments of Trommershäuser et al. effectively fixed the length of time a subject had to perform the task. As time taken varied little from trial to trial, their model treated motor variability as exogenously fixed for each subject[1]. This is, however, an artificial constraint. In most tasks, a person gets to choose how long they take over a movement, and so the accuracy of that movement. The degree of motor variability that a person faces therefore becomes an endogenous choice variable. In an environment in which there are rewards for both quick and accurate movements, finding the optimal point on the speed accuracy trade-off can

[1] Though motor variability was modelled as varying between subjects.

be a non-trivial problem. The aim of this paper is to incorporate this choice into an optimizing model, and compare its predictions to behavior. In short, we wish to answer the question: 'Do people optimally trade off speed against accuracy?'.

Our experimental design has two sections, A and B. In section A, we ran a sequence of treatments in which subjects were rewarded for performing a pointing task within various time limits. In section B, 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 section A and B 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 section B. 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.

1 Materials and Methods

1.1 Overview

The environment in which we explored the trade-off between speed and accuracy is similar to that previously used to examine optimality in motor tasks (Trommershäuser et al., 2003ab). 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 experimental sessions, which we labeled session A and session B. Session A was designed to elicit the subject's speed-accuracy trade-off. Each subject was presented with 4 treatments in which they were asked to perform the pointing task within 4 different time limits. In each treatment, the subjects received a monetary reward if they hit the target within the time limit, but a large monetary penalty if they failed to make contact with the screen before the time limit has expired. Session B was designed to test whether the subjects selected the right point on the speed-accuracy trade-off in environments where there were benefits associated with both quick and accurate movements. We therefore presented each subject with environments in which the reward for successfully hitting the target decreased with the time taken. Unlike session A, 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 treatments - a 'fast' decay in which the value of hitting the target decreased quickly, and a 'slow' decay in which it decreased more slowly. Subjects were divided between two experiments, which varied only in the speed of the fast decay condition. Experiment 1 had a modest difference in speed between the fast and slow conditions. The results from experiment 1 motivated us to run a second experiment, experiment 2, in which the fast decay condition was significantly faster than in experiment 1, to see if subjects could perform optimally on this more difficult task.

1.2 Apparatus

A touch monitor (Elo IntelliTouch 17in. LCD monitor) was mounted vertically on the Structural Framing System (McMaster Carr Inc.). A double-square 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. At the beginning of every experimental session, a calibration procedure on touch location was performed to optimize endpoint accuracy for each subject.

1.3 Stimuli

1.3.1 Session A

The subjects were asked to hit a green target (radius = 7mm / 24 pixel) 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 1.5s and was followed by a presentation of a blue square region (74 mm by 74mm) indicating the possible locations of the target. Shortly after the blue region was presented, the target and a 'time' bar, displayed horizontally on top of the blue region, was presented simultaneously. The length of the time bar decreased to indicate how much time had left before the time limit was reached. To prevent subjects from pre-planning the movement, we perturbed both the timing of target onset and target location. The target was presented after 500ms plus an amount drawn from a uniform distribution with range of 100ms after the onset of the blue region. Target location was perturbed in both the x and y direction with a range of $\pm$23mm relative to the screen center.

1.3.2 Session B

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

1.4 Procedure

There were two sessions in the experiment, session A and B. Each subject took part in both sessions. A subject would participate in the two sessions on separate days but within 48 hours of each other.

In session A, we implemented four time constraints for the task. For 12 of the 16 subjects, these time constraints were 400ms, 525ms, 650ms and 775ms. For 4 of the subjects we used with a slightly different schedule of 400ms, 600ms, 800ms and 1000ms. The timer started upon target presentation and the time that the subject had left in a trial was indicated by a white bar (we called it 'time bar') whose length decreased continuously over time. The bar was horizontally 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. 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. Hitting the screen but missing the target within the time limit earned the subject 0 points. 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 experimental trials. Prior to the session, the experimenter explicitly instructed subjects to be consistent in speed within a block and to use the time as well as they can. 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 longer breaks (3 min) between blocks and shorter breaks (20s) in between trials, especially within the block of shorter time constraints, to minimize the impact of fatigue. Session A took approximately 90 minutes to complete.

In session B, instead of receiving a fixed reward (100 points) for hitting the target, target value decreased linearly in value over time after its presentation. There was no time constraint after which the subject we penalized. We implemented two experiments, with different sets of decay conditions and assigned half of the subjects to each set. Figure 1 summarized the decay conditions of both experiments.

In the first experiment, 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 770ms after its presentation, while target value decreased to zero at 1000ms after presentation in the slow-decay block. In the second experiment, we drastically increased the faster decay rate. To balance incentives, we also increased the initial value to 200 points for the fast-decay condition which resulted in a decrease to zero at 606ms. Subjects started each block with 40 practice trials with no monetary consequences, and performed 120 experimental trials. The order of fast and slow blocks was balanced across subjects. Session B took approximately 45 minutes to complete.

Figure 2 provided a visual depiction of the task, 1A for session A, 1B for session B.

1.5 Subjects and instructions

Sixteen subjects, unaware of experimental purpose, participated. Among them, seven were male and nine were female. Nine of them were graduate students from the Department of Economics in New York University. The remaining subjects were graduate students in Psychology and Neural Science, or students from the Law School. Subjects were not aware of the purpose of the experiment. All but one subject 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 $36 ($24 from session A, $12 from session B) plus the additional bonus they earned in the experiment. Total payment ranged from $40 to $55 across subjects.

1.6 Model of optimal movement planning

Subjects' performance in the task was compared to an optimal movement-planning model based on statistical decision theory (Berger, 1985; Maloney, 2002). The model is an extension of

Trommershäuser et al. (2003ab). Here, we use 'optimal' to refer to the maximization of expected monetary reward.

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

Choose s $\in$ S to maximize

$$\Gamma(s) = \int_G R(g) f(g|s) dg \tag{1}$$

where s is a motor strategy, S is the set of all possible strategies, G is a set of outcomes (i.e. realized movements), $R : G \to \mathbb{R}$ 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 they are 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. (2003ab), a strategy consists of selecting a target point on the screen representing the target for the end point of the movement. Thus a strategy could be represented as a tuple of $\bar{x}$ and $\bar{y}$ target coordinates.

An innovation of this paper 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 $\bar{y}$ 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, the reward that the agent receives for a particular movement will be a function only of the point at which the agent touches the screen (which we will represent as the coordinates x and y) and the time when the touch takes place (which we will represent as t). We can therefore rewrite problem (1) as

$$\text{Choose } \bar{x} \in (-\infty, \infty),\ \bar{y} \in (-\infty, \infty) \text{ and } \bar{t} \in [t^*, \infty) \text{ to maximize}$$

$$\Gamma(\bar{x}, \bar{y}, \bar{t}) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{t^*}^{\infty} R(x, y, t) f(x, y, t | \bar{x}, \bar{y}, \bar{t}) dt\ dx\ dy \tag{2}$$

where t^* is some lower bound on the time taken. Note that the *actual* time taken to complete a movement is also stochastic: the agent gets to choose a *target* movement time, but the actual time taken will vary randomly due to motor uncertainty.

In the case of the specific environment dealt with in this experiment we can simplify the problem further. First, the choice of x and y is trivial. Assuming that the marginal distributions of x and y about $\bar{x}$ and $\bar{y}$ are symmetric, then the best one can do is to aim for the center of the target. Thus we do not model the choice of $\bar{x}$ and $\bar{y}$ explicitly. Second, the location of the end point of the movement only matters to the extent to which it falls within or without the target circle. Let A be the event that the endpoint of the movement falls within the target circle. We can once again rewrite the objective function of the agent as:

$$\Gamma(\bar{t}) = \int_{t^*}^{\infty} \left(R(A,t)p(A|t,\bar{t}) + R(A^c,t)p(A^c|t,\bar{t}) \right) f(t|\bar{t})dt \tag{3}$$

where A^c is the event that the endpoint of the movement falls outside the target circle, and $p(.|t,\bar{t})$ is the probability of an event given that the target time for the movement is $\bar{t}$ and the actual time taken is t.

So far, the model we have written down allows the probability of the agent hitting the circle to depend both on the target time that the subject chooses and the actual time taken to perform the movement. As we will report in the Results section, we found that the probability of hitting depends only on the *planned* movement time. In other words, the stochastic error in timing does not affect spatial accuracy. Thus we can replace $p(A|t,\bar{t})$ with $p(A|\bar{t})$.

In order to be able to solve for the agent's optimal choice of $\bar{t}$ in session B we need to determine the nature of $p(A|\bar{t})$ and $f(t|\bar{t})$, both of which we do on an individual-by-individual basis. We estimated the function $p(A|\bar{t})$ using the data gathered from both session A and session B. In order to do so, we make two assumptions: First, we assume that $E(t|\bar{t}) = \bar{t}$, 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. Formally, index trials by i and let $J(x) \subseteq \mathbb{N}$ be the trials for which the treatment is x. We assume that $\bar{t}_i = \bar{t}(x)$ for $\forall\ i \in J(x)$. Under these two assumptions, we can identify for each of the 6 treatments in sessions A and B the following two values:

$$\bar{t}(x) \approx \frac{1}{|J(x)|}\sum_{i\in J(x)} t_i(x) \tag{4}$$

$$p(A|\bar{t}(x)) \approx \frac{\sum_{i\notin J(x)} A_i}{|J(x)|} \tag{5}$$

where A_i is an indicator variable which is equal to 1 if the outcome of trial i is A.

Thus sessions A and B provide 6 points on the speed-accuracy trade-off. We use these points to estimate the full trade-off using a cumulative Weibull function with three parameters that characterize different properties in the speed-accuracy trade-off. This is a standard model commonly used in the reaction time literature (see, e.g. McElree and Carrasco, 1999).

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

We are now in a position to calculate the optimal choice of $\bar{t}$ in session B. In this experiment, the reward for missing the circle is always zero, while the reward schedule for hitting the circle, $R(A,t)$, varies between treatment. 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 $f(t|\bar{t})$ we find the optimum using numerical techniques.

1.7 Data analysis

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

1.7.1 The effect of planned time and actual time on spatial accuracy

In order to determine the optimal arrival time for a subject, we needed to determine the exact nature of the speed-accuracy trade-off that the subject faces. To begin with, we needed to determine whether it was *planned* or *actual* movement time which determines accuracy. In order to do so, we analyzed how spatial accuracy is 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 session A and decay condition in session B, we computed mean distance between movement end points and target center and mean response 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.

1.7.2 Estimating speed-accuracy trade-off (SAT)

As we explain in the Results section below, we found that it is planned rather than actual arrival time which determines accuracy. We therefore used our data to estimate a relationship between the probability of hitting the target and planned arrival time. To do so, for each time constraint in session A and each decay condition in session B, we computed the mean arrival time and the probability of hitting the target. As a result, each subject had six data points (4 from session A, 2 from session B). Note that we relied on the assumption that within a time block or decay condition, subjects used the same planned arrival time, and that the mean arrival time reflected planned time. We selected the following functional form (McElree and Carrasco, 1999) to characterize the trade-off between speed (mean arrival time) and accuracy (probability of hitting the target):

$$p(\bar{t}) = \beta(1 - e^{-(\bar{t}-\delta)/\lambda}) \tag{6}$$

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, λ using maximum likelihood for each subject[2].

Figure 3A gave an example of the SAT function from subject MA.

[2] We place restrictions on the parameter space such that $\beta \leq 1$, $\delta \geq 50$ and $\lambda \geq 50$ in order to ensure a reasonable shape for the speed-accuracy trade-off. All our estimates lie on the interior of the restreicted parameter space.

1.7.3 Model prediction

Given each subject's estimated speed-accuracy trade-off, we computed $EG_i^j(t)$ for each agent and decay condition in session B. This refers to the expected reward of agent i choosing planned movement time t in decay condition j. We then searched $EG_i^j(t)$ to find $(\bar{t}_{ij}^{MEG}, MEG_i^j)$, where $\overline{t_{ij}}^{MEG}$ is the target movement time which maximizes expected gain $EG_i^j(t)$, and MEG_i^j is the maximal value for $EG_i^j(t)$. Figure 3B gave an example of the calculation of EG and MEG for subject MA in the slow decay condition. .

1.7.4 Efficiency

We defined efficiency as the actual average score a subject achieved divided by MEG for that subject. We computed efficiency for each subject and each decay condition in session B. We computed the 99% confidence interval of efficiency using Bootstrap methods (Efron and Tibshirani, 1993) as follows. We assumed that the distribution of arrival time in each condition is a truncated Gaussian $(\bar{t}, \sigma_{\bar{t}})$ and estimated the mean and standard deviation from choice 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.

2 Results

2.1 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, while Table 2 shows the results of the regression of accuracy on actual movement time, controlling for planned movement time. While not unanimous, the regressions suggested that changes in planned movement time have an important effect on accuracy, while changes in actual movement time *conditional* on planned movement time did not. 10 of the 16 subjects showed a relationship between accuracy and planned movement time which is significant at the 5% level. In comparison, only 5 subjects show a statistically significant relationship between accuracy and actual movement time once planned movement time has been controlled for. This is true despite the fact that there are two orders of magnitude more observations 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.

2.2 Estimating speed-accuracy trade-off (SAT)

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

2.3 Model comparison

Experiment 1. Figure 4A plotted mean arrival time against MEG timing in Experiment 1. If subjects were close to optimal in their timing, the points should lie on the 45 degree diagonal line. We regressed actual timing against the 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, while the intercept was not significantly different to zero at the same level. The first of these results means that MEG timing is significantly and positively related to actual timing, while the second means that we cannot reject the hypothesis that these points are distributed around the 45 degree line. While 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 condition, with the extent to which subjects were slower than the model prediction was more pronounced in the fast decay condition in 6 out of 8 subjects. This meant that, for most subjects, the change in planned timing between the slow and fast decay conditions was smaller than predicted by the model.

Experiment 2. In Experiment 2, we made the value decrease 3.3 times faster in the fast condition than the slow condition while controlling for incentives by raising the initial value to 200 points. Based on simulations obtained from all subjects in Experiment 1, MEG was approximately equal between the two conditions for experiment 2. One of the reasons for running this second experiment is that the large difference in decay speeds led to a much larger difference in optimal arrival time between the two conditions. As shown in Figure 5A, subjects' timing was close to the model prediction across conditions. Regression analysis again shows both that actual time is 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 Experiment 1, we observed that subjects tended to be slightly slower than predicted. However, subjects did speed up in the fast decay condition in response to the much faster decay rate, as can be seen in Figure 5B. The tendency for slowness was little more marked in the extremely fast decay condition in Experiment 2 than in the fast decay condition of Experiment 1.

2.4 Efficiency

Figure 6 showed subjects' efficiency along with the 99% confidence interval. In Experiment 1, 7 of the 8 subjects achieved efficiency which was indistinguishable from 100% at the 1% confidence level in both conditions. The point estimates for efficiency were also above 90% for all but one subjects. It is clear that while subjects tended to be slower than the model prediction in the fast decay condition, this did not seriously affect their performance as most subjects' efficiency was not discernible from optimal.

In Experiment 2, 100% efficiency lay within the 99% confidence interval for all subjects in the slow decay condition. To our surprise, six of eight subjects were indistinguishable from 100% efficiency in the fast condition. This suggested to us that subjects shifted their timing to achieve near-optimal performance even in much more difficult tasks where the optimal speed is close to the fastest they have performed in session A.

3 Discussion

In this study, we have tried to determine whether humans choose to trade off speed of movement against accuracy of movement in an optimal way. This aspect of movement planning, which is clearly important in a host of everyday movement tasks, has not been explored in previous studies testing optimality of human motor control (Körding and Wolpert, 2004; Miyazaki et al., 2005; Trommershäuser et al., 2003ab, 2005), nor in studies of speed-accuracy trade-off (for review, see Meyer et al., 1988).

Underlying our work is the assumption that everyone is endowed with an individual-specific constraint which relates the speed at which s/he can make a movement and the accuracy of that movement. People can then choose to trade off speed against accuracy by choosing a point on this constraint. This can be done by either choosing a speed of movement, or by choosing a desired accuracy. Assuming that subjects operate on their constraint, the choice of one will determine the other.

Our first task was to determine the nature of this constraint in our pointing task. Past research on speed-accuracy trade-off in rapid movements have focused on (1) characterizing the fundamental feature of the trade-off (Fitts, 1954; Fitts and Peterson, 1954; Schmidt et al., 1979; Zelaznik et al., 1988), (2) building models on the grounds of neurophysiology and biomechanics that could explain the empirical observations of speed-accuracy trade-off (Keele, 1968; Kvalseth, 1974; Meyer et al., 1982; Meyer et al., 1988; Meyer et al., 1990; Schmidt et al., 1979) and (3) examining the classical Fitts Law under different experimental constraints or different types of movements (Hoffmann, 1991a; Keele and Posner,1968; MacKenzie et al., 1987; Murata and

Iwase, 2001; Plamondon and Alimi, 1997).

In previous experimental investigations of the relationship between speed and accuracy of movement, subjects have either been asked to move as quickly as possible with high spatial accuracy to a target (spatially constrained tasks) or to move to the target at a pre-specified time (temporally constrained tasks). In the former, mean movement time has been found to be a logarithmic function of movement distance divided by target width. In the latter, the relation between speed and accuracy appears to be linear. This was potentially something of a worry for our study, as we assume that there is a single constraint, which should be invariant to this change in emphasis. However, Meyer et al. (1988) showed how both of these experimental findings are consisted with a single biomechanical trade-off. For our own study, we relate the time taken to perform a task to the probability of hitting a target using a cumulative Weibull function with three parameters. This choice allowed for a flexible, potentially non-linear relationship.

In determining the relationship between speed and accuracy, we have differentiated between *planned* movement time and *actual* movement time. This difference comes about because, in our model, the time that a particular movement takes is itself a random variable with a mean determined by the planned time for the movement. Our analysis indicates that it is planned movement time which determines accuracy. For most subjects, planned time is positively related to accuracy, while actual time is not related to accuracy once planned time is controlled for. 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.

Having determined individual-specific speed-accuracy relationships, we were in a position to compare optimal movement time to actual movement time in an environment which rewarded both speed and accuracy of movement. Each subject performed the pointing task in a fast and slow decay condition. In Experiment 1, we tested 8 subjects in a paradigm in which the difference between the fast and slow conditions was modest: the value for hitting the target decreased 1.3 times faster in the fast decay condition. Our main result gave strong support to the hypothesis that subjects were selecting movement times which were very close to optimal. However subjects seemed to move slightly slower than predicted. This was particularly true in the fast decay condition, which meant that subjects changed their planned movement time slightly less than predicted between treatments. It was also true that, for the estimated speed accuracy relationship for most subjects, the modest difference in decay speeds lead to only a modest difference in optimal movement time between treatments. These two factors led us to run a second experiment, Experiment 2, in which we increased the speed of the faster decay condition to 3.3 times that of the slower condition. We found that, while subjects remained slightly slower than suggested by our model, the faster decay condition did not exacerbate the problem. In fact, 6 of the 8 subjects in the fast decay condition in Experiment 2 achieved efficiency indistinguishable from 100%. Moreover, the changes in planned movement time between treatments was much more marked in Experiment 2 than in Experiment 1, in line with our model. This suggests to us that, subjects choose movement time broadly in line with optimizing principles both in the relatively benign conditions of Experiment 1 and the more challenging conditions of Experiment 2.

4 References

Berger JO (1985) Statistical decision theory and Bayesian analysis, 2nd ed. NY: Springer.

Fitts PM (1954) The information capacity of the human motor system in controlling the amplitude of movement. J Exp Psychol 47:381-391.

Fitts PM, Peterson JR (1964) Information capacity of discrete motor responses. J Exp Psychol 67:103-112.

Hoffmann ER (1991a) Capture of moving targets: A modification of Fitts' law. Ergonomics 34:211-220.

Keele SW (1968) Movement control in skilled motor performance. Psychol Bull 70:387-403.

Keele SW, Posner MI (1968) Processing visual feedback in rapid movements. J Exp Psychol 77:155-158.

Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244-247.

Kvalseth TO (1974) A preview-constraint model of rotary arm control as an extension of Fitts' law. J Exp Psychol 102:696-699.

MacKenzie CI, Mateniuk RG, Dugas C, Liske D, Eickmeier, B (1987) Three-dimensional movement trajectories in Fitts' task: Implications for control. Q J Exp Psychol A 39:629-647.

Maloney LT (2002) Statistical decision theory and biological vision. In: Perception and the physical world (Heyer D, Mausfeld R, eds), pp145-189. NY: Wiley.

McElree B, Carrasco M (1999) The temporal dynamics of visual search: Speed-accuracy tradeoff analysis of feature and conjunctive searches. J Exp Psychol Hum Percept Perform 25:1517-1539.

Meyer DE, Smith JE, Wright CE (1982) Models for the speed and accuracy of aimed movements. Psych Rev 89:449-482.

Meyer DE, Abrams RA, Kornblum S, Wright CE, Smith JE (1988) Optimality in human motor performance: Ideal control of rapid aimed movements. Psychol Rev 95:340-370.

Meyer DE, Smith JEK, Kornblum S, Abrams RA, Wright CE (1990) Speed-accuracy tradeoffs in aimed movements: toward a theory of rapid voluntary action. In: Attention and Performance XIII, pp173-226. Hillsdale, NJ: Lawrence Erlbaum.

Miyazaki M, Nozaki D, Nakajima Y (2005) Testing Bayesian models of human coincidence timing. J Neurophysiol 94:395-9.

Murata A, Iwase H (2001) Extending Fitts' law to a three-dimensional pointing task. Hum Mov Sci 20:791-805.

Plamondon R, Alimi AM (1997) Speed/accuracy trade-offs in target-directed movements. Behav Brain Sci 20:279-303.

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

Trommershäuser J, Gepshtein S, Maloney LT, Landy MS, Banks MS (2005) Optimal compensation for changes in task-relevant movement variability. J Neurosci 25:7169-7178.

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

Trommershäuser J, Maloney LT, Landy MS (2003b) Statistical decision theory and the selection of rapid, goal-directed movements. J Opt Soc Am A Opt Image Sci Vis 20:1419-1433.

Zelaznik HN, Mone S, McCabe GP, Thaman C (1988) Role of temporal and spatial precision in determining the nature of the speed-accuracy trade-off in aimed-hand movements. J Exp Psychol Hum Percept Perform 14:221-230.

5 Figure and Table Legends

Figure 1. Summary of decay conditions. Target gain, G, (in points) was plotted as a function of time (ms). A. Decay conditions of Experiment 1. Target value started at 100 points but decreased in two different rates. In the slower decay condition, value decreased to zero at 1000ms. In the faster decay condition, value decreased to zero at 770ms. Values started decreasing as soon as the target appeared. B. Decay conditions of Experiment 2. The slow condition had the same decay rate as that in Experiment 1. However, in the fast decay condition, target value initiated at 200 points but decreased 3.3 times faster than the slow condition.

Figure 2. Stimulus configurations. A. Stimulus configuration in session A. Subjects saw the circular target (radius=7mm) 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 session B. Unlike session A, there was no time constraint in session B. Instead, target value decreased rapidly over time. The time bar in session A was replaced by a money bar to provide continuous feedback on value.

Figure 3. Calculation of maximum expected gain for subject MA in the slow decay condition. 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 session A (4 points) and session B (2 points). Blue curve was the estimated speed-accuracy trade-off. Green line: points gained for hitting the target plotted as a function of time in the slow decay condition. B. Expected gain EG, plotted as a function of time. The orange line was $EG(t)$. This line can be thought of as

representing the product of the two lines shown in panel A. The green diamond indicated the maximum.

Figure 4. Model comparison for Experiment 1. A. Model comparison on timing. Actual mean arrival time was plotted against MEG timing across 8 subjects. The decay conditions were color coded. Orange indicated fast decay, green indicated slow decay. B. Model comparison on timing difference between the two decay conditions. Actual difference in mean arrival time between the two conditions was plotted against the MEG timing difference.

Figure 5. Model comparison for Experiment 2. A. Model comparison on timing. Actual mean arrival time was plotted against MEG timing across 8 subjects. The decay conditions were color coded. Orange indicated fast decay, green indicated slow decay. B. Model comparison on timing difference between the two decay conditions. Actual difference in mean arrival time between the two conditions was plotted against the MEG timing difference.

Figure 6. Efficiency. A. Performance efficiency for subjects in Experiment 1. B. Efficiency for subjects in Experiment 2. 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 methods.

Table 1. This table reports the ordinary least squares (OLS) results for the regression of accuracy on planned movement time for each subject. The dependant variable is the mean distance of movement end point from target within each treatment. This is regressed on the mean movement time within each treatment and a constant. The 4 treatments from session A and the 2 treatments from session B provide 6 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.

Table 2. This table reports the OLS results for the regression of accuracy on actual movement time controlling for planned movement time. The dependant variable is the distance of movement endpoint from target for each trial. The dependant variables are 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 above.

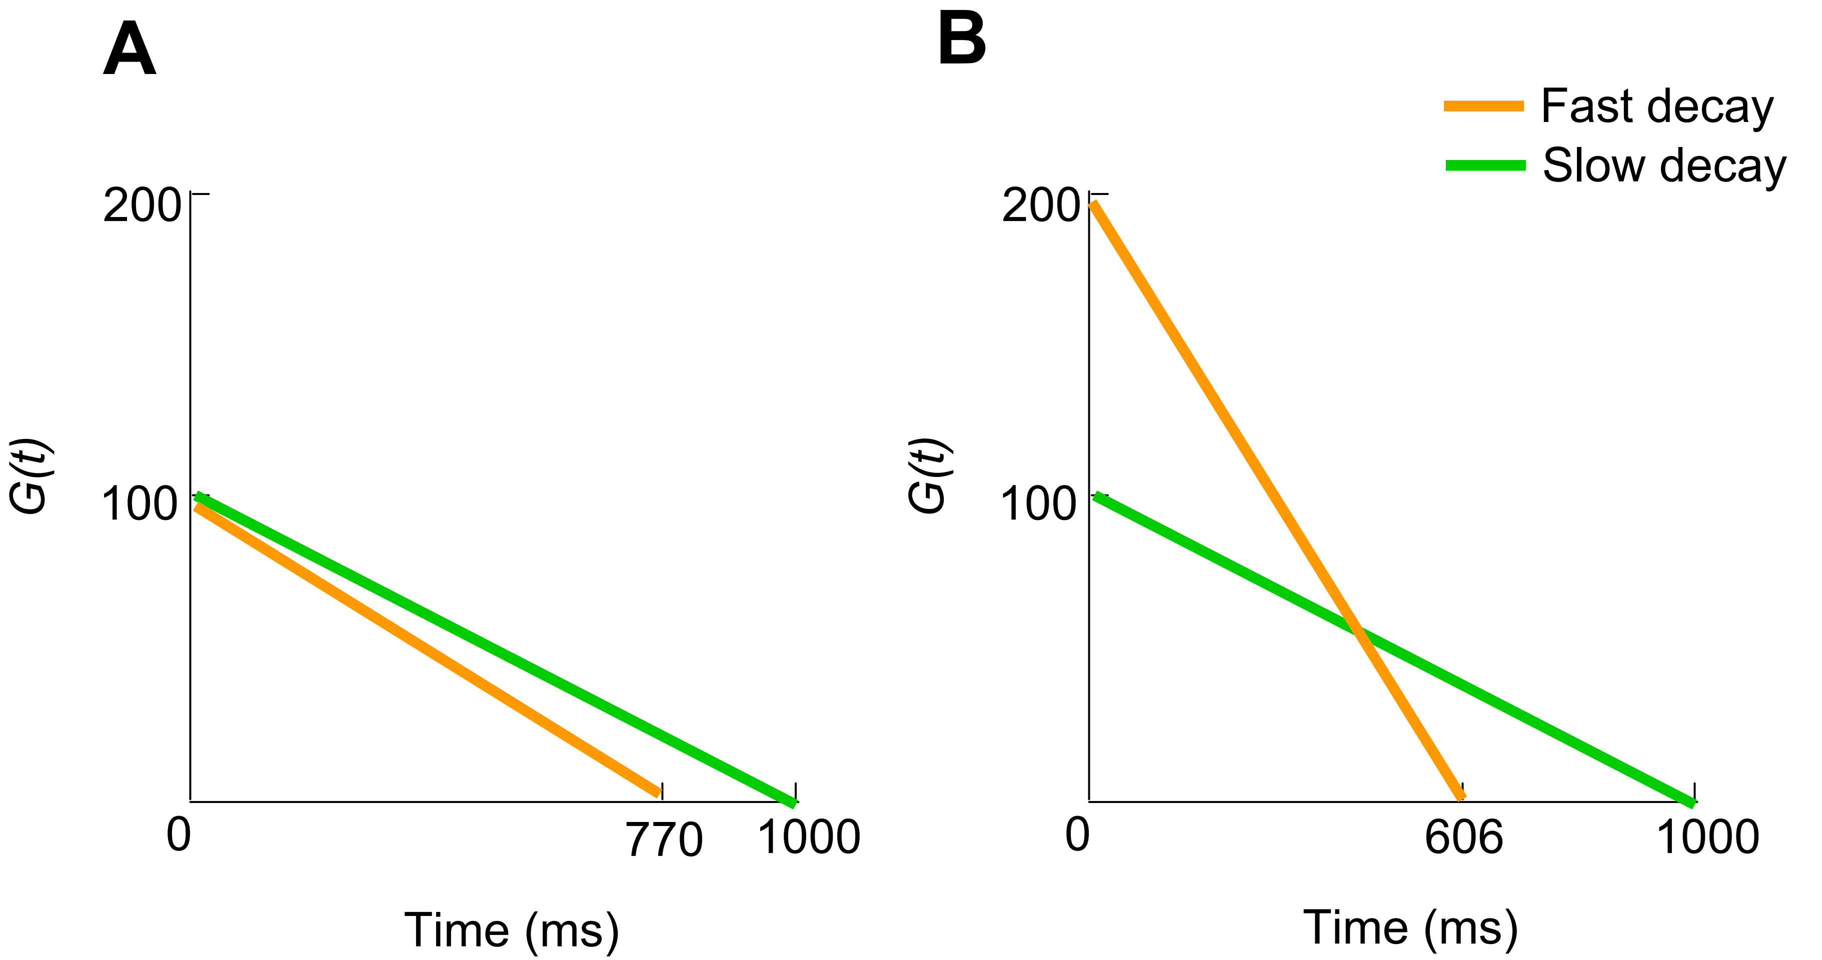

Dean et al: Figure 1

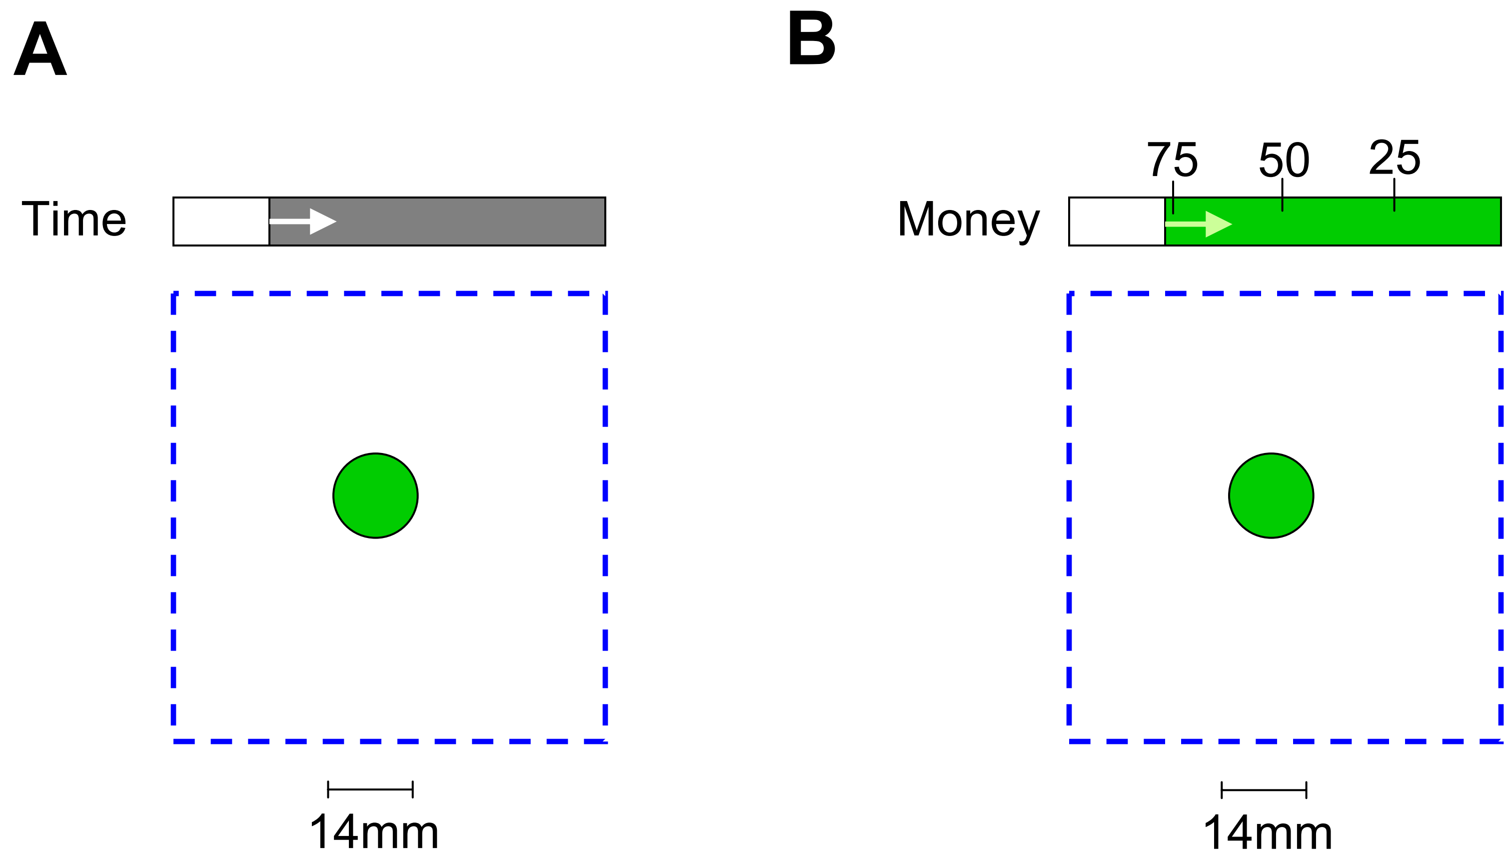

Dean et al: Figure 2

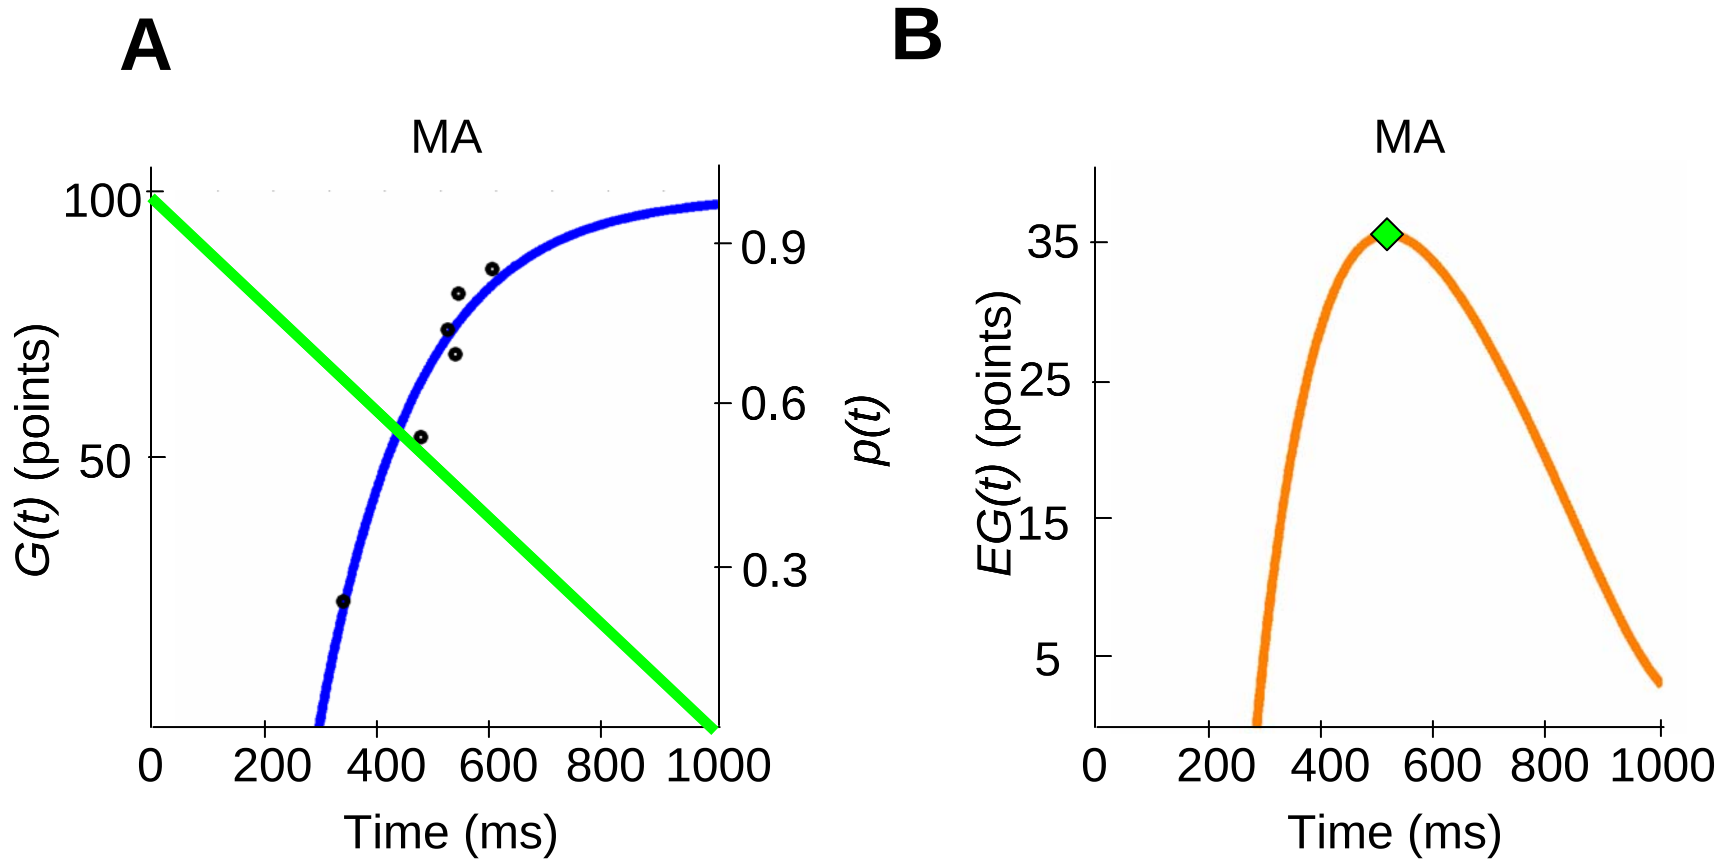

Dean et al: Figure 3

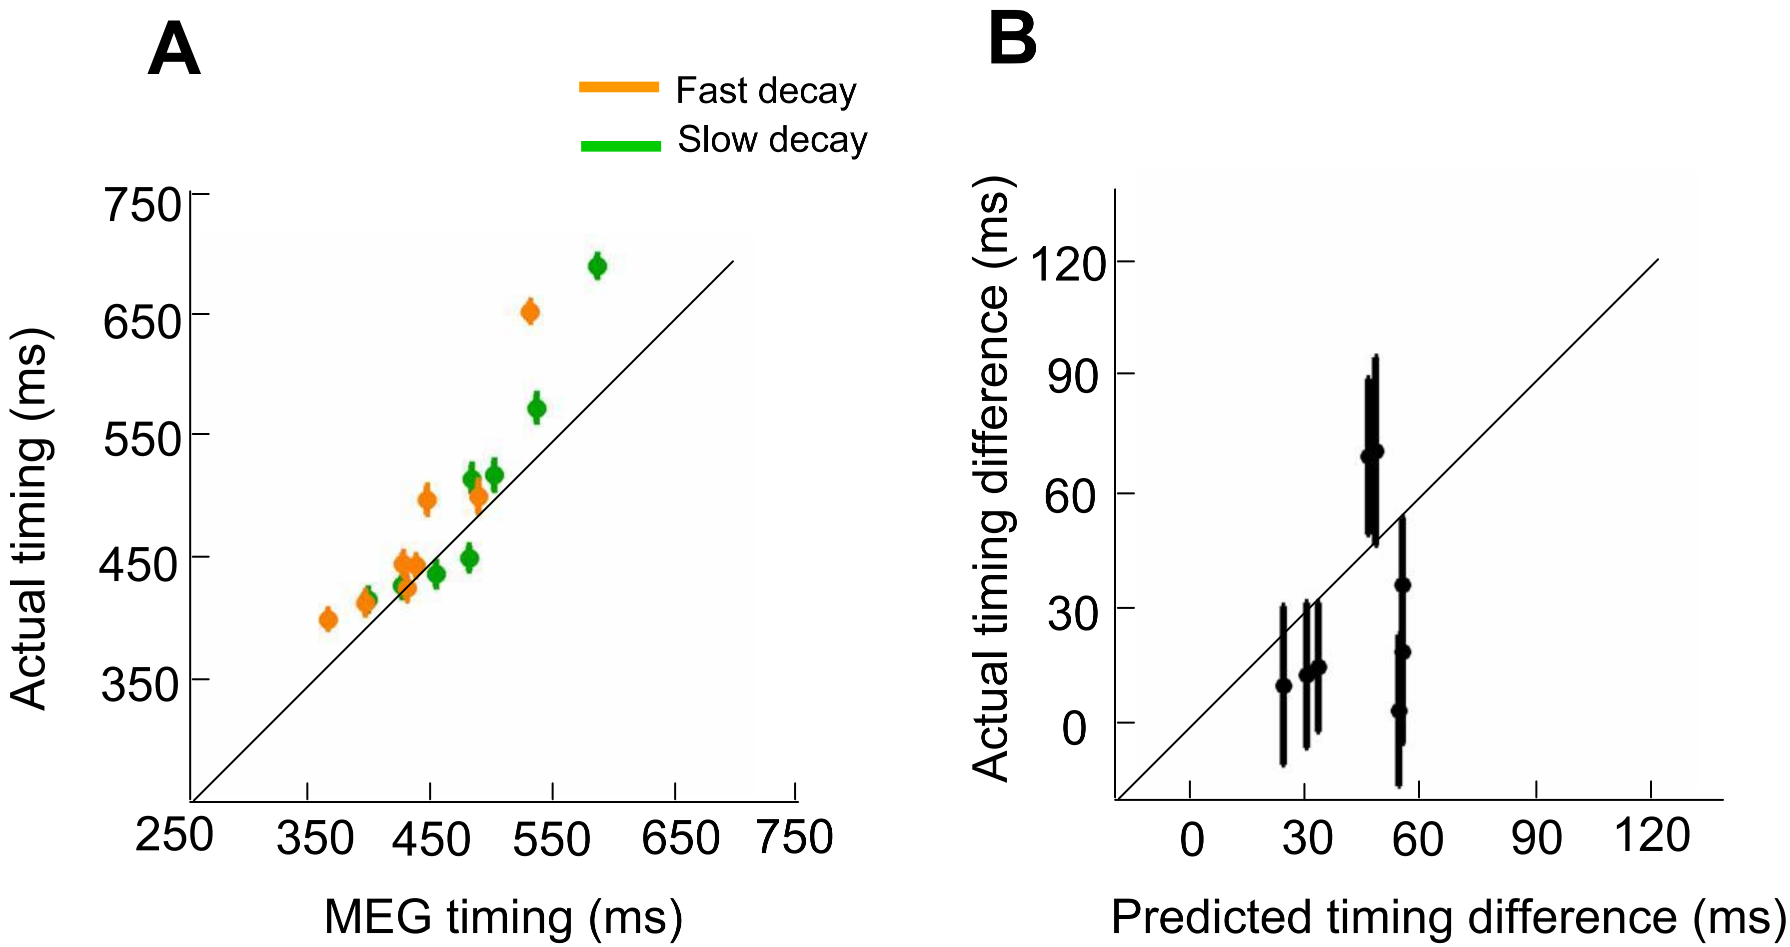

Dean et al: Figure 4

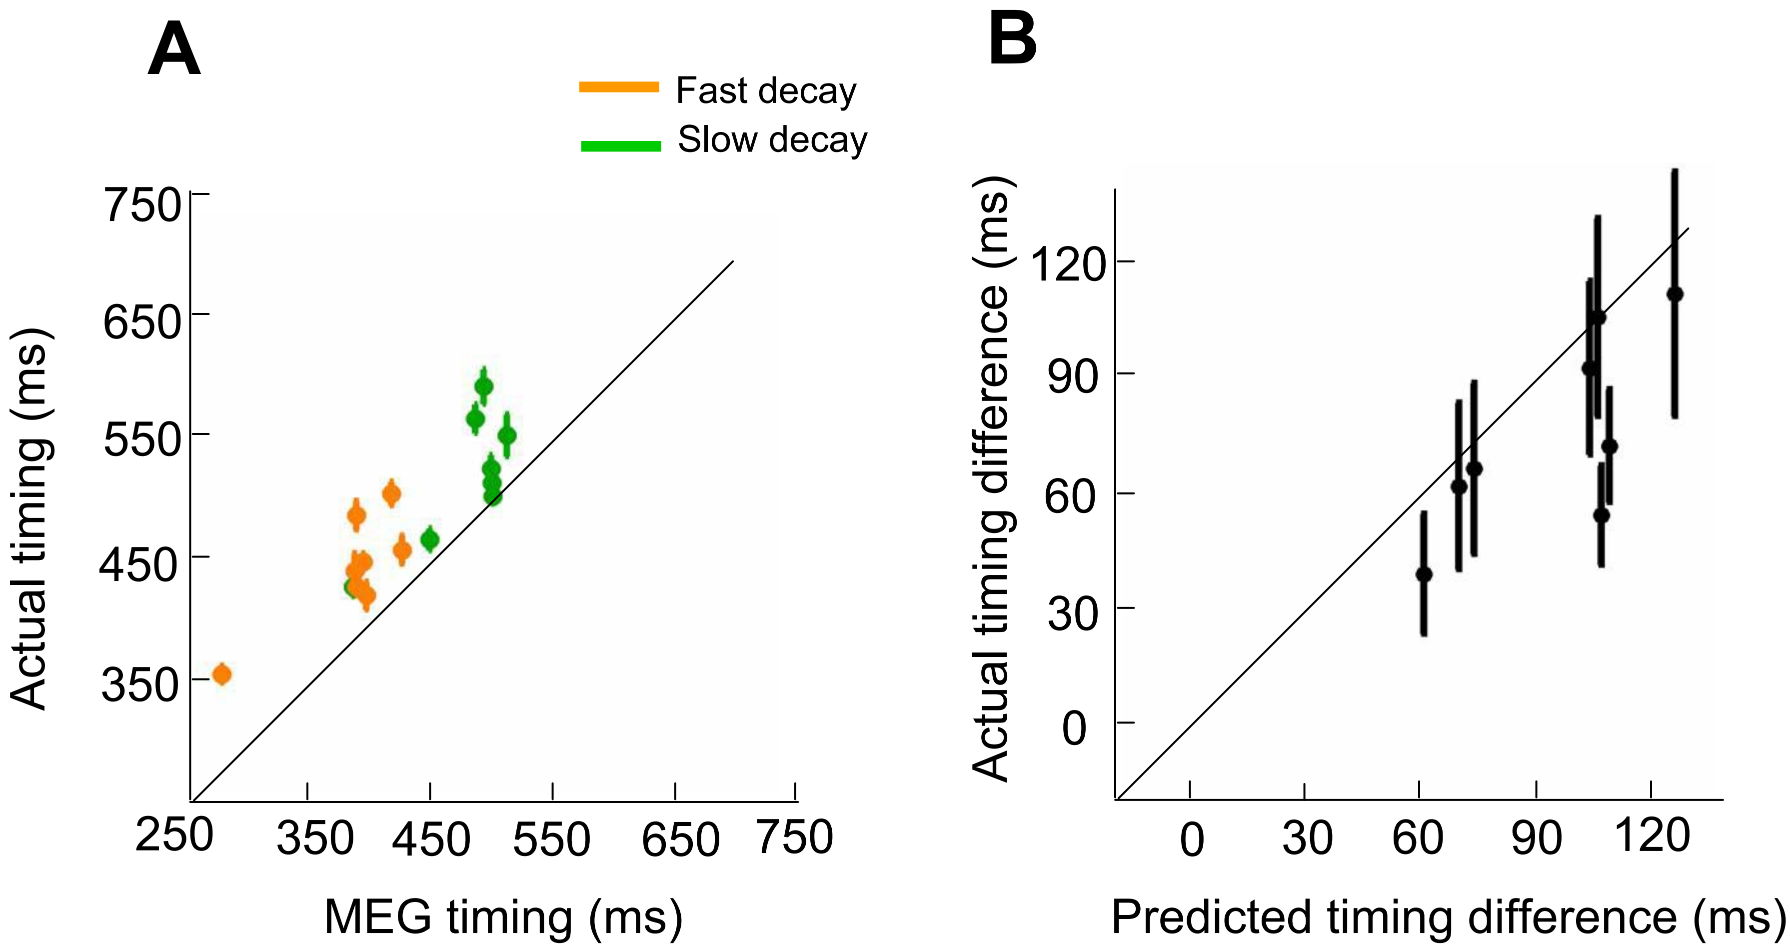

Dean et al: Figure 5

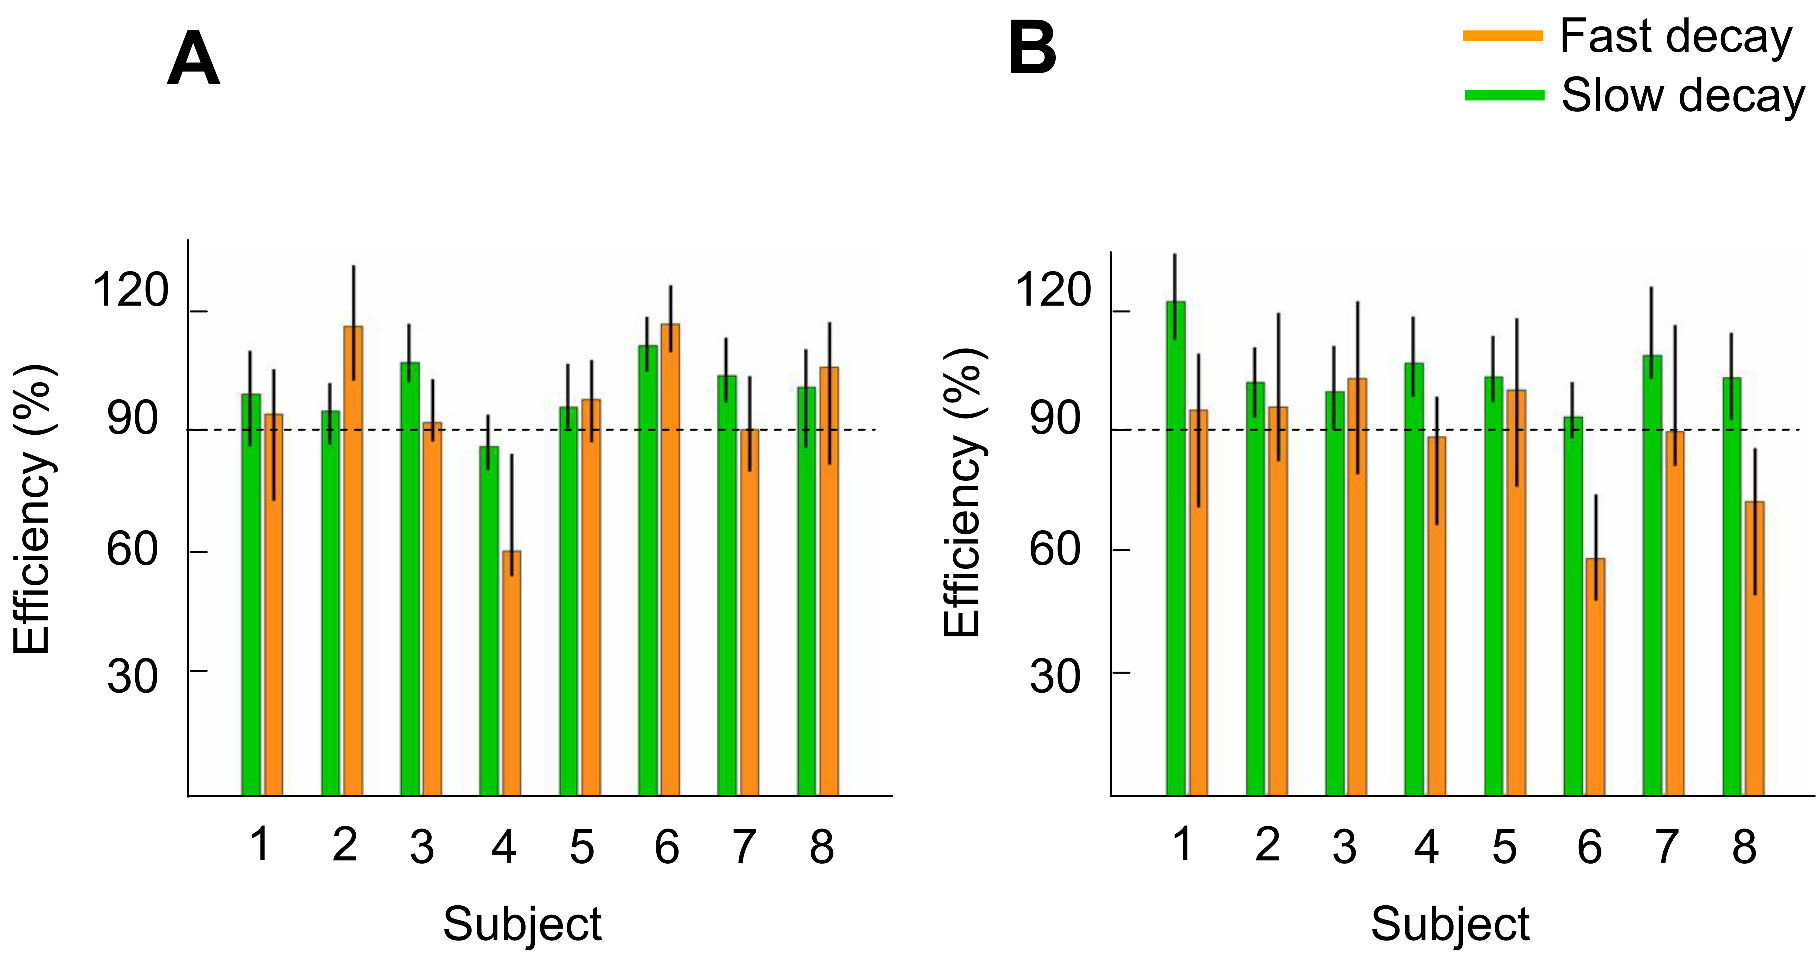

Dean et al: Figure 6

Table 1. Regression of Accuracy on Planned Arrival Time

Subjects	Coefficient	Standard Error	p value	95% Confidence Interval	
1	-0.0384	0.0093	0.02*	-0.0643	-0.0125
2	-0.0317	0.0092	0.03*	-0.0573	-0.0062
3	-0.0026	0.0023	0.31	-0.0089	0.0036
4	-0.0318	0.0058	0.01*	-0.0479	-0.0158
5	-0.0078	0.0040	0.12	-0.0188	0.0032
6	-0.0211	0.0025	0.00*	-0.0280	-0.0141
7	-0.0244	0.0065	0.02*	-0.0426	-0.0063
8	-0.0315	0.0040	0.00*	-0.0426	-0.0204
9	-0.0225	0.0064	0.03*	-0.0404	-0.0046
10	-0.0376	0.0048	0.00*	-0.0511	-0.0242
11	-0.0609	0.0226	0.06	-0.1238	0.0019
12	-0.0201	0.0063	0.03*	-0.0375	-0.0027
13	-0.0197	0.0093	0.10	-0.0455	0.0060
14	-0.0109	0.0054	0.12	-0.0259	0.0042
15	-0.1914	0.0618	0.04*	-0.3629	-0.0198
16	-0.0078	0.0042	0.13	-0.0194	0.0038

* Indicates $p<.05$

Table 2. Regression of Accuracy on Actual Arrival Time

Subjects	Coefficient	Standard Error	p value	95% Confidence Interval	
1	-0.0038	0.0029	0.18	-0.0095	0.0018
2	-0.0137	0.0027	0.00*	-0.0191	-0.0083
3	-0.0034	0.0018	0.06	-0.0069	0.0002
4	-0.0147	0.0028	0.00*	-0.0202	-0.0092
5	-0.0036	0.0029	0.21	-0.0092	0.0021
6	-0.0035	0.0025	0.17	-0.0084	0.0015
7	-0.0074	0.0026	0.01*	-0.0126	-0.0022
8	-0.0039	0.0042	0.36	-0.0121	0.0044
9	-0.0045	0.0028	0.12	-0.0100	0.0011
10	-0.0058	0.0032	0.07	-0.0121	0.0005
11	-0.0004	0.0065	0.95	-0.0132	0.0124
12	-0.0124	0.0037	0.00*	-0.0197	-0.0052
13	-0.0044	0.0033	0.18	-0.0108	0.0021
14	-0.0102	0.0019	0.00*	-0.0140	-0.0065
15	-0.0086	0.0092	0.35	-0.0095	0.0267
16	-0.0038	0.0029	0.18	-0.0095	0.0018

* Indicates p<.05

Experiment 1

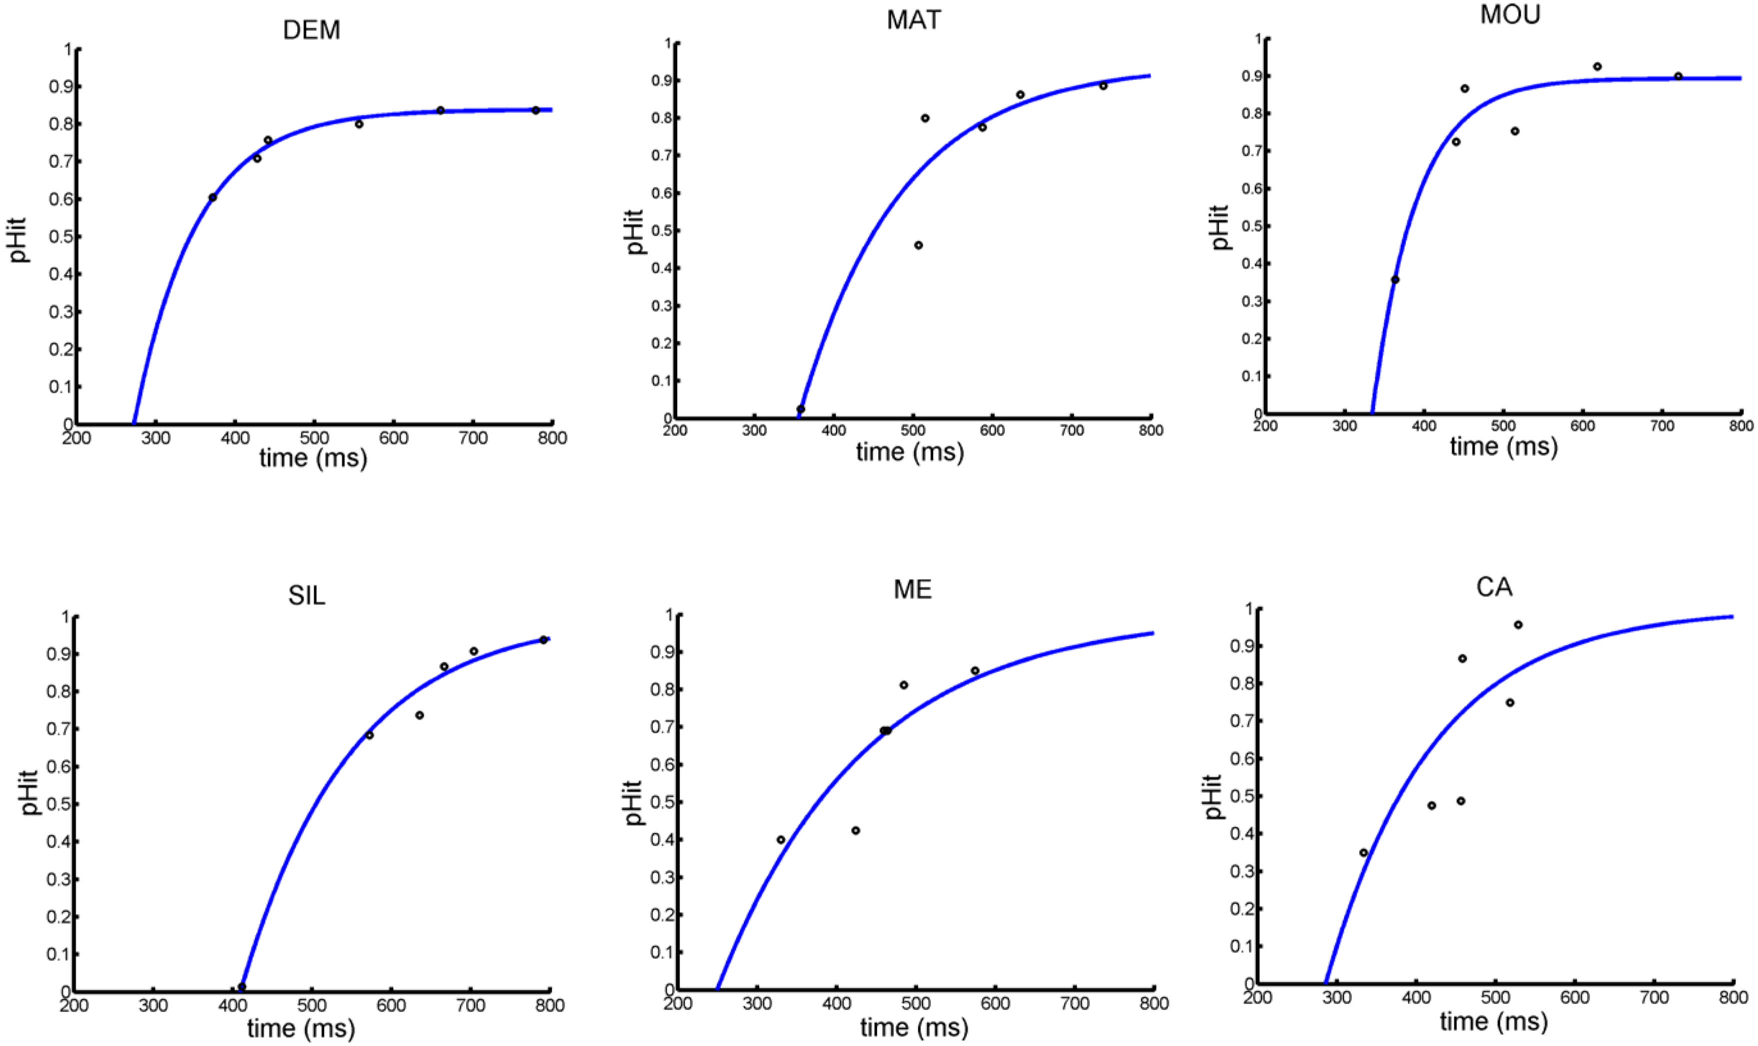

x: arrival time

pHit: probability of hitting the target

Experiment 1

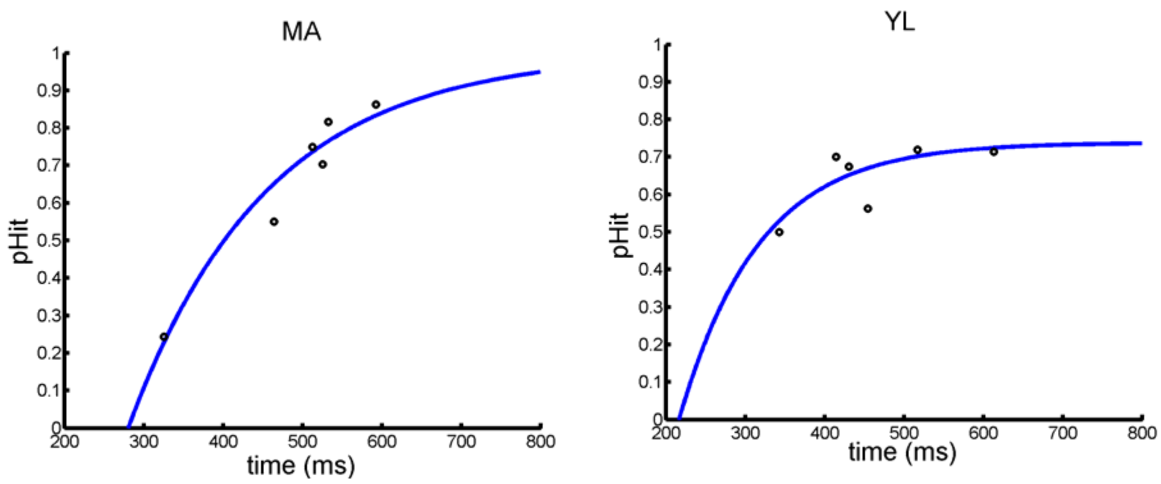

Experiment 2

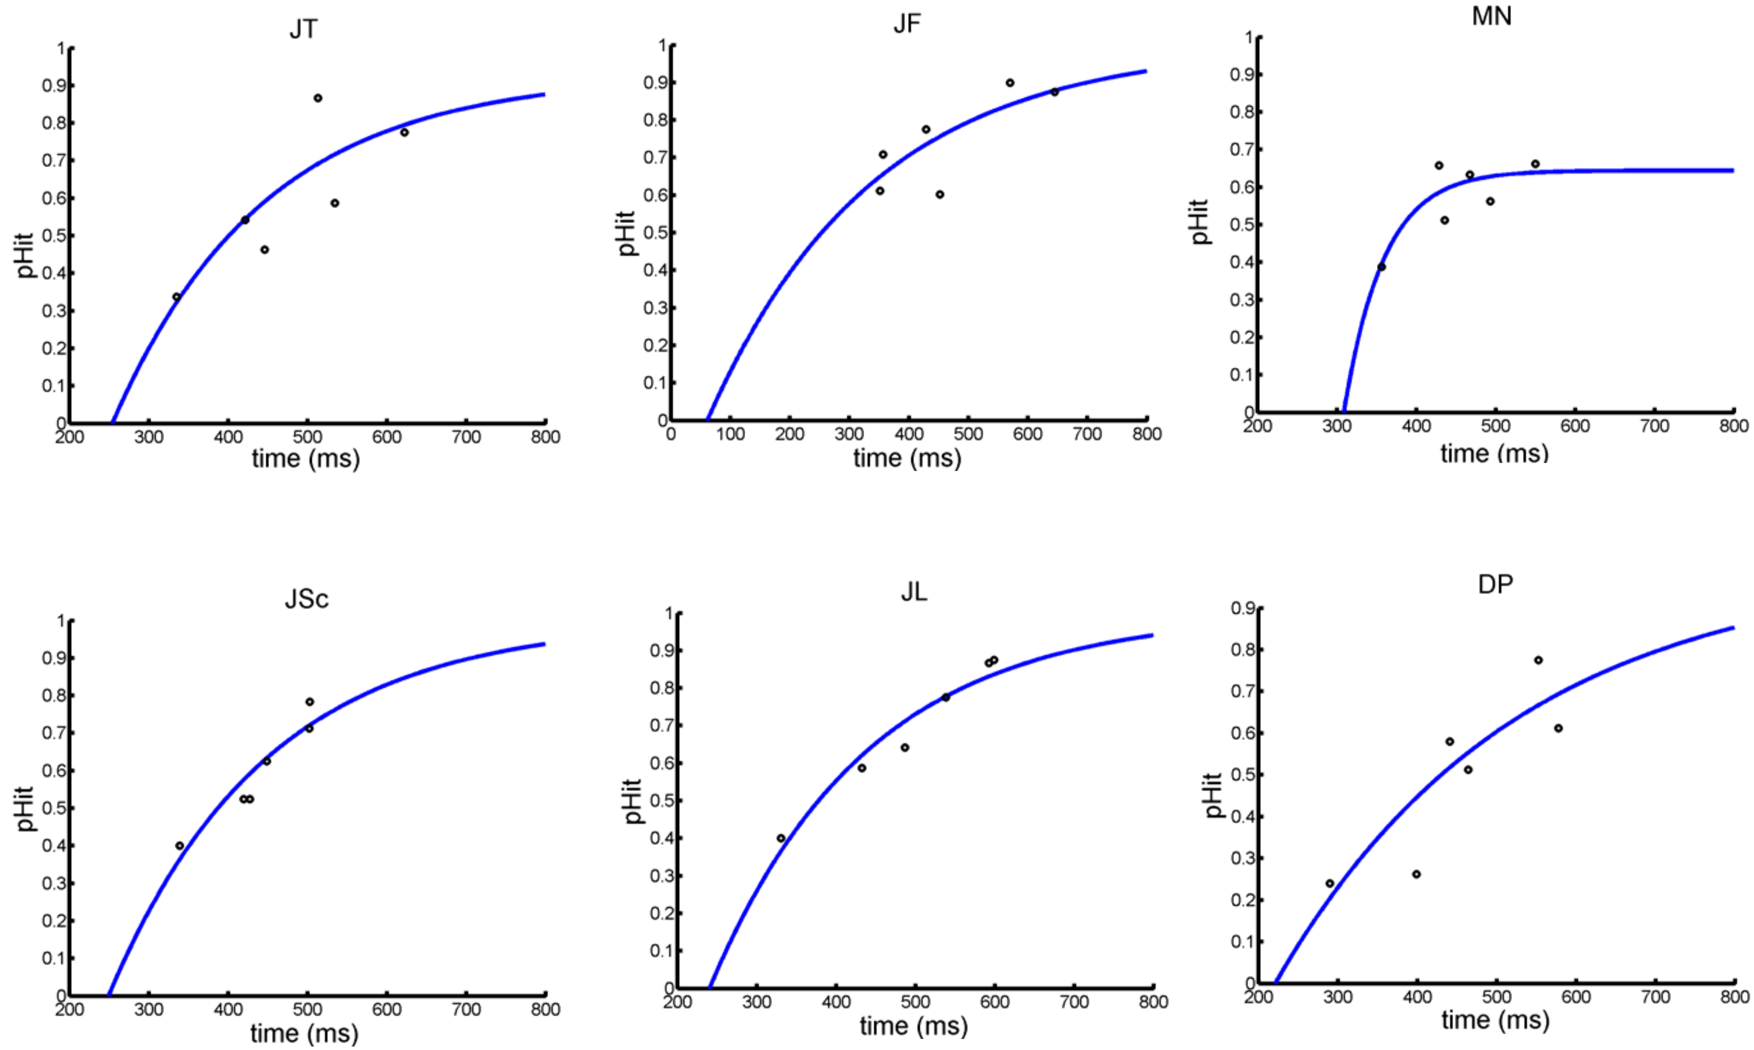

Experiment 2